Camera-Based Calibration Techniques for Seamless Multi-Projector Displays

Michael Brown*, Aditi Majumder†, Ruigang Yang‡

***Abstract*— Multi-projector, large-scale displays are used in scientific visualization, virtual reality and other visually intensive applications. In recent years, a number of camera-based computer vision techniques have been proposed to register the *geometry* and *color* of tiled projection-based display. These automated techniques use cameras to "calibrate" display geometry and photometry, computing per-projector corrective warps and intensity corrections that are necessary to produce seamless imagery across projector mosaics. These techniques replace the traditional labor-intensive manual alignment and maintenance steps, making such displays cost-effective, flexible, and accessible.**

In this paper, we present a survey of different camera-based geometric and photometric registration techniques reported in the literature to date. We discuss several techniques that have been proposed and demonstrated, each addressing particular display configurations and modes of operation. We overview each of these approaches and discuss their advantages and disadvantages. We examine techniques that address registration on both planar (video walls) and arbitrary display surfaces and photometric correction for different kinds of display surfaces. We conclude with a discussion of the remaining challenges and research opportunities for multi-projector displays.

***Index Terms*— Survey, Large-Format Displays, Large-Scale Displays, Geometric Alignment, Photometric Alignment, Graphics Systems, Graphics.**

I. Introduction

Expensive monolithic rendering engines and specialized light projectors have traditionally made projector-based displays an expensive "luxury" for large-scale visualization. However, with advances in PC graphics hardware and light projector technology, it is now possible to build such displays with significantly cheaper components. Systems, such as Li et al.'s Scalable Display Wall [22], and displays, constructed using Humphreys et al.'s *WireGL* [18] and *Chromium* [17] PC-cluster rendering architecture, have demonstrated the feasibility of cost-effective, large-format displays constructed by assembling many commodity projectors.

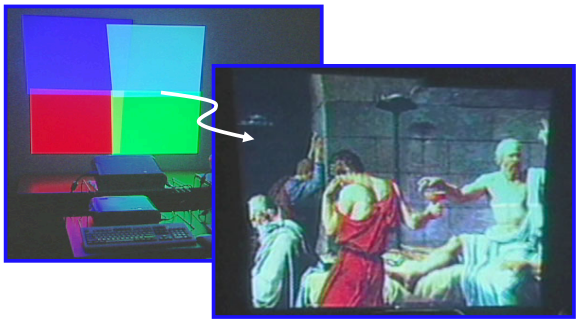

Fig. 1. Camera-based geometric registration is used to calculate image-based corrections that can generate a seamless image from several (unaligned) overlapping projectors.

Images from a multi-projector display must be seamless, i.e., they must appear as if they are being projected from a single display device. This involves correcting for geometric misalignment and color variation within and across the different projectors to create a final image that is both geometrically and photometrically seamless. This correction process is commonly referred to as "calibration". Calibration involves two aspects: *geometric registration* and *color correction*. Geometric registration deals with geometric continuity of the entire display, e.g., a straight line across a display made from multiple projectors should remain straight. Photometric correction deals with the color continuity of the display, e.g., the brightness of the projected imagery should not vary visibly within the display.

Calibration can be achieved through mechanical and electronic alignment, a common approach adopted by many research and commercial systems [12], [22], [18], [14], [34]. Such alignment procedures often require a specialized display infrastructure and a great deal of personnel resources, both to setup and maintain the system. This significantly increases the cost and effort needed to deploy such large-scale, high-resolution displays. Often, half of a displays' total cost is related to the display infrastructure, including the mounting hardware and display surface. In addition, most reasonably sophisticated

*M. S. Brown is with the Hong Kong Univ. of Science and Technology.

†A. Majumder is with the University of California, Irvine.

‡R.-G. Yang is with the University of Kentucky, Lexington, KY.

Fig. 2. Left: This image illustrates the geometric misalignment problem at the boundary of two overlapping projectors. The geometry is noticeable unaligned. Right: This image shows the final seamless imagery of the same projector alignment. Such registration that is the goal of geometric registration methods.

mounting hardware does not have the capability or the precision to correct non-linear distortions like projector radial distortion and intensity non-linearities. Further, manual methods tend to be unscalable. Calibrating even a four-projector system can be severely time consuming.

