Color Balancing Experimental Projection Displays

Maureen C. Stone
StoneSoup Consulting, Los Altos, CA

I. Abstract Heading

Experimental, tiled displays made of commodity projectors provide a relatively easy and cost effective way to explore "on the wall" viewing and interaction. To color balance the display, each projector must be characterized and mapped to a common gamut. Projectors with three imaging elements and three filters can be characterized by a simple extension to the monitor calibration model. However, projectors with a single micro-mirror array and a color wheel may include a "white printer" to increase the system luminance. This makes the characterization more complex.

Introduction

Experimental, tiled displays made of commodity projectors are becoming common in computer graphics research.[9] They provide a relatively easy and cost effective way to explore "on the wall" viewing and interaction. Displays can be either front or back projected. Each "tile" is a single, projected image. There are many challenges to making such a display appear as a single, seamless surface.[12] This paper will focus on the color characteristics of such systems, and of projection displays in general.

Figure 1 shows such a display, called the Stanford Interactive Mural, that has been constructed by tiling twelve Compaq MP1800 projectors. These are small, lightweight projectors based on the Digital Light Processing (DLP) imaging technology from Texas Instruments.[6] The mural is used for research in parallel computing, visualization and interaction. It is far from seamless. Most of the variation is caused by the high-contrast screen, which does not diffuse the light uniformly in all directions. This effect, which accounts for most of the appearance of dark edges around each tile, is viewing angle dependent. There are screens with more uniform distribution, but they have less contrast and tend to blur the image, at least at the pixel densities we use on the mural (64 pixels per inch, which is 2 to 3 times that of a commercial display wall).

Figure 1: A 4x3 projection array, the Stanford Interactive Mural. Displayed are images, sketches, 3D models and a virtual desktop. The user interacts with the content using an eBeam pen

Projectors are essentially additive devices, and those that use three imaging elements, one per separation, can be characterized using an extension of the CRT characterization model.[1,2,3] However, small DLP projectors use a single imaging element constructed from an array of micro-mirrors (DMD) and a color wheel. Along with the expected red, green and blue filters, there is a clear segment that is used to increase the maximum luminance of the system, much as a black printer is used to increase the density of a print.

This paper will first describe the general process we propose for color balancing a tiled display. Then, it will discuss the problem of applying the monitor characterization model to projectors. Key to this approach is correcting for the high black level (light displayed at pixel value (0,0,0)) that is common to all digital projectors. It then will describe the effect of the white segment on the DLP projector gamut and its characterization. Finally, it will discuss overall strategies and issues for characterizing tiled projection displays.

Process Overview

The process we propose is similar to any device-independent color management problem[5]. First, create an invertible characterization for each projector that maps from input RGB pixel values to a perceptually based space such as tristimulus values. Then, define the standard gamut for the tiled display. Ideally, this would be contained within all the projector gamuts to avoid gamut mapping. For a homogeneous array of projectors, such a constraint should not be too limiting. Finally, compute the transformation that maps input RGB pixels for a specific projector to the standard gamut. That is, 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 square matrix, and M_s is the equivalent matrix for the standard gamut, then $M_s M_d^{-1}$ describes the transformation needed. This assumes, of course, that pixel values are correctly transformed to linear intensity. Even if the characterization is not a matrix, the form of the solution will be similar. We need a forward transformation to the standard space, and an inverse transformation to the specific device space.

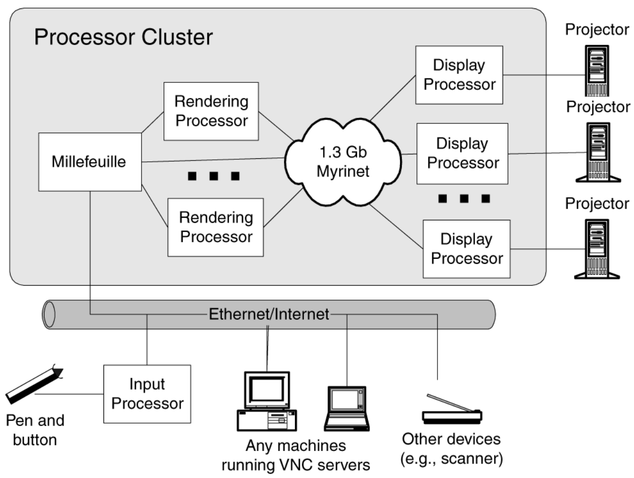

Figure 2: System architecture for the Interactive Mural. Color balancing would best be implemented in the display processors

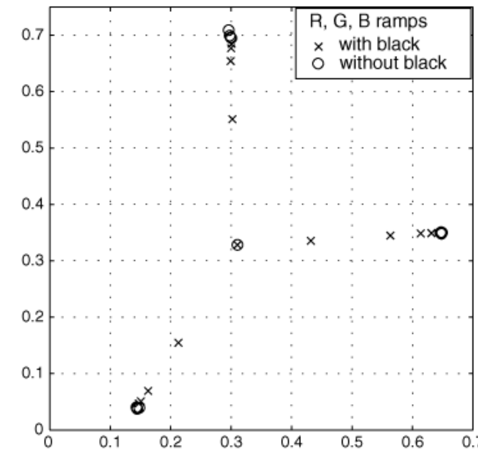

Figure 3: Comparison of RGB ramp data with and without subtracting the black value from the color measurements. Subtracting black produces a nearly constant color for each primary.

For an interactive display wall, the color balancing must be performed in real time without degrading the system performance. Figure 2 shows the system architecture for the processor cluster that drives the Interactive Mural. Millefeuille acts as the window manager and handles input to the system. Rendering is done on a 32 processor graphics cluster.[7] These processors are connected by a high-speed network, the Myrnet. Twelve of these processors contain graphics cards and are connected to the projectors using a digital video interface (www.ddwg.org).

At the level of Millefeuille and the higher levels of rendering in the cluster, the Mural is treated as a single, large display surface (roughly 4000 x 1500 pixels). Logically, color balancing should occur in the display processors, as this is the point in the system where the image is split into individual projected tiles. Modern high-performance graphics cards, such as the NVidia G-Force series, should have sufficient power and flexibility to implement the color transformations.

Applying the CRT model to projectors

The most common form of digital projector contain three liquid crystal (LCD) imaging elements and a dichroic mirror to create the red, green and blue separations. How well these projectors adhere to the CRT characterization model depends on the quality of the LCD panels in the projector, the projector optics, plus characteristics of the output screen.

Brainard[2] reports that the CRT model has been used successfully to characterize LCD projectors. However, some LCD panels, both in flat panel displays and projectors, do not exhibit channel constancy, at least, not to the levels required for vision science.

All digital projectors (in contrast to CRT-based projectors) emit a significant amount of light when displaying "black" pixels. As a result, the black point must be included in the characterization. Figure 3 demonstrates this problem. The raw measurements, shown as small crosses, shift dramatically towards the black point, emphasizing the need to subtract the black from the measurements for these devices. The open circles are the chromaticity coordinates computed from the same measurements after subtracting the tristimulus value for black (X: 0.657, Y: 0.695, Z: 0.765). These are nearly constant, as is required for applying the CRT model.

An earlier version of the mural used eight NEC MT1030 LCD projectors. At the time, we found that independently balancing the red, green and blue primaries significantly improved the appearance of the display wall, suggesting an additive model could be applied. Recently, we have taken some more detailed measurements of two of these projectors to evaluate them with respect to the CRT characterization model. Figure 3 is data from one of these projectors.

The projectors were displayed, side-by-side on a rear projection screen. The measurements were taken with an X-Rite DTP92 colorimeter placed on the surface of the screen, which minimizes viewing angle effects. For each projector, we measured 31 levels of brightness, in equal pixel steps, for each primary and for the grayscale. Each measured color filled most of the displayed area, as scattering from other displayed colors is a known problem with this projector/screen combination.

The coefficient of variation (standard deviation/mean) for each primary is shown in table 1. These value can be read as percentages, so the maximum error shown is 6.6%.

Table 1: Coefficient of variation of the primary colors

	Left Projector		Right Projector	
	x	y	x	y
red	0.005	0.006	0.009	0.013
green	0.009	0.004	0.011	0.006
blue	0.010	0.062	0.012	0.066
white	0.039	0.042	0.040	0.064

To test for additivity, we compare the sum of the tristimulus values for red, green and blue with the measured grayscale. The average error in X, Y and Z was 4.75%, which is rather high compared to monitors and high-quality flat panel displays.[4] We note, however, that performing the same comparison on the raw measurements (without subtracting black) gives an error of over 25-30%. We also measured a 5 x 5 x 5 set of colors, spaced uniformly in pixel value throughout the gamut. We then compared these measured values to the sum of red, green and blue, using values interpolated from the 31-step primary data. While 80% of the computed values were within 5% of the measured values, the remaining 20% showed significantly higher errors. These were distributed primarily along the red-green plane for both projectors. It is not clear whether this is an artifact of the measurement and evaluation process, or a specific flaw in this particular type of projector.

Because the black level is so high, and varies from projector to projector, it may be important to include it in the calculation of the characterization matrix. The mapping from pixel to intensity values is not necessarily a gamma curve, so it is best to represent it as a sampled function. More formally, if p is an input pixel value, there is a function $ITF(p)$ that maps p to normalized intensity. Let c be the color intensity vector computed from the *ITFs*, and $[X_R, Y_R, Z_R]$, $[X_G\ Y_G\ Z_G]$, and $[X_B\ Y_B\ Z_B]$ be the measured 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:

$$t = cM + t_K$$

$$M = \begin{bmatrix} X_R - X_K & Y_R - Y_K & Z_R - Z_K \\ X_G - X_K & Y_G - Y_K & Z_G - Z_K \\ X_B - X_K & Y_B - Y_K & Z_B - Z_K \end{bmatrix} \quad (1)$$

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

$$c = (t - t_K)M^{-1} \quad (2)$$

These transformations can be defined as a single 4x4 homogeneous transformation matrix as shown in equation 4, and its inverse.

$$\begin{bmatrix} R & G & B & 1 \end{bmatrix} \begin{bmatrix} X_R - X_K & Y_R - Y_K & Z_R - Z_K & 0 \\ X_G - X_K & Y_G - Y_K & Z_G - Z_K & 0 \\ X_B - X_K & Y_B - Y_K & Z_B - Z_K & 0 \\ X_K & Y_K & Z_K & 1 \end{bmatrix} = \begin{bmatrix} X & Y & Z & 1 \end{bmatrix} \quad (3)$$

This convenient representation is commonly used in graphics systems and hardware, and is a convenient way to include the black offset in the characterization

.Most commercial projectors are optimized for displaying video, so a typical intensity transfer function will approximate a gamma curve, such as the ones shown in figure 4. These are the result of image processing hardware within the projector, as the native response of the LCD imaging element is not a power function. To get the best results, these curves should be set to smoothly cover the entire output range, with no "clipping" at either the white (contrast too high) or black (brightness too low) ends.

Characteristics of DLP projectors

Small DLP projectors, like the Compaq MP 1800 used in the Interactive Mural, are not simple RGB systems because they include a "white" filter on the color wheel. This makes the white about 145% of the sum of the RGB primaries. This can be seen in figure 5 which plots output luminance for an 11-step gray ramp. The input values have been normalized to the summed RGB luminance. This is equivalent to applying the pixel to luminance ITF's for each primary

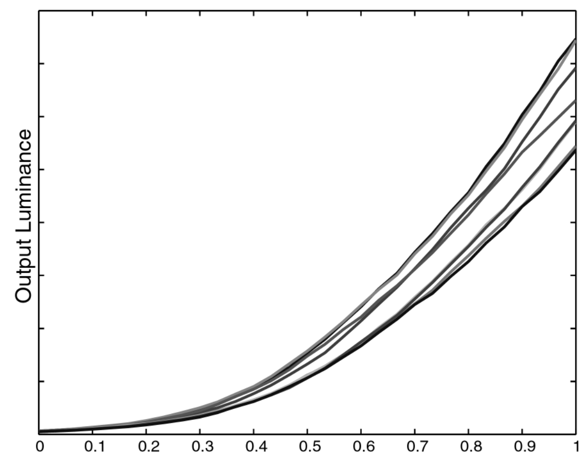

Figure 4: ITFs for eight NEC MT1030 LCD projectors. These curves can be changed by adjusting brightness (raises whole curve) and contrast (raises the maximum value) using the projector menu

The algorithm for applying the white filter was published by Texas Instruments in 1998.[7] The white filter is added in three fixed amounts. The first step occurs at the maximum brightness for the sum of R, G and B. These are reduced at the same time the white is added, to maintain a constant brightness. When this sum reaches its maximum, more white is added and the primaries are again reduced, similarly for the final step. The R, G and B values are not reduced uniformly. There is a calibration step where the balance of the RGB values is determined to maintain a gray ramp without hue shifts. The published specs are variation under 3 ΔE in u*v*, and 1 ΔE in L*.Figure 6 shows an XYZ scatter plot of a full gamut of an MP1800. Each edge of the color cube has been highlighted by overlaying a line plot of the colors at the edge. The white point that would be achieved by summing R, G and B is shown as dashed black lines. The scatter plot is a 9x9x9 array, but the line plots come from an 11x11x11 set of data. The effect of the white segment is clearly visible as an extension of the white point

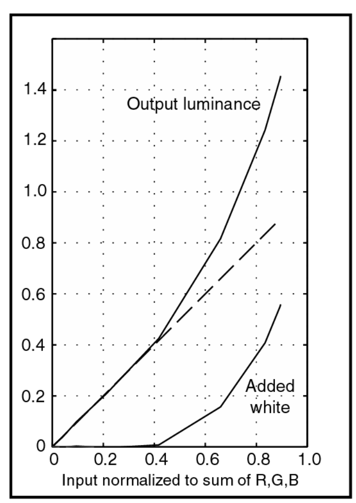

Figure 5: Plot of output luminance vs. summed RGB luminance. The added white increases the output luminance 43%

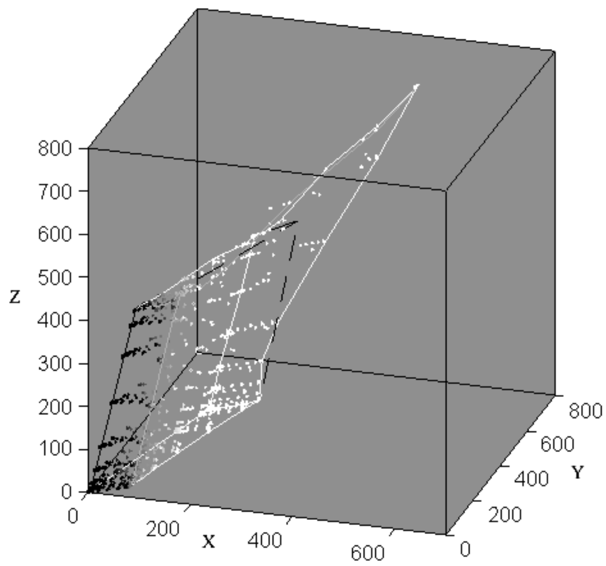

Figure 6: A scatter plot in CIE XYZ tristimulus space of data taken from a Compaq MP1800 projector. The extended white point is clearly visible

These data were taken with an X-Rite DTP92 colorimeter, which is designed to measure monitors. Therefore, the various bumps and wiggles in the figure probably should not be taken as significant. However, the general shape of the gamut reflects what the model predicts: an additive gamut with an extrusion at the white point.

Figure 7 is a plot of the spectral distribution for the red, green, blue, black and white light of a DLP projector, again showing the effect of the added white. The dashed lines show the sum of the R, G, B spectra, and the difference between this sum and white.-

Because the white filter simply adds another additive component, the transformation from RGB to XYZ can be characterized as follows:

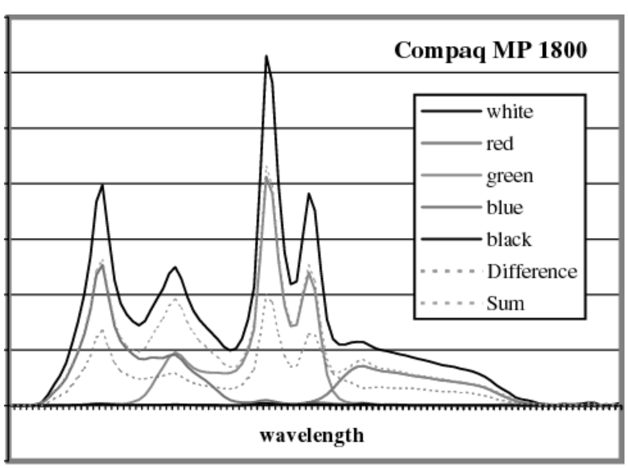

Figure 7: Spectral distribution curves for the Compaq MP 1800

$$\begin{bmatrix} R' & G' & B' & W & 1 \end{bmatrix} \begin{bmatrix} X_R - X_K & Y_R - Y_K & Z_R - Z_K & 0 \\ X_G - X_K & Y_G - Y_K & Z_G - Z_K & 0 \\ X_B - X_K & Y_B - Y_K & Z_B - Z_K & 0 \\ X_W - X_K & Y_W - Y_K & Z_W - Z_K & 0 \\ X_K & Y_K & Z_K & 1 \end{bmatrix} = \begin{bmatrix} X & Y & Z & 1 \end{bmatrix} \quad (4)$$

However, the input R, G and B must be modified using logic that compares the summed luminance to the target luminance, which is indicated by R', G' and B' in equation 4. The logic is published, and is conceptually simple. But, it includes parameters that must be calibrated for each projector. Furthermore, the matrix in equation 4 is not invertible.

The ITF's we measured for the MP1800 are similar to those for the NEC projectors, except that they roll-off at white. Again, these are manufactured curves. Grayscale in a DLP projector is created by pulsing the mirrors, which are binary devices.

Cross-projector comparisons

Figure 8 shows the red, green, blue primary colors plus white plotted on the 1931 CIE chromaticity diagram for both the NEC (small crosses) and the Compaq projectors (small dots). Variation can occur both from the filters and, more significantly, from the color of the bulb. Bulb color can vary significantly, and changes with age, to the point that some groups keep a set of carefully matched "demo bulbs" for special occasions.

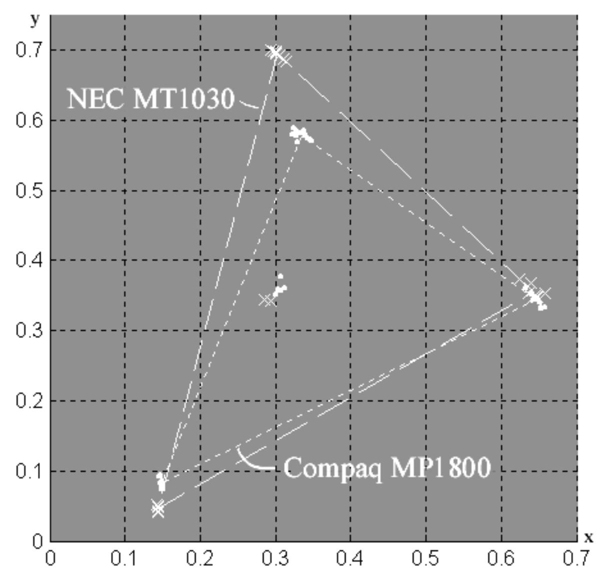

Figure 8: Chromaticity plot of full red, green, blue and white fo[r] eight NEC LCD projectors (crosses) and twelve Compaq DLP pr[o]jectors (dots). The triangles show the average gamuts.

Table 2 shows the average deviation for the red, green and blue primaries and white for both types of projectors. of the primaries and white for both types of projectors. This data represents twelve Compaq projectors and eight NEC projectors. These variations are all quite small, comparable with the variation shown within a single projector (see table 1). This suggests that the primary cause of visible color variation is simply the difference in relative brightness.

Table 2: Average deviation of the chromaticity coordinates

	Compaq MP1800		NEC MT1030	
	x	y	x	y
red	0.011	0.024	0.013	0.027
green	0.026	0.010	0.022	0.008
blue	0.008	0.078	0.002	0.096
white	0.013	0.024	0.016	0.025

Using small DLP projectors results in a smaller gamut and the characterization problems introduced by the white filter. However, they have superior contrast and substantial size advantages over similarly priced LCD projectors. The smallest DLP projectors are approaching 3 pounds, and advertise a 400:1 contrast ratio. The image is crisp and bright and can easily be viewed with the lights on. Their small size makes it easy to build structures to hold and align them. There is substantial brightness variation, as in all projection displays, but no visible hue variation across the image.

LCD projectors have larger gamuts and a potentially simpler characterization model. However they may have a pronounced color "mottle" across the image. The old mural, with the NEC projectors, had purple and yellow blotches on each tile. This is caused when light from the image strikes the dichroic filters at different angles. It can be avoided with careful engineering, but small, inexpensive projectors often demonstrate it. LCD projectors also emit polarized light, which can visibly interact with other polarizing elements in the system.

There are also DLP projectors with three imaging elements and dichroic filters like the LCD projectors, though these tend to be much larger than the ones we used in our system. These can be characterized with the CRT model,[10] and are the technology used in digital cinema. As these become smaller and cheaper, they may provide the best performance of all the projection technologies.

A few notes on measurement

The easiest way to measure a back projection display is to use an instrument that can be attached to the screen. Because the brightness of projection displays varies both spatially and with viewing angle, an instrument with a lens must be positioned precisely in front of each tile to ensure accurate comparison between each tile. Over a wall-sized display, this is difficult to achieve without special alignment hardware. We have heard, for example, of a group that uses a digitally controlled telescope mount to position a PR-650 for such applications.

The X-Rite DTP92 monitor colorimeter used for the measurements in this paper was purchased and integrated into our color measurement software several years ago. It is far from ideal. Projection displays are now much brighter (around 1000 cd/m2 for the Compaq projectors as we use them) than monitors. To compensate, we either measure a dim portion of the display or use a neutral density filter. Comparing the X-Rite DTP92 monitor colorimeter with a PhotoResearch PR-650 spectroradiometer, we have seen significant variation in the X and Z measurements. It is possible that the additivity errors we observed can be attributed to the colorimeter. In the future, we plan to try the Spectrostar Spectrocam, which is a desktop spectroradiometer.

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 also eliminates any contributions from extraneous light as long as the measurement conditions are kept constant. This is especially convenient for research displays, which are not usually sealed into light-tight "cubes" like commercial systems.

The black measurement in the characterization, however, is an absolute measurement. Fortunately, it is fairly bright compared to monitor and display "flare." 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. This is such an abnormal condition, however, that it might be perceptually more accurate to include all the light visible during normal operation.

Conclusions

Tiled projection displays that appear virtually seamless are a rapidly growing commercial market. They generally include large, customized projectors with special electronics for color balancing and matching brightness characteristics at the edges. These systems are generally very expensive, and not very flexible.

Groups that create projection displays out of commodity projectors are primarily experimenting with applications of large-format displays, some of which even wrap around the walls. While a seamless, uniform display would be nice, it is not essential. What is more important to these groups are procedures that are simple and automatic, ideally using cameras, rather than specialized color measuring equipment.

The number of applications of projection displays is growing as digital projectors become smaller, cheaper, and higher quality. Understanding their color characteristics will become more important as their uses multiply. Creating an accurate ICC profile, for example, requires at minimum including the black point. The best way to characterize a small DLP projector may be sampling. But, the narrow white point may provide a challenge for existing sampling and gamut mapping algorithms.

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

References

[1] R. S. Berns, "Methods for characterizing CRT Displays," *Color Research and Application,* 16, 1996, pp. 173-182.

[2] D. H. Brainard, Denis G. Pelli, Tom Robertson. Display Characterization, *The Encyclopedia of Imaging*

Science and Technology, John Wiley and Sons, New York, to be published in 2001.

[3] W. B. Cowan, "An Inexpensive Scheme for Calibration of a Colour Monitor in Terms of CIE Standard Coordinates," *ACM Computer Graphics*, Vol 17. No. 3, 1983, pp. 315-322.

[4] J. E. Gibson and M. D. Fairchild, Colorimetric Characterization of Three Computer Displays (LCD and CRT), Munsell Color Science Laboratory Technical Report, January, 2001.

[5] E. Giorgianni and T. Madden, *Digital Color Management*, Addision-Wesley, Reading, MA, 1998.

[6] L.J. Hornbeck, "Digital Light Processing and MEMS: Timely Convergence for a Bright Future," *Micromachining and Microfabrication '95: Part of SPIE's Thematic Applied Science and Engineering Series,* Austin, TX, October 1995.

[7] G. Humphreys, M. Eldridge, I. Buck, G. Stoll, M. Everett, P. Hanrahan, WireGL: A Scalable Graphics System for Clusters, SIGGRAPH 2001 Conference Proceedings, ACM, New York, NY, 2001, pp. 129-140.

[8] W. Kunzman, G. Pettitt, "White Enhancement for Color Sequential DLP," *SID Conference Proceedings,* 1998.

[9] K. Li and T. Funkhouser, eds., "Large Displays," *IEEE Computer Graphics and Applications*, Vol 20, No. 4, July/Aug 2000.

[10] A. Majumder, Z. He, H. Towles, G. Welch. "Achieving Color Uniformity Across Multi-Projector Displays," *Proc. IEEE Visualization 2000,* IEEE Press, Piscataway, NJ, Oct 2000 pp

[11] O. Packer, L. C. Diller, J. Verweij, B. B. Lee, J. Pokorny, D. R. Williams, D. M. Dacey, and D. H. Brainard, "Characterization and use of a digital light projector for vision research," *Vision Research,* 41, 2001, 427–439.

[12] M. C. Stone. "Color and Brightness Appearance Issues in Tiled Displays," *IEEE Computer Graphics and Applications,* Vol 20, No. 5, September/October 2001, pp. 58-66.

[13] E. Stupp and M. Brennesholtz, *Projection Displays*, John Wiley & Sons, New York, NY, 1999.