Multi-Channel Display Systems for Data Interpretation and Command and Control

Robert M. Clodfelter and Yoav Nir

Barco Simulation Products, 600 Bellbrook Avenue, Xenia, OH 45385

ABSTRACT

Military, academic, and commercial decision makers need to review large quantities of high-resolution visual information in real time. Data sets ranging from millions to billions of pixels or voxels are not uncommon and the number of simultaneous users many vary from one to tens of users. A large area display of 2m by 4m, not an uncommon size, requires 95 million pixels to achieve near eye limiting resolution when viewed a 1m distance. The required size and resolution of these images far exceeds the ability of a single projection display. It is therefore necessary to devise techniques to seamlessly tile images from multiple projectors to meet this need.

This paper discusses the difficulties in tiling multiple projectors, and the methods currently used to overcome them. Large area images exhibit geometric distortions caused by projector misalignment, imperfect lenses and folding mirrors and those from the displays themselves. The images, if uncorrected, generally have luminance and color variations both between, and within, each projector channel. Finally, the images must be tiled in such as way as to have smooth transitions between each of the projector channels.

Keywords: Large-format projection displays, high resolution, projection, tiled displays

BACKGROUND

Conventional display systems cannot achieve the size and resolution required for many complex decision making tasks that rely upon multiple complex datasets for multiple users. Several factors contribute to this fact. The location of a projector is determined by the throw ratio of its lens. This is the ratio of the width of the image divided by the throw distance from the projector to the screen. The lower this ratio, the shorter the throw distance for a given screen size. The lowest practical throw ratio is 0.8:1. Using this value, a 4m wide image requires a throw distance of 3.2m. By using two projectors in this example, the depth of the system can be reduced by half to 1.6m. It can be concluded that for simple optical reasons, tiled projection solutions will use less floor area than a single projector solution, even if such a projector were possible.

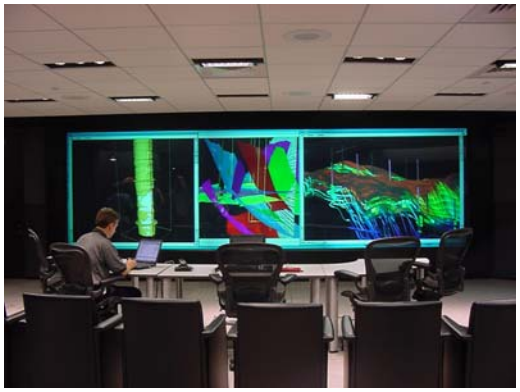

Figure 1: Three channel tiled display system for visualizing complex data sets.

The images to be displayed are delivered to the projectors via video cables. The signals may be in an analog or digital form dependent upon the source and capabilities of the projector. The signals carry information for each pixel in the image at a typical rate of sixty images per second. This means a 95 million-pixel image is transmitted to the display system at 5.7 billion pixels per second. Such a rate is impractical today for many reasons, including the ability of a single video source to create and generate such an image, the ability to practically transmit the image from the source to the projector, and the ability of the projector to process and display the image. However, the ability to generate such complex data sets is clearly possible, and is being done using clusters of PCs or other graphics engines. Such a configuration for the source data matches nicely to a tiled configuration in the display system. The total data rate to the display system remains the same, however, it is divided evenly over multiple parallel paths, thereby easing the transmission and processing requirements on any one particular path.

Should the optical and electrical problems be overcome, a final difficulty remains. The cost per pixel of a projector rises dramatically with the total number of pixels. A QXGA (3.14M pixel) display costs ~$71 per 1000 pixels while a SXGA

(1.3M pixel) display costs only ~$38 per 1000 pixels, almost half the cost per pixel. At commodity XGA (786k pixel) resolutions, the cost per 1000 pixels may be as low as $3-4. There are solid physical reasons for these cost differences: projectors with greater resolutions necessarily use larger light modulators, and possibly smaller feature sizes on the modulators leading to lower yields and therefore higher unit costs. CRT projectors were not considered in this cost comparison, as most large format displays require lumen outputs that exceed CRT capabilities. However, it is interesting to note that at QXGA resolutions, the cost per 1000 pixels of a CRT displays is less than one fifth of completing matrix technologies.

The convergence of the desire to examine large data sets, the availability of very low cost projectors, and the use of PC clusters to generate these sets has led to multiple academic, military, and commercial efforts to develop, understand, and use tiled projector systems. Development systems exist at many locations around the world including Princeton University – Princeton Display Wall, University of North Carolina at Chapel Hill – Pixel Flex, University of Minnesota - PowerWall, University of Illinois at Chicago – Infinity Wall, Stanford University – Interactive Workspaces Project, MIT, Fraunhofer Institute (Germany), Lawrence Livermore National Laboratory, Argonne National Laboratory - ActiveMural, Sandia National Laboratories, Air Force Research Laboratory Information Directorate – Interactive DataWall, Sydney VisLab – High Density Display.

INTRODUCTION

The requirements for a large area high-resolution display are diverse. Applications may require a number of different system configurations such as front or rear projection; different screen shapes from spherical, to cylindrical, to conic, to flat, for example. Screen sizes may range from 1m x 1m to 3m x 20m flat to a 30m diameter hemispherical dome. The system may operate in mono or stereo modes using active or passive means to deliver the proper image to the eyes. Users may require high brightness, high contrast or both. The system may be for a single user or for multiple simultaneous users. However, the tiling problem remains substantially the same across all these variations.

A perfect projector does not exist. Even if it did, no one could afford it. However, modern projection devices maintain an exquisite balance between performance, features, and cost. When used singularly, projected images appear very good and design compromises are virtually unnoticeable even to the trained observer. However, when tiled, the eye becomes a harsh diagnostic instrument revealing all. When two projectors are used side by side, the most casual observer notices that the images do not align, their colors do not match, and they are probably not the same brightness. Many of these conditions are caused by the projectors themselves and not how they are used.

The first-order solution towards these tiling problems is to select projectors that have the appropriate mix of features and design trade-offs for tiling. To correct geometry problems, look for projectors that have low distortion lenses. If a zoom lens is used, operate it at the point of least distortion. If the projector is to be used off axis, meaning the axis of the lens is not perpendicular to the center of its tile, look for a design that allows for optical lens shift, thereby eliminating keystoning of the projected image. Some models have lens shift in both the horizontal and vertical axis, allowing displacements in both of these directions.

Color and luminous variations between matrix projectors are caused by variations in the color spectrum of the lamps used, dichroic mirror tolerances, lack of uniformity in the illumination system, poor gamma correction in the modulators, and vignetting in the lens. The tolerance problem is one of cost. By using lower cost dichroic mirrors, projector costs can be reduced, however, in a tiled display, this tradeoff may not be the best. Typically, color correction for this type of error requires more than a simple 1-dimensional lookup table for each color to properly match the color spaces between projectors. The quality of the design of the illumination system and lens are, again, strongly constrained by cost; leading to compromises in the projector's design that may not be ideal in a tiled projector display.

SYSTEM CONSIDERATIONS

How many projectors should a tiled system use? The simple answer is as few as possible. Pixels are expensive, even commodity pixels. The desired screen size and shape must be taken into account, as well as the tiling possibilities. Most projectors currently create 3:4 (XGA) or 4:5 (SXGA) aspect ratio images. The aspect ratio is simply the ratio of image height to width. UXGA projectors, which are just becoming available in quantity, return again to the more typical 3:4 format. However, the long-range commodity sweet spot for aspect ratios may be 16:9, the HDTV format. The projectors should be arranged to optimally fill the required area with an overlap that ranges from 5% to 20%. Typical overlaps are 12.5% and 20%. This means that for every pixel purchased, up to 20% will be wasted in the overlap areas used to create a smooth transition from one projector to the next. Projectors are typically tiled in rectangular grid when possible. Such an arrangement has the unfortunate characteristic that blend areas may include up to four overlapping corners. However, it does fill typical screen arrangements nicely.

The number of projectors is strongly dependent upon the resolution required of the system. Keep in mind the degree of overlap anticipated in the design. In our example of a 95 Mpixel wall, using UXGA projectors with 1600x1200 resolution, nearly fifty projectors are required if no overlap were used. However, anticipating a 12.5% overlap suggests at least 65 projectors. Further increases may be necessary to properly tile these projectors into a regular grid that meets the system requirements.

The light output of the tiled projectors is also a concern. The desired screen brightness should be determined from the user's requirements according to how the display is to be used, the expected ambient light, and the desired system contrast. Each projector's brightness can be calculated based on the resulting number of channels and screen gain characteristics. True high-resolution displays generally have a surplus of light available allowing the use of low gain screens to improve the system contrast and simultaneously increase the viewing volume of the display.

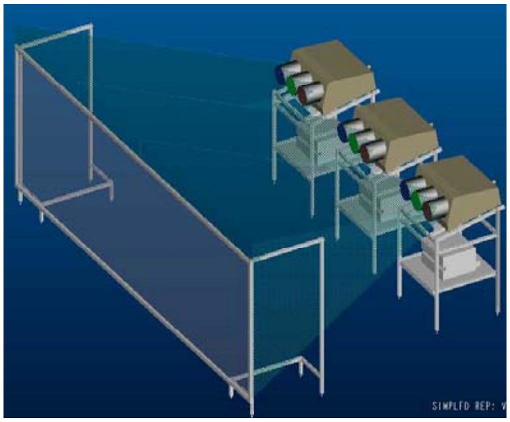

Figure 2: Three channel tiled display system showing projectors, stands, screen, and overlap areas.

The viewing volume of the display is defined as the volume where the image formed from the tiled projectors appears acceptable to the observer. It is not immediately obvious that a viewing volume would exist or that the display would not appear correct from all viewing positions. In fact, it doesn't. This is because the light reaching a viewer depends upon the intervening geometry between the viewer, the screen, and the projectors. As the viewer moves around the viewing volume, these relationships change causing changes in the apparent brightness of the individual projector channels. When these differences become great enough, the user is able to distinguish the individual tiled channels forming the image. To reduce the effect of these geometric changes, and therefore increase the viewing volume, low gain screens are used.

An ideal unity gain screen will emit incident light equally in all directions regardless of the angle of incidence. With such a screen, the viewing volume is unlimited. Such screens do not exist. However, using a low gain, high quality screen is a must when a large viewing volume is desired.

The screen may limit the system in other ways. Screen materials can only be manufactured to certain finite sizes. Flexible screens are generally the largest, but may be unsuitable because of their propensity to move under the influence of air currents or the desire to interact with the screen by touching or pointing. Large, rigid screens are limited in size based on manufacturing limits but, in a more practical way, can be limited by the physical possibility of getting the screen to the display system for installation.

As mentioned previously, the depth of a rear projected tiled system is controlled by the throw ratio of the projection lens. As multiple projector channels are used, the depth of a tiled system will be much less than a non-tiled system using an equivalent lens. It is possible to further reduce system depth by using folding mirrors in the optical path from projector to the screen. The mirrors should be high quality front surface mirrors, very flat, and carefully mounted to prevent geometric or focus problems in the resulting image. Using mirrors to reduce system depth is a difficult and surprisingly challenging task to do successfully.

PROJECTOR MODULATORS

Many projection technologies exist to create tiled displays. The oldest is uses the cathode ray tube or CRT. CRT projectors offer extremely high quality images at moderate costs. Their drawbacks are that they are not very bright; they may require day-to-day tuning to maintain a high quality image; they tend to be large; and they do not have uniform luminance across their display field. Their cost per pixel is very competitive at UXGA and QXGA resolutions. Most current research is not using CRT projectors. However, it is interesting to note that much of the research on automatically correcting geometric distortions inherent in matrix based display solutions could be applied to CRT projectors, eliminating the day-to-day tuning problem. As an added bonus, they would reduce the processing overhead in the image generator and reduce the latency of the total system.

Most tiled projection systems use matrix light modulation technologies such as poly-silicon Liquid Crystal Display (LCD) or a Micro-Electro-Mechanical Systems (MEMS) using mirrors. Poly-silicon solutions are available from many suppliers. All commercially available MEMS based projectors are based on Texas Instruments Digital Light Processing (DLP) technology. Either technology has its strengths and weaknesses depending on how it is used and the types of images it is required to display.

LCD projectors can be divided into transmission and reflective categories. Light passes through transmisive panels where the polarization of the light is modified on a pixel-by-pixel basis according to its desired brightness. Reflective panels reflect light while modulating its polarization on a pixel-by-pixel basis. The polarization state of the light can be continuously varied so that each pixel's brightness is also continuously variable from off to full on. In either category, three panels are used, one for each of the red, green, and blue colors used in the projector to create a full color image. Resolutions are possible up to QXGA. The switching time of a pixel from on or off or off to on can be long enough to cause smearing in an image where fast moving objects appear. Special processing is required to reduce this effect.

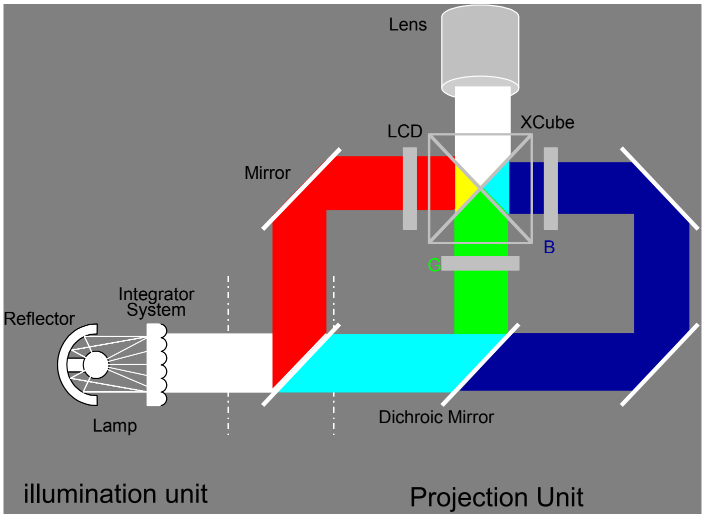

Figure 3: Three panel LCD Projector showing lamp, color separating dichroic mirrors, LCD panels, color combining X-cube and lens.