Recently, techniques have been developed that use one or more cameras to observe a given display setup in a relaxed alignment, where projectors are only *casually* aligned. Using camera-based feedback obtained from a camera observing the display setup, the necessary adjustments needed to register the imagery, both in terms of geometry and color, can be automatically computed and applied through software [42], [40], [37], [10], [50], [9], [20], [38], [28], [27], [30]. The key idea is to use cameras to provide *closed-loop control*. The geometric misalignments and color imbalances are detected by a camera (or cameras) that monitor the contributions of multiple light projectors using computer-vision techniques. The necessary geometric- and color-correction functions necessary to enable the generation of a single seamless image across the entire multi-projector display are determined. Finally, the image from each projector is appropriately pre-distorted by the software to achieve this correction (see Figure 1). Thus, projectors can be casually placed and the resulting inaccuracies in geometries and color can be corrected automatically by the camera-based calibration techniques in minutes, greatly simplifying the deployment of projector-based, large-format displays. In comparison with traditional systems relying on precise setups, camera-based calibration techniques provide the following advantages in particular:

- *More flexility*. Large-format displays with camera-based calibration can be deployed in a wide variety of environments, for example, around a corner or a column or on a poster. These irregularities can cause distortions that traditional systems may find difficult to work with.
- *Easy to setup and maintain*. Camera-based calibration techniques can completely automate the setup of large-format displays. This is particularly attractive for temporary setups in trade-shows or a field environment. Labor-intensive color balancing and geometric alignment procedures can be avoided and automated techniques can be used to calibrate the display in just minutes.
- *Reduced costs*. Since precise mounting of projectors is not necessary, projectors can be casually placed using commodity support structures (or even as simple as laying the projectors on a shelf). In addition, it is not necessary to hire trained professionals to maintain precision alignment that keeps the display functional. Further, since the color variations can also be compensated, expensive projectors with high quality optics (that assure color uniformity) can be easily replaced by inexpensive commodity projectors.

While camera-based calibration techniques require cameras and support hardware to digitalize video signals, these costs are amortized by savings from long-term maintenance costs. Overheads like warping and blending at rendering time to correct for various distortions are reduced or eliminated by the recent advances in graphics hardware (Section V).

In this paper, we present a survey of different camera-based calibration techniques. Our goal is to provide potential developers of large-format displays a useful summary of available techniques and a clear understanding of their benefits and limitations. We start with geometric registration techniques in Section II. We organize different approaches by the types of configurations addressed and the modes of operation accommodated. We discuss techniques for planar or arbitrary display surfaces with stationary or moving viewers. We compare these approaches and point out their positive and negative points for given environments and tasks. In Section III, we focus on color correction techniques. Until recently,

expensive color measurement instruments were used to calibrate the color of such displays and the spatial variation in color was largely ignored. Recent camera-based techniques address photometric variation across such displays for surfaces with different reflectance properties. We discuss all these methods along with their advantages and disadvantages. In Section IV, we list several representative systems that use camera-based calibration and discuss their features. In Section V, we introduce some interesting technological hardware advancements that reduce or even remove the rendering overhead that is typical in performing geometric and photometric corrections. Finally, we conclude in Section VI with a discussion of the remaining challenges that still need to be addressed for projector-based displays.

II. Geometric Registration

When building a multiple-projector display, two types of geometric distortion must be addressed – *intra-projection* and *inter-projection* distortion. *Intra-projector* distortions are distortions within a single projector caused by off-axis projection, radial distortion, and in some cases, display on non-planar surfaces. *Inter-projector* distortions are found between adjacent projectors where edge boundaries do not match. Geometric registration techniques are used to detect and correct both types of distortion (see example in Figure 2).

Camera-based display registration techniques can be divided into two categories based on the type of display surfaces addressed, either *planar* or *non-planar*. We first discuss techniques that assume a planar display surface. These are used to construct large-scale video walls. Later, we extend the discussion to arbitrary display surfaces, for example, multiple planar walls or semi-spherical screens. These scenarios are particularly suited for immersive displays.

A. Planar Display Surfaces

When the display surface is planar, each projector P_k's image can be related to a reference frame, R, on the display surface, via 2D planar homography. This projector-to-reference frame homography is denoted as ${}_R\mathbf{P}_k$ where k is the index of the projector, and the subscript R denotes that the homography maps P_k image to the reference frame R (notation adopted from [9]). To compute the homography, it is necessary to establish four point-correspondences between coordinate frames. Using more than four point-correspondences allows a least-squares fit solution which is often desirable in the face of errors and small non-linearities. In practice, most techniques project many known features per projector to compute the homographies [10], [22], [39], [9], [50], [38].

The alignment of the projected imagery is achieved by pre-warping the image from every projector, P_k, using the homography, ${}_R\mathbf{P}_k^{-1}$. This pre-warp can be performed directly in the rendering pipe-line [50] or by using a post-rendering warp [40].

Thus, the key is to determine the correct, ${}_R\boldsymbol{P}_k$ for each projector P_k. In essence, we need to establish point-correspondences between each projector and the display's reference frame, R. This can be accomplished by using a camera (or cameras) to observe the projected imagery, as shown in Figure 3.

