

# Color and Brightness Appearance Issues for Tiled Displays

Maureen C. Stone

StoneSoup Consulting, Los Altos, CA

## Abstract

Large format displays are often created by combining tiled imaging elements. Ideally, these would give the appearance of a single, seamless display, but in practice, this is difficult to achieve. Even if physical alignment is perfect, differences in color and brightness will make the tiles distinct. Not only are the brightness and colors different, the appearance also varies over the face of the tile and with viewing angle. This paper will describe the primary appearance problems, their causes, and possible solutions.

## Keywords

Display walls, color matching, projection displays

## I. INTRODUCTION

Large format displays are becoming common. They are available commercially, and are being custom-made in various research organizations.[1] In an ideal world, we would have digital wallpaper of unlimited size. For the moment, however, most wall-sized displays are made by tiling multiple, smaller displays, most often projection displays.

Figure 1 shows a back-projected, tiled display where the tiles are physically aligned to sub-pixel resolution. However, the display is far from seamless. The tiles overlap, creating bright lines at the seams. The tiles vary in color and brightness, not only from tile to tile, but within each tile. The goal of this paper is to describe what causes these variations and what can be done about them. Unfortunately, it cannot provide a simple, low-cost way to make this wall look perfect--the problems are too hard for that. But, it will provide a clear description and analysis of the various problems and their solutions. These include removing the bright lines at the seams, color balancing additive color spaces, and reducing the spatial variation within tiles.




Figure 1: A 4x3 projection array, the Stanford Mural, showing the seams and the color variation between tiles. The color variation has been somewhat exaggerated by the photographic process.

## II. Overview

Each rectangular *tile* in figure 1 is the output of a single digital projector. The tiles overlap slightly, causing the bright lines or seams. Each projector has a slightly different *color gamut*, caused by variations in the bulb, the color filters, and the digital processing (contrast, brightness and gamma) for the projector. The spatial variation in brightness, or *vignetting*, has two causes. First, the light from a projection system does not strike the screen with a uniform brightness. Second, the light does not scatter uniformly out the front of the screen, making the perceived brightness depend on viewing angle.

Figure 2 shows a commercial display wall with many of these problems solved. Each tile is a sealed “cube” containing a projector, a screen, and a screen-sized Fresnel lens. The units are stacked to create the display wall. There are physical metal seams between the cubes, just visible in the figure as thin black lines. The projectors in the cubes contain special electronics to aid in color balancing them.




Figure 2: Synalec ([www.synalec.com](http://www.synalec.com)) installation at Columbia energy. The thin, black seams are just visible

There are problems using such a system for research, however. First, they are quite expensive. Second, they are inflexible. Most are optimized for displaying video to viewers standing back several feet, and cannot easily be modified to display custom software. Finally, the surface is not smooth, which limits the type of input devices that can be used. To write on the surface, for example, you would need to put a smooth sheet over the front of the display wall, which would introduce parallax.

## III. Seams

The place where two tiles meet is called a *seam*. The tiles on the display can either abut or overlap. The Synalec display wall shown above has abutting seams with a small gap caused by the physical separation of the cubes. To make abutting tiles truly seamless, the gap between the tiles must be no more than

the inter-pixel gap within the tile. This is difficult to achieve with any mechanical system, and is why seamlessly tiling flat panel displays is so difficult.

Projection systems sharing a common display can abut to the extent that they can be aligned precisely. However, the edge of projected images are often not square, an effect called “pincushion.” Pincushion and other distortions are a common problem in optical systems, and are a property of the lenses used in the system.

Most single-screen projection displays overlap the tiles slightly, which creates bright lines where the light from two projectors falls, and a bright square where four projectors overlap as can be seen in figure 1. A number of methods for removing these have been developed.

The most effective is to mechanically block the light, either by casting soft shadows from masks placed out-of-focus, near the projection lens[2], or by casting sharp shadows from masks placed near the screen. The latter can be seen in figure 3. There is a double layer of screen; the projection screen itself, and a clear acrylic screen mounted about 6” behind it. At this point in space, the images do not overlap, so overlapping pixels can be eliminated simply by placing tape along the seam boundary on each image. This gives the effect of abutting the images without having to rely on the squareness of the edge pixels.




Figure 3: Tape is used to mask the seams for the Stanford mural. The projectors in the foreground are about 3 feet from the taped seams.

Instead of physical masking, some systems minimize the seams by adjusting the brightness in the overlap region. This can be done in software [3], or electronically in the projection system [4]. There are several problems with this approach. The most critical is that it cannot fix the seams at black. Projectors still emit a visible amount of light when off, in the range of 2 to 5 cd/m<sup>2</sup> for the projectors used in the Stanford mural<sup>1</sup>. This is twice as bright on the seams, and four times as bright where four projectors overlap at a corner. Without physical masking, the only way to eliminate these bright areas is to raise the overall system black level to match the brightest point. This color is equivalent to a 25-30% gray, and reduces the contrast by a factor of 4. Also, projectors with imaging elements that reflect light (DLP and DLA) typically have a visible halo around the viewing area as shown in figure 4. We

1. A monitor or laptop black level is less than 0.5 cd/m<sup>2</sup>

believe this is caused by scatter within the projection system, and only physical masking can eliminate it.



Figure 4: Halo caused by scattering in a DLA projector. A similar phenomenon can be found in DLP projectors.

Finally, it turns out to be difficult to feather smoothly. While smooth changes in brightness are not very visible, the eye is good at detecting changes in the derivative. So, if the feathering is not perfect, it can introduce different, visible artifacts at the seams.

#### IV. Color Balancing

Projectors are essentially additive color devices like monitors and liquid crystal displays. White light from the bulb is imaged and filtered to produce the red, green and blue color separations. There are two main forms of imaging elements in commercial projectors: liquid crystal (LCD) and micro-mirrors (DLP) ([www.dlp.com](http://www.dlp.com)). If the projector uses LCDs, the light is split using dichroic mirrors into its red, green and blue components. There are three imaging elements, one for each separation. DLP projectors use tiny mirrors, one for each pixel. Light is directed either out the lens, or into a black cavity. To get grayscale, the light is pulsed. For all but the largest DLP projectors, there is a single imaging element and a color wheel. Therefore, the separations are combined temporally instead of spatially. It is also common in DLP projectors to include a “white” filter on the color wheel, along with the red, green and blue primaries, which makes DLP color more complex. There is a third commercial system, by JVC-Hughes, called DLA ([www.jvc.com](http://www.jvc.com)), which uses a reflecting liquid crystal display. It has three imaging elements, and combines some of the perceptual characteristics of both the LCD and DLP systems.

There are two basic approaches to matching the colors in a tiled display. The first is to make all the critical components match. That is, use exactly the same filters and light sources. In this case, matching two projectors is simply a matter of matching the brightness of the red, green and blue components at all points in space. This is the approach used by most commercial systems.

Precisely matching these components is surprisingly hard to do, however. Even for projectors of all the same type, there is visible variation in the filter color, most noticeably in the red filter. Some groups resort to hand-picking their projectors to get matching filters. There is also variation in bulb color, which varies with age. Again, some groups hand pick their

bulbs and take care to match their ages. There are experiments being performed at the San Diego Supercomputing Center VisLab (vis.sdsc.edu) with using a single light source and sharing it among projectors. However, such a solution will be difficult to scale.

The other approach is to characterize the projector color in a device-independent color space and adjust the input pixels to make the colors match. This is an application of the principles used in color management systems [5][6], and is the approach we are investigating with the Stanford Mural.

### A. Characterizing RGB projectors

Projectors using only red, green and blue filters (LCD and DLA), can be characterized like a monitor [7] except that the black level cannot be assumed to be negligible.

The color gamut is defined by the CIE tristimulus values of the full red, green and blue primary colors, plus the tristimulus value for black. This black value needs to be subtracted from all color measurements to reveal the basic additive system underneath. Figure 5 shows this effect. The purple data are the chromaticity coordinates of the three primaries measured at four levels of brightness, from 0.25 to 1.0. For an additive system, the chromaticity should be independent of brightness, but these data shift towards the black point. The orange data are the chromaticity coordinates computed from the same measurements after subtracting the tristimulus value for black. These are nearly constant.

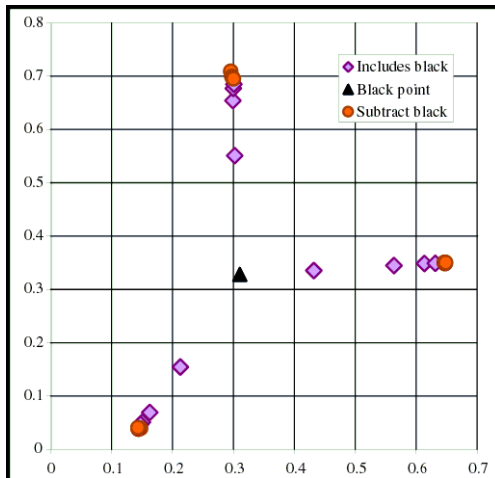


Figure 5: Tristimulus values of the red, green and blue primaries measured at 4 brightness levels, plus black. Subtracting black produces a constant chromaticity for each primary.

The tristimulus values of the primaries defines a matrix that, when combined with the black offset, can be used to transform from device color intensity values to tristimulus values and back again.

There are two sets of measurements that need to be taken to create this transformation. The first defines the relationship between input pixel values and normalized intensity values for each primary. This is the intensity transfer function, or *ITF*, and must be measured for each primary independently. For commercial projectors optimized for displaying video (i.e.

most of them), this will be a gamma curve, such as the ones shown in figure 6.

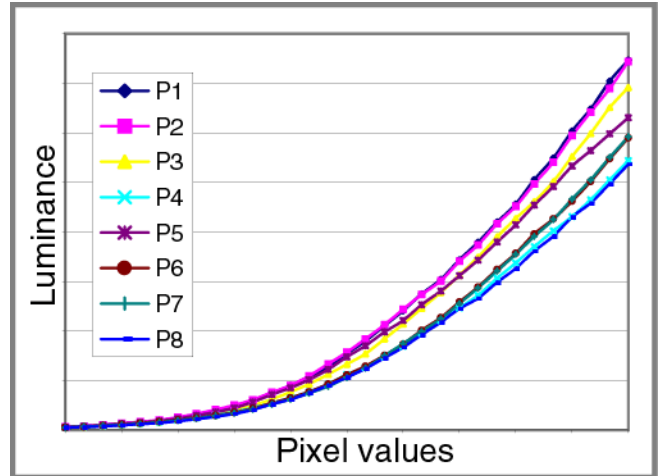


Figure 6: ITFs for eight NEC MT1030 LCD projectors. These curves can be changed by adjusting brightness (raises whole curve) and contrast (raises the maximum value) using the projector menus.

More formally, if  $p$  is an input pixel value, there is a function  $ITF(p)$  that maps  $p$  to normalized intensity. Therefore,

$$\begin{aligned} R &= ITF_R(p_R) \\ G &= ITF_G(p_G) \\ B &= ITF_B(p_B) \end{aligned} \quad (1)$$

The *ITF* is defined by the “gamma” settings in the graphics card, in the operating system, and in any application used to display the colors, as well as by the contrast and brightness settings in the projector. Therefore, it very important to be sure that these are constant during measurement and don’t change once the system has been characterized

The second set of measurements gets the tristimulus values for black, and for the maximum brightness of the red, green, and blue primaries. Let  $c$  be the color intensity vector computed from the *ITFs*,  $[X_R, Y_R, Z_R]$ ,  $[X_G, Y_G, Z_G]$ , and  $[X_B, Y_B, Z_B]$  be the tristimulus values for the primaries, and  $t_K = [X_K, Y_K, Z_K]$  be the tristimulus values for black. Then the tristimulus values  $t$  corresponding to  $c$  can be computed from:

$$\begin{aligned} cM + t_K &= t \\ M &= \begin{bmatrix} X_R - X_K & Y_R - Y_K & Z_R - Z_K \\ X_G - X_K & Y_G - Y_K & Z_G - Z_K \\ X_B - X_K & Y_B - Y_K & Z_B - Z_K \end{bmatrix} \end{aligned} \quad (2)$$

To convert from CIEXYZ to RGB, invert the matrix and rearrange, giving:

$$c = (t - t_K)M^{-1} \quad (3)$$

These transformations can also be defined as a single 4x4 homogeneous transformation matrix common to most graphics systems.

## B. Characterizing DLP projectors

DLP projectors, like the Compaq MP 1800 ([www.compaq.com](http://www.compaq.com)) used in the Stanford Mural, are not simple RGB systems because they include a “white” filter on the color wheel. This makes the white substantially brighter than the sum of the RGB primaries, about 145-150% for the ones we have measured. This can be seen in Figure 7, which is a plot of the spectral distribution for the red, green, blue, black and white light. The dashed lines show the sum of the R, G, B spectra, plus the difference between this sum and white. The difference curve reveals an additional source of light that is spectrally similar to the summed white. We have confirmation from a Compaq engineer that this is indeed what is going on.

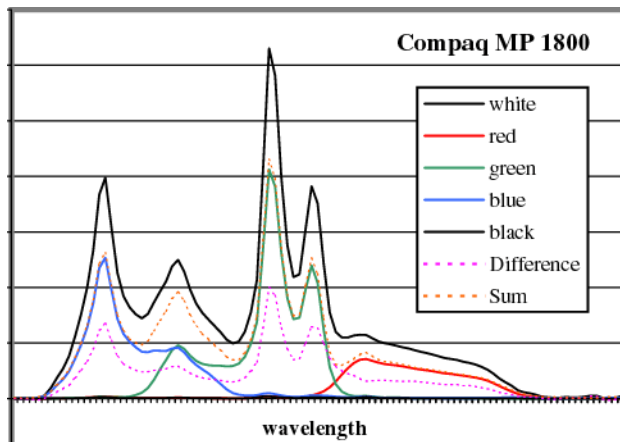


Figure 7: Spectral distribution curves for the Compaq MP 1800

Selectively adding an additional white light to the system simply adds another row to the matrix in equation 2 and another value to the vector  $c$  that indicates what percentage of white is added. The unknown at this point is when the white filter is included. [note to reviewers--if this paper is accepted, I would expect to have a better description by the time the final version is written]

## C. Conversion to a standard gamut

Given the characterization for each projector, the next step is to create a common gamut that is the target for all of them. If this gamut is contained in all the projector gamuts, then the conversion will introduce no clipping. Call this common gamut the *standard gamut*. To make the tiles all look the same, color, we want to modify the RGB input colors so that full red, for example, becomes the full red of the standard gamut instead of the full red of the device.

If  $M_d$  is the transformation from a projector’s color intensity values to tristimulus space represented as a  $4 \times 4$  matrix, and  $M_s$  is the equivalent matrix for the standard gamut, then  $M_s M_d^{-1}$  describes the transformation we want. The full algorithm is:

- 1 Take the RGB pixel values and find the corresponding intensity values. That is, pass them through the *ITFs* in equation 1.
- 2 Use  $M_s M_d^{-1}$  to transform them into the standard gamut and back into color intensity values.

- 3 Map them through the inverse of the *ITFs* to convert them back to pixel values.

It may be tempting to make the *ITFs* linear to avoid the mapping step, especially since many computer graphics students have been taught to “gamma correct” to a linear intensity ramp for rendered graphics. However, a linear *ITF* is perceptually quite non-uniform, and conflicts with common practice in most imaging applications. A 2.2 gamma curve is perceptually a much better use of an 8-bit brightness mapping, and is what is expected for video and scanned imagery. Windows desktops run at a gamma of 2.2. Not even SGI workstations default to a linear ramp (they have a default of around 1.3). Only consider a linear *ITF* if all applications are going to be high-end rendered imagery designed for it.

## D. Finding the standard gamut

The gamut of an additive device based on three primaries is a geometric shape called a *rectangular parallel piped*, which is simply a cube where the faces are rectangular parallelograms.<sup>1</sup> The standard gamut, therefore, should also be such a shape, and should fit completely inside of all the other gamuts. Finding such a shape could be treated strictly as a computational geometry problem. The following process, however, incorporates some colorimetric principles, and will generate a usable, if possibly not optimal, gamut.

- 1 Compute the chromaticity coordinates of the primary colors. Plotted on a chromaticity diagram, they form a triangle around the black/white point, as can be seen in figure 5. Pick an R, G and B chromaticity whose triangle lies completely inside of all the points.
- 2 Select a black point by averaging the chromaticities for black. Set the luminance (Y) for black at the maximum of all the projectors.
- 3 Select a white point by selecting a Y for each of red, green and blue. Again, using an average for each primary will be a good starting point if the projectors are similar. Or, select a white point that matches the black point chromaticity.
- 4 Compute the matrix  $M_s M_d^{-1}$  for each projector. If the bottom row contains negative numbers, the corresponding black is out-of-gamut, so adjust the standard black until none are significantly negative. Raising the black Y slightly is often sufficient.
- 5 Sum the columns of the matrix. If any sum to greater than 1.0, there are bright colors out of gamut. Scale the Y of the primaries of the standard gamut to correct it. A uniform scale maintains the white point.

To date, we have simply used Excel and done this by hand, but it could be written into a program fairly easily.

---

1. This will not be true for a DLP projector, but may be similar shape.



## E. Measurement technique

Color characterization is only as good as the data that drives it. There are a few tricks to measuring projection displays worth mentioning here.

For a front projection system, the only choice is an instrument with a lens like the Photoresearch PR-650 ([www.photoresearch.com](http://www.photoresearch.com)). This is a spectroradiometer, where the incoming light is split into spectral bands and measured. Colorimetric information is then calculated from the spectral distribution. Instruments that measure spectra give the most accurate color measurements. However, an instrument with a lens will “see” all the ambient light and be sensitive to viewing angle. And, these are expensive instruments. Typical prices for a spectroradiometer range from \$12,000 to \$40,000, and they only operate accurately if given regular maintenance.

For a rear-projection system, it is possible to use an instrument with an aperture that is placed directly onto the screen. Most of these instruments are colorimeters, though there are also a few spectroradiometers that work this way, such as the X-Rite ColorTron ([www.x-rite.com](http://www.x-rite.com)). Colorimeters are relatively inexpensive instruments (under \$1,000) that use specially designed filters to represent the color matching functions. However, care must be taken to be sure the colorimetric information is accurate enough for the application. We have tried using an X-Rite DTP92, which is designed for measuring monitors. We have had two problems using this colorimeter on the current Mural. First, the projectors are so bright it is easy to exceed the maximum luminance rating for the device. Applying neutral density filters can correct this problem. More seriously, monitor colorimeters are optimized for the sort of spectra expected from a CRT display. The spectra coming from a projection display are quite different. We have seen errors as large as 50% in the X and 100% in the Z measurements of the colorimeter with respect to the spectroradiometer. The normalized luminance measurements, however, are accurate within 5% over most of the range.

For most of the measurements needed to characterize the projectors, relative measurements are sufficient. That is, first measure the tristimulus values for black, then subtract that value from all subsequent measurements. This eliminates any contributions from extraneous light as long as the measurement conditions are kept constant.

The black measurement in the characterization, however, is an absolute measurement. It is an interesting question what is “correct” for the black in this characterization. The most stable measurement is to measure only the light leaking from the projector when it is displaying “black.” This is achieved by eliminating all ambient light, included that generated by adjacent projectors, before measuring. However, it might be better to measure the black for normal viewing conditions, possibly even including typical room lighting. Further work is needed to understand this aspect of projector characterization.

## F. System integration and performance

Adding two table lookups and a matrix multiply per pixel in software has a noticeable impact on rendering speed. How-

ever, this type of operation could easily be implemented in hardware on the graphics card. Given the trend in graphics hardware like the NVidia GeForce 3 ([www.nvidia.com](http://www.nvidia.com)), this seems the most effective and general solution.

For systems without hardware acceleration, the following approaches can be used to improve the performance.

- Use a cache to reuse the corrected values. For many applications, this will have a substantial impact.
- Ignore the color variation and correct only the relative brightness of the primaries. This reduces to three table lookups, though it care must be taken to avoid quantization.
- Correct higher in the graphics pipeline. OpenGL and similar systems do not compute the lighting and shading equations for each pixel. Instead, the vertices of small triangles are computed and the intervening pixels are interpolated. Similarly, the color at the vertices of the triangles could be corrected and the rest interpolated. Unlit textures must be corrected as well. This introduces complexity, but potentially corrects fewer pixels.

## V. Spatial variation

Figure 8 shows the Stanford Mural with the seams taped and a background color that minimizes the color variation. However, the tiles are still clearly visible due to the variation in brightness, sometimes called *vignetting*. This section will discuss its causes. Most of the topics in this section are described in detail in the book “Projection Displays” [8].

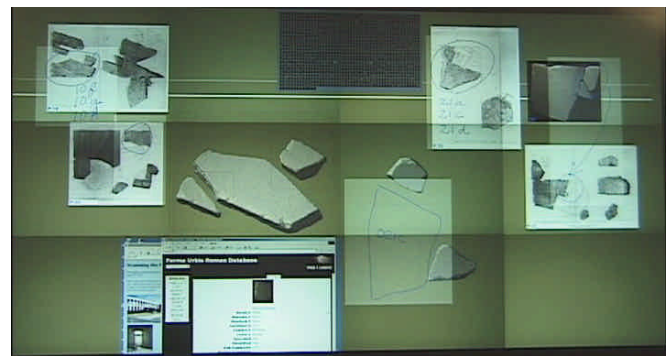


Figure 8: Stanford mural with the bright seams eliminated by masking them as in figure 3. The vignetting is clearly visible.

### A. Non-uniform illumination

There are two causes for the vignetting. One is that the pattern of light produced by the projector is non-uniform. The brightest spot is near the bottom edge of the tile, and the upper corners are the darkest. The bright spot is not in the middle because these projectors contain *keystone correction*. When the projector is set on a table, the bottom edge of the image will remain at roughly a constant height as the image expands. This is achieved by placing the imaging element off-center with respect to the lens.

The luminance of the dark corners of the tiles in the figure is about 45% of the luminance of the brightest point. In com-

parison, a monitor or laptop screen drops to about 70% of maximum in the corners. For a simple projection system, the brightness away from the center of projection obeys the  $\cos^4$  law,  $Y = \cos^4(\alpha)$ , where  $\alpha$  is the angular displacement. For our projectors, however, this will account for only about a 13% fall-off ( $\alpha=15$  degrees). Another possible cause is variation in illumination of the imaging element inside of the projector.

### B. View-dependent variation

The other cause of vignetting is illustrated in figure 9(a). Each projector generates a cone of light. While the projection screen provides some scattering, the light is still strongly oriented. A viewer at point A sees less of the light at the edges than a viewer at point B, who sees less of the light in the middle. The light is brighter in the middle than the edges. Therefore, the brightness varies with the position of the viewer.

To correct for this, a Fresnel lens is typically placed directly behind the screen to redirect the light, as shown in figure 9(b). Commercial display walls include such a lens in each cube. It is much more difficult to include on a single screen, as an array of perfectly abutting lenses is needed. If the seams are feathered, light falling on an adjacent lens will create artifacts. For the abutting arrangement in the Stanford Mural, however, we might be able to get very good results from such an array.




Figure 9: Effect of adding a fresnel lens to remove viewing angle dependence.

Figure 10 shows the effect of adding a Fresnel lens. The figure shows is a white, rectangular tile viewed at an oblique angle to the screen. Without the Fresnel lens, it is visibly brighter closer to the viewer. Moving to the other side would maintain this relationship--the near side is always brightest. With the Fresnel lens, the pattern is much more uniform. Note, however, that the Fresnel lens does not change the brightness fall-off caused by the variation in illumination.




Figure 10: The effect of adding a Fresnel lens. (a) without a lens. (b) with a lens. This is a different projector than those used in the Mural

### C. Screen gain

Most projection display screens are designed for audiences sitting or standing well back from the screen. Therefore, care is taken to avoid “wasting” light by directing it away from the sides. *Screen gain* is the measure of how much the screen differs from a pure lambertian surface with respect to scattering the light. A gain of 1.0 means uniform scatter, greater than 1.0 directs more of the light forward. Most commercial systems use high-gain screens. If, however, multiple projectors share a screen, it is important to have a gain of 1.0. Otherwise, the screen gain will amplify the viewing dependent brightness variation at the seams. Front projection screens with unit gain are not difficult to find, but such back-projection screens are rare. The Jenmar ([www.jenmarvs.com](http://www.jenmarvs.com)) screen we use on the Mural has a gain near 1.0, but only when used with a Fresnel lens. A simple translucent material like drafting vellum will have a gain of 1.0, but will be low-contrast and can introduce blur if the material is thick. This is sometimes called a “white screen” in contrast to the Jenmar BlackScreen.

### D. Other spatial variations

LCD imaging elements polarize the light passing through them. If there are any other polarizing elements in the system, this polarization will create visible artifacts. Figure 11 shows a white “X” on a supposedly uniform gray background. This is a picture of a DLA projector, but it is also typical of what we have seen with LCD projectors. We believe that this is caused by a slightly polarized filter in the system, probably the infrared filter that keeps heat away from the lens. Stretching such a filter could polarize it. Placing a polarizer in front of the screen and rotating it changes the coloring, verifying that this is a polarization effect. The picture exaggerates the colorization somewhat, and it is most noticeable at the small tile sizes we use for the Mural (21” diagonal as opposed to a more common 30-40” diagonal).

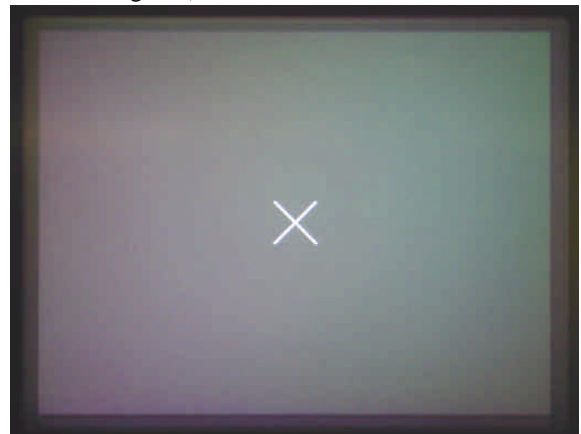


Figure 11: Polarized color mottle caused by LCD imaging systems.

DLP projectors do not produce polarized light. However, because of the color wheel, it is possible to see RGB “rainbows” along edges if you move your head too quickly. That is,

each separation is imaged on a slightly different part of the retina due to the head motion. Based on an informal study, we found that a few people find this extremely annoying, but most do not find it significant.

*Speckle* can be seen on high-contrast back projection displays, and looks like tiny sparkling dots in the image. Speckle is a phenomenon associated with coherent light, and is most commonly seen on lasers. However, the light from a small aperture lens like that in a small digital projector is also somewhat collimated. The Jenmar screen we use has sufficient contrast to show speckle. Reducing the contrast will eliminate the speckle, but degrade the overall image appearance.

## VI. Summary and Conclusions.

In summary, it is very difficult to build a tiled display that has perfectly uniform appearance. However, this should not discourage research groups from building them and experimenting with them. There is something inherently exciting about pixels “on the wall.”

The Stanford Mural, in spite of its appearance flaws, creates excitement whenever we demonstrate it. Not only is it visually commanding, but it enables novel and high-quality systems and applications research. This is the point of building experimental display walls.

Generally speaking, front projection systems are easier to make “seamless” than back projection systems because there is little view-dependent vignetting. It is also easier to have a bright, high-contrast screen without encountering speckle because the light is more completely scattered. Similarly, polarization-induced color shifts should be less of a problem.

Color balancing is a theoretically straight-forward process for additive systems. In practice, care must be taken to correctly account for the high black level of projection displays, and to ensure that the measured values are accurately taken. To color balance with at high rendering speeds, however, requires hardware assistance. Practically speaking, balancing the ITFs so that all tiles have the same minimum and maximum brightness makes a big improvement in appearance. This can be done at a variety of places in the system, or sometimes even in the projector itself.

For rear-projection systems, the problem of vignetting, especially view-dependent vignetting, is the most significant appearance feature after spatially aligning the tiles. While application of an alpha-mask could be used to compensate for non-uniform illumination, only a Fresnel lens will remove the view-dependent vignetting. Commercial systems generally sacrifice complete seamlessness for the visual uniformity created with the Fresnel. The thin black line at the seam may also help minimize the visual impact of small, tile-to-tile, color variations, as it eliminates the sharp edge between them.

In conclusion, the goal of this paper is to describe those features of tiled displays that affect their color and brightness appearance. While there is no single, simple solution to any of these problems, understanding them avoids wasted effort and leads to better designs.

## VII. Acknowledgements

Thanks to Pat Hanrahan and the Interactive Workspaces project at Stanford for giving me the opportunity to work on this problem. Thanks to Brian Wandell for letting me borrow his PR-650. Special thanks to François Guimbretière, my collaborator in building and evaluating the Mural. This work was partially funded by the Department of Energy

## VIII. References

- [1] K. Li and T. Funkhouser, eds., “Large Displays,” *IEEE Computer Graphics and Applications*, Vol 20, No. 4, July/Aug 2000.
- [2] K. Li and Y. Chen, “Optical Blending for Multi-Projector Display Wall System,” *Proc. 12th Lasers and Electro-Optics Society 1999 Ann. Mtg*, IEEE Press, Piscataway, NJ, Nov. 1999.
- [3] R. Rasker et al., “Multi-Projector Displays Using Camera-Based Registration,” *Proc. IEEE Visualization 99*, IEEE Press, Piscataway, NJ, Oct 1999, pp. 161-168
- [4] P. Inova, Seamless Video Display, US Patent #4,974,073, filed 14 Jan. 1988, issued 27 Nov. 1990.
- [5] E. Giorgianni and T. Madden, *Digital Color Management*, Addison-Wesley, Reading, MA, 1998.
- [6] M. Stone, W. Cowan, J. Beatty. “Color Gamut Mapping and the Printing of Digital Color Images,” *ACM Transactions on Graphics*, 7(3), October 1988, pp. 249-292
- [7] W. Cowan, “An Inexpensive Scheme for Calibration of a Colour Monitor in Terms of CIE Standard Coordinates,” *ACM Computer Graphics*, Vol 17. No. 3, pp. 315-322.
- [8] E. Stupp and M. Brennessholtz, *Projection Displays*, John Wiley & Sons, New York, NY, 1999.