DLP technology uses a small mirror for each pixel in the display. An inherently digital process controls the brightness of each pixel. A mirror or pixel can only be fully on or fully off. Gray scales are created by a method known as Pulse Width Modulation (PWM). By using PWM techniques, the average time a pixel (or mirror) is on during a complete frame is proportional to the desired brightness of the pixel. While such a digital approach can lead to a very uniform and linear display, it can also result in potentially disturbing image artifacts. PWM sequences must be carefully designed to reduce this effect. DLP processing requires a full frame time before the image is displayed. This delay adds another frame time giving a total system latency due to the modulator of two frames, or 33mS at 60Hz refresh. However, DLP projectors are known for their excellent uniformity and color linearity resulting from the PWM method used to create gray scales on the panels.

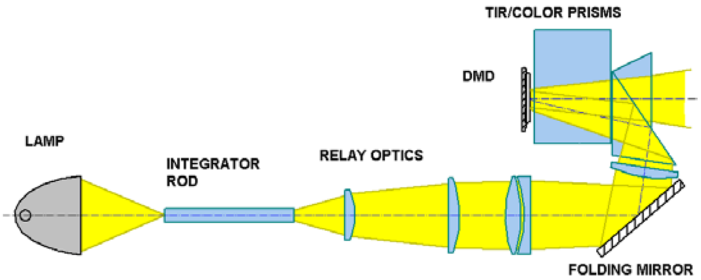

Figure 4: Three-Panel DLP Projector showing lamp, integrator rod, relay optics, and TIR prism.

DLP projectors are divided according to the number of DLP modulators they use. Single chip DLP projectors use a single modulator to reduce cost. The display works by flashing red, green, and blue images in rapid succession. The repetition rate is sufficiently high that the individual color images are fused into a single full color image. A variant on the single chip projector flashes red, green, blue, and white images. The white image is added to increase the total brightness of the display.

The second design uses three DLP modulators much the same as the LCD display. Each modulator receives a continuous flow of red, green, or blue light that it modulates and is then projected to the screen via the lens. Three chip designs have the advantage of no color break up that can be a complaint of the single chip designs. They also have the ability to sustain much high brightness and superior color fidelity.

COLOR MATCHING

In all projectors, regardless of modulator, it is important that the color and quantity of the red, green, and blue light is matched between projectors. Systems such as the High-Density Display at the Sydney VisLab achieve this goal by using a single light source for all projectors. The light source is connected to each projector using three fiber optic cables. This method has the advantage that all variations in lamp brightness and color, if properly designed, are carried equally to each of the projectors. This means that although the image may vary, it varies equally among all the tiled channels. Efficient coupling between the light source and fiber bundles used to carry the light to the individual projectors is difficult to achieve resulting in lost light and increased lamp power to overcome these losses.

Another solution is to electronically or optically stabilize the lamp output in each projector. This method also reduces the effect of light output variations due to lamp aging. It does not require additional cables and has the advantage that if a lamp fails, the remainder of the wall remains operational. Since each projector has its own lamp, the power of the individual lamps will be less than that of a central lamp leading to longer lamp life and greater reliability.

A further enhancement of lamp stabilization is to link the stabilizing electronics between projectors using a simple electrical connection. This link allows the stabilization circuits to track the weakest (dimmest) lamp in the system. The system can continue to operate with no deterioration in tiling even when a lamp has passed its ability to be stabilized at full light output.

Once the light source is stabilized, the individual colors must be matched. The white light from the lamp is divided into red, green, and blue light via dichroic mirrors. Dichroic mirrors alternatively reflect or transmit light based upon its wavelength. The wavelength of the transition region from transmission to reflection must be carefully controlled to maintain consistent color between projectors. Typical tolerances range from 2-5nm. Additional optical adaptations are possible to further match channel-to-channel colors.

The color of the projector's primaries can be further corrected electronically. By using a 3x3 matrix transform in linear color space, the color primaries can be mapped consistent color coordinates from channel to channel. This transform can be applied in either the projector or the image generator. As negative light is not possible, the resulting color gamut can only be smaller than the original. The color gamut for the entire system is therefore the smaller than all the color gamuts combined.

IMAGE WARPING

Tiled systems using flat screens can be designed and built without the need for image warping capabilities. By avoiding warping, system cost and latency may be reduced. Systems employing curved screens or off-axis optical designs beyond the capability of lens shift to compensate necessarily require some form of image warping. When warping is available, the tolerances required to position and optically match the projectors are drastically relaxed. Images may be warped inside the image generator, by an external video processing box, or by the projector itself.

Low cost systems perform image warping in the image generator. Such rendering is usually a two-pass operation for every frame. The desired image is computed and stored in texture memory. In a second pass, the image, in the form of a texture, is mapped on to a warping surface to form the final warped image. The advantage of this approach is its low cost and that it operates in linear luminance space. The disadvantage is that it adds additional frame delays in the image generation chain increasing the total system latency. The quality of the warped image may be significantly lower than that of a dedicated hardware solution due to the use of simple but fast bi-linear interpolation in typical IG warping methods.

Dedicated image warping solutions may be implemented in external hardware or internally to the projector. The quality of warped images varies considerably between systems as it depends on the warp algorithms used. Many commodity systems warp using a two-pass process that distorts the image sequentially in the horizontal and vertical directions. This approach reduces the amount of hardware required but, as in IG based distortions, adds at least one additional frame delay of latency to the system. Commodity systems also tend to use simple bi-linear interpolation methods. While this interpolation method produces images superior to nearest neighbor algorithms, much of the image's original resolution is lost. Such processing is clearly contrary to the goal of producing a very high-resolution display system.

It is possible to design warping systems that overcome all these obstacles. Delaying the warping function only long enough to insure the arrival of all necessary pixels to begin processing minimizes latency. The video inputs must be linearized to eliminate the video gamma function typically used in the IG output look up table (LUT). Simple bi-linear sampling is replaced with bi-cubic or greater interpolation methods that utilize at least sixteen adjacent pixels for interpolation. Adaptive filters may also be used to further enhance the final image. Although not strictly necessary, integrating the warp electronics directly into the projector has the additional benefits of avoiding unnecessary digital to analog conversions and gamma conversions of the video stream thereby maintaining its dynamic range.

When warping is available inside the image generator, in an external box, or by the projector itself, the projectors should be positioned and aimed according to the system design but their placement will be nowhere near as critical as that of a non-warp system. The quality of the system alignment will depend principally on the adjustment of the image warping parameters, and not on the placement and aiming of the projectors.

PROJECTION OPTICS

All projection lenses introduce small geometric distortions into their images. When used stand-alone, this distortions are sufficiently small to be unnoticeable. However, when multiple projectors are arranged one to the other, differences in these distortions become very apparent. Distortions as large as 1% are possible. This means the ideal and actual locations of an image of a pixel may be displaced by a distance of 1% of the image diagonal. On a SXGA (1280x1024) projector, this error is greater than 10 pixels. Clearly, a misalignment of 10 pixels would be very noticeable across a seam. These errors can be partially cancelled by using the same type of lens and same lens shift setting (if available) on all the projectors in the tiled array. When operated in this manner, the lens distortions should match at the periphery of each image and only the variances of the distortions will remain. Having said that, barrel or pincushion distortions will cause the width of blend zones to vary along the length of the seam while projected lines parallel to the seam will not properly overlap.

PROJECTOR POSITIONING

Complex interactions of lens distortions, off-axis projection angles, and lens shift all conspire to make proper projector positioning a difficult and time-consuming task in systems without warping capabilities. The system geometry must be carefully followed. The geometry of the resulting system is completely dependent upon the screen shape, projection optics, and positioning of the projectors. Projector alignment becomes critical to the success of the display as it provides the only degrees of freedom available to geometrically co-align adjacent images in the system. Each projector must be positioned very accurately in six degrees of freedom. To illustrate, a one-pixel error is caused by an angular error of only 2.7 arc-minutes on a SXGA projector using a 1:1 lens. Several adjustable projector mounts providing very precise positioning and aiming of tiled displays are discussed in the literature. The use of templates to locate the design location of the projectors and reference grids attached to the screen also aid in the alignment.

Proper lens shift is critical to properly match adjacent channels and to avoid incorrectly pointing the projector towards the screen in an off axis way. Lens shift moves the image up/down left/right much in the same way as rotating the projector in azimuth and elevation making these degrees of freedom difficult to separate. If used incorrectly, these shifts and rotations will cause keystone distortions in the image that are difficult to detect and properly cancel.

BLENDING

Once the projectors are positioned and aligned, the images must be blended. Edge blending techniques overlap the edges of adjacent projected images and blend the overlapped regions to smooth the luminance and chromaticity transition from one image to the other. Blending methods fall into two major categories: optical, and electronic. Optical or Aperture Modulation blending uses either hard edges or gradients to produce the luminance roll off required. The edges or gradients may be internally or externally mounted. Optical blending solutions have the advantage that both the white and black light levels are blended. Optical blending solutions may or may not be field adjustable. For simple two or three sided blends, almost all optical solutions can be made adjustable. However, complex blends between three or more channels may prove to be more difficult. Hard edge blending solutions are easy to manufacture and are relatively low cost, especially when they are used externally. They are limited in their ability to generate blends of variable width. The width of a hard edge blend is determined by the location of the edge and the exit pupil size of the lens. The resulting penumbra coincides with the region to be blended. Hard edge solutions also suffer when the exit pupil of the lens is not uniformly filled by the illumination system within the projector. In this case, rather than producing a smooth gradient, a hard edge produces a stair stepping effect.

Gradient blending solutions may use continuous gradient patterns or dither patterns to produce a blended edge. Gradient solutions can easily create variable and very wide blend regions. However, such solutions are generally not easily adjustable and may have lengthy iteration cycles between trials. Gradient solutions are also less affected by problems caused by non-uniform illumination within the exit pupil. Gradient solutions using dither patterns may suffer from diffraction problems caused by regular patterns in the dither. Projector light can diffract through these patterns leading to color splitting and dim shifted images in the blend areas.

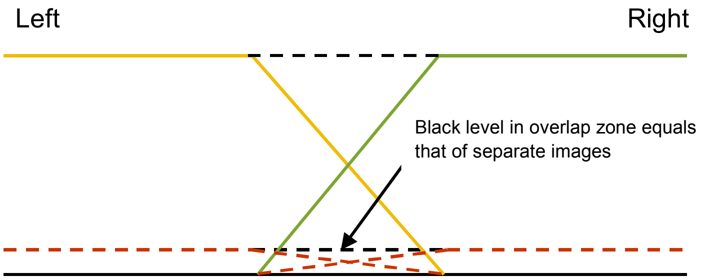

Figure 5: Optical Blend Situation showing roll off of white and black levels.

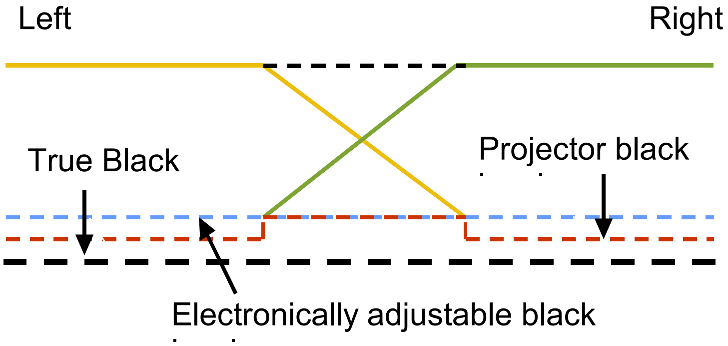

Figure 6: Electronic Blend Situation showing roll off of white levels and compensation of black levels as described in the text.

Electrical blending can be applied in the projector, if so equipped, in an external video processing box, or in the image generator itself as an alpha channel. Electrical blending is desirable as it is easily controlled, conducive to automation, and produces generally good results. However, electrical blending fails miserably for very dark, near black images. This is because all matrix based projectors leak a little light when in the black state. This light cannot be stopped electrically. This means that for black or near black images, the electrical blending does essentially nothing. The result on the screen is that for black images, areas where two projectors overlap are twice as bright as that of a single projector. Areas where three projectors overlap are three times as bright and so on. It is possible to electrically add a compensating black level to non-overlapped areas to equalize the entire black field. While successful, careful consideration should be made concerning the resulting loss of system contrast.

In applications where a significant amount of ambient light will always be present to wash out the projector's minimum black level thereby rendering it imperceptible, electrical blending is acceptable. Similarly, where mostly bright images are to be used, electrical blending may also be acceptable. Electrical blending's flexibility and automation possibilities make it compelling; however, the system designer should take care for how the system will be used and under what conditions before proposing its use.

EXAMPLE SYSTEMS

Figure 7: A two-channel CadWall using 3-chip DLP projectors operating at 110Hz for active stereo applications. External optical blending and folding mirrors are shown. System resolution is 2400 x 1024. Brightness is 10,000 lumens, Contrast 800:1.

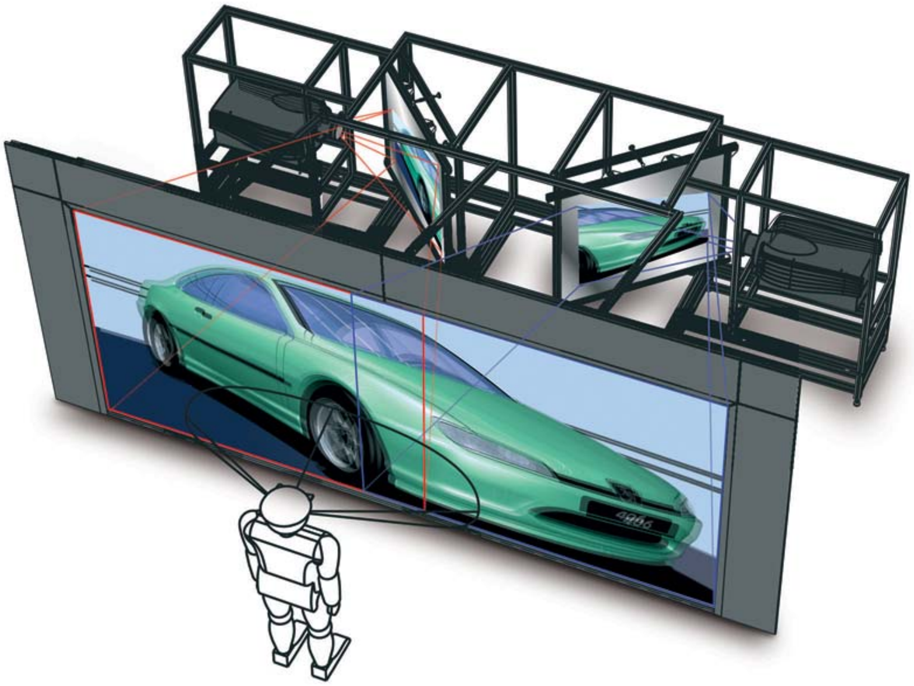

Figure 8: A three-channel front projection cylindrical screen system using single-chip DLP projectors with internal warping and external optical blending. System resolution is 3520 x 1024. Brightness is 2,500 lumens. Contrast 1,000:1.

Figure 9: A three-channel front projection cylindrical screen system using three-chip DLP projectors with internal warping and external optical blending. System resolution is 3520 x 1024. Brightness is 15,000 lumens. Contrast 800:1.

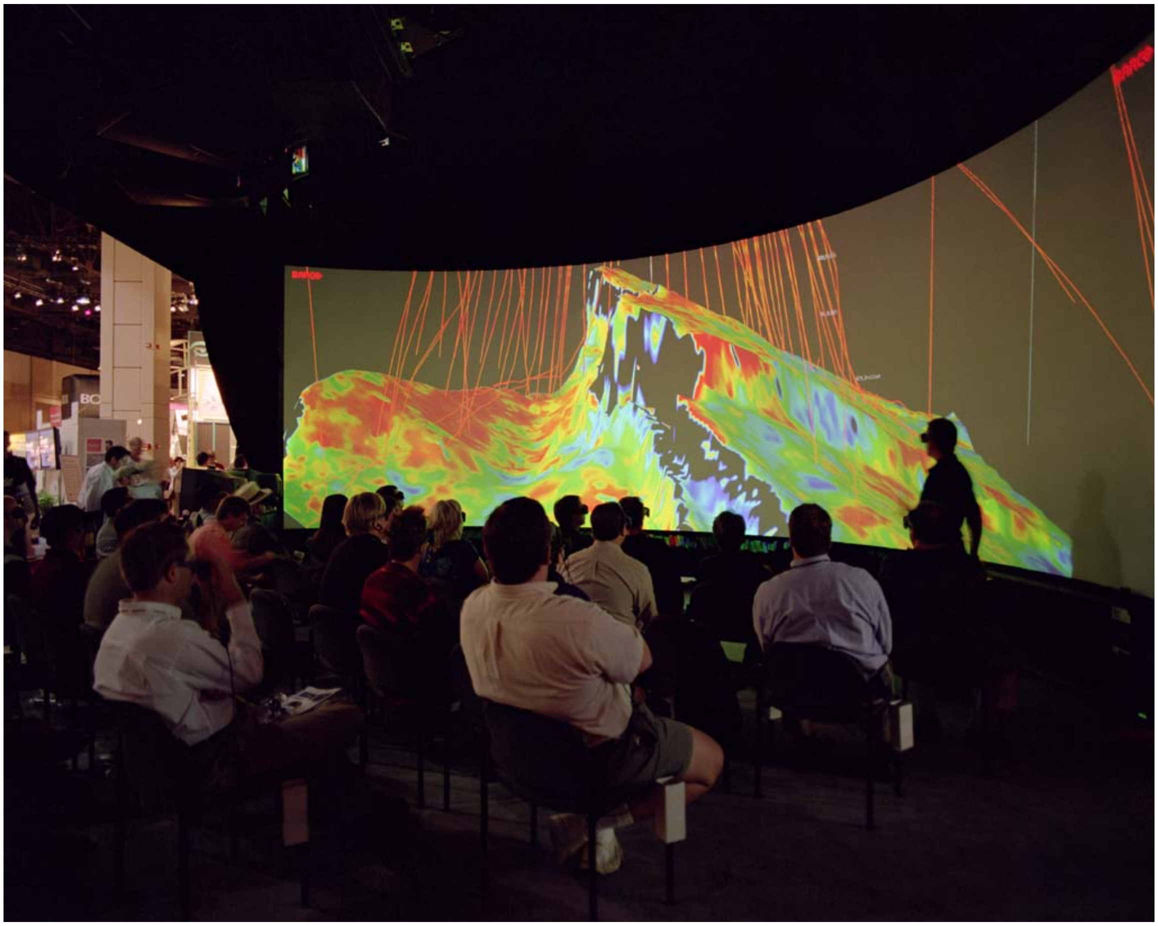

CONCLUSIONS

The confluence of CPU performance, low cost IGs and PC clusters, low cost projection displays, and high quality graphics has produced a plethora of mutli-channel tiled display solutions in universities across the world and are now within the reach of academic, scientific, military and corporate decision makers. Tiled flat wall, cylindrical, conic, spherical dome, and cube systems are now commonplace. Software libraries that allow applications to transparently use multi-channel tiled display systems without modification further drive this phenomenon. Projection displays designed specifically for tiling are now available with features such as luminance stabilization, color matching, distortion correction (warping), edge blending, and luminance and color uniformity correction. These devices, when coupled with appropriate rendering architectures and algorithms, human interface devices, and cluster management tools, make collaboration and visualization tools with unsurpassed performance and cost effectiveness.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the support by Barco Simulation Products in the creation of this paper.

REFERENCES

1. K. Li, H. Chen, Y. Chen, D. W. Clark, P. Cook, S. Damianakis, G, Essl, A. Finkelstein, T. Funkhouser, T. Housel, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. Pal Singh, G. Tzanetakis, Jiannan Zheg, "Early Experiences and Challenges in Building and Using a Scalable Display Wall System", IEEE Computer Graphics and Applications, 20(29-37):671-680, 2000.
2. R. Clodfelter, "*Introduction to Flight Simulator Visual System Displays*," Robert M. Clodfelter, Joe A. Mays, and Richard E. Holmes, Tutorial session, NEACON 1992.
3. R. Yang, D. Gotz, J. Hensley, H. Towles, M. Brown, "Pixel Flex: A Reconfigurable Multi-Projector Display System",
4. Y. Chen, H. Chen, D. W. Clark, Z. Liu, G. Wallace, K. Li, "Automatic Alignment of High-Resolution Multi-Projector Displays Using an Uncalibrated Camera", IEEE Visualization 2000, Salt Lake City, UT, October 8-13, 2000.
5. R. Surati, "Scalable Self-Calibration Display Technology for Seamless Large-Scale Displays", PhD Thesis, Massachusetts Institute of Technology, 1999.
6. R. Van Belle, "High Quality Geometry Distortion Tool for Use with LCD and DLP Projectors", Display Special Interest Group, IMAGE 2000, The IMAGE Society, Inc., Phoenix, Arizona, USA, July 2000.
7. R. Clodfelter, "*New Directions for Light Valve Projectors*," presented at Displays Special Interest Group, IMAGE, 1995.
8. Edited by Durrett, H.J., Interactive Systems Laboratory, San Marcos, Texas, "Color and the Computer", Academic Press, INC, 1986
9. R. Clodfelter, "*Non-Linear Screen Mapping*," presented at Displays Special Interest Group, IMAGE 1997.
10. R. Raskar, "Seamless Projection Overlaps using Image Warping and Intensity Blending", Fourth International Conference on Virtual Systems and Multimedia, Gifu, Japan, November 1998.
11. A. Majumder, Z. He, H. Towles, G. Welch, "Achieving Color Uniformity Across Multi-Projector Displays", Department of Computer Science, University of North Carolina at Chapel Hill.
12. K. Li, Y. Chen, "Optical Blending for Multi-Projector Display Wall System", Proceedings of the 12th LASERS and Electro-Optics Society, Annual Meeting, 1999.
13. R. Yang, D. Gotz, J. Hensley, H. Towles, M. Brown, "Pixel Flex: A Reconfigurable Multi-Projector Display System",
14. R. Surati, "Scalable Self-Calibration Display Technology for Seamless Large-Scale Displays", PhD Thesis, Massachusetts Institute of Technology, 1999.
15. R. Clodfelter, "*Estimating Display System Performance*," presented at Display Systems session, IMAGE 1996.
16. R. Clodfelter, "*Modulation Transfer Function for the Display Engineer*," presented at Advances in Display Technology VI session, O-E LASE, 1986.
17. C. Lloyd, "Quantifying edge blend quality: Correlation with observer judgments", presented at Display Systems session, IMAGE 2002.