1) Using a Single Camera: We first consider the case when only one camera is used. Figure 3 (left) shows an example of this setup. A homography between the camera and the display reference frame, R, denoted by ${}_R\boldsymbol{C}$, is first computed. Typically, manually selected point-correspondences between the camera image and known 2D points on the display surface are used to calculate ${}_R\boldsymbol{C}$. Techniques that do not require manually points on the display surface to be selected have been introduced [32]. After ${}_R\boldsymbol{C}$ has been computed, projected imagery from each P_k is observed by the camera and a projector-to-camera homography for each projector k, is calculated, and denoted as ${}_R\boldsymbol{C}_k$. The projector-to-reference frame homography, ${}_R\boldsymbol{P}_k$, is then derived from ${}_R\boldsymbol{C}$ and ${}_R\boldsymbol{C}_k$ as:

$$ {}_R\boldsymbol{P}_k = {}_R\boldsymbol{C} \times {}_R\boldsymbol{C}_k, \tag{1}$$

where the operator $\times$ represents a matrix multiplication.

Raskar et al. [39] presented a system using a single camera that could compute this mapping for a 2×2 projector array in roughly then seconds. In this work, a camera was first registered to the display's reference frame manually. Each projector, P_k, projected a checkerboard pattern that was observed by the camera. Corners on the checkerboard pattern where determined in the camera's image plane, establishing the point correspondences between the camera an the projector. From this information, projector-to-camera homographies, ${}_R\boldsymbol{C}_k$, were computed. Next, the necessary ${}_R\boldsymbol{P}_k$ could be computed using Eq. 1. This allowed the projected imagery to be correctly aligned to the display reference frame. Raskar et al. reported that this approach could align the projected imagery with sub-pixel accuracy [39].

While this approach is effective, the use of a single camera limits the scalability of this technique to displays composed of a larger number of projectors. Such large displays made of $40 - 50$ projectors are used in many national institutes like Sandia, Lawrence Livermore National Laboratories and the National Center

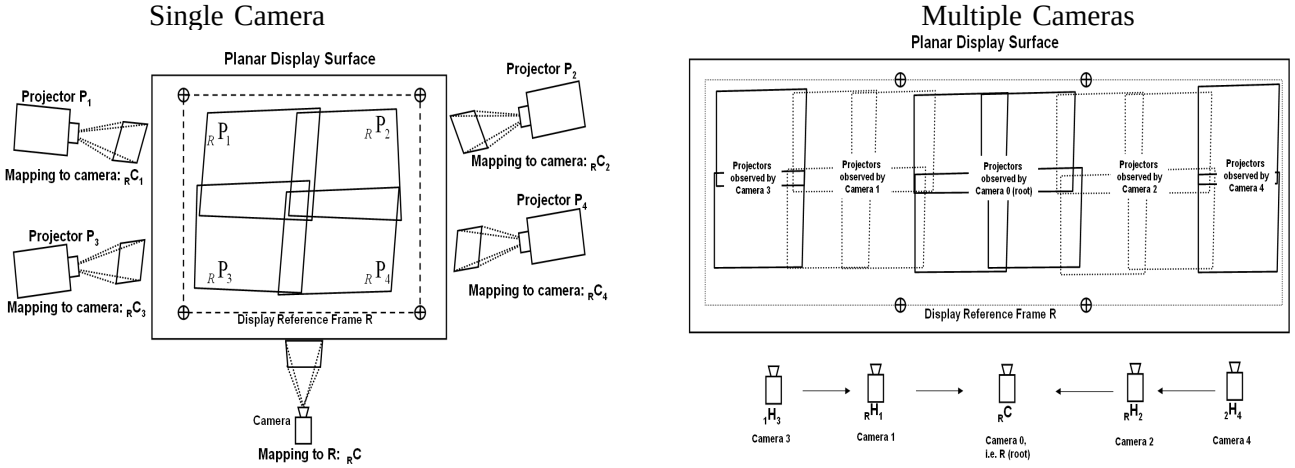

Fig. 3. (Left) 3×3 linear homographies are computed that relate the projectors to the display reference frame, R. A camera is used to observe projected imagery from each projector. (Right) For large field-of-view projector arrays, multiple cameras are used. Each camera observes a region of the display. Projector-to-camera homographies concatenated with camera-to-reference frame homographies are used to compute the necessary projector-to-reference frame mapping.

for Supercomputing Applications (NCSA) at University of Illinois at Champagne Urbana (UIUC).

To address this scalability issue, Y. Chen et al. [10] proposed a method that used a single camera mounted on a pan-tilt unit (PTU). The PTU allowed the camera to move such that it could see a very large field-of-view. Controlled by a PC, the camera could be automatically moved to observe points and lines projected by the individual projectors (the experiment in [10] registered eight projectors). The camera could relate points and lines from each projector to the display's global reference frame, R. A rough alignment of projectors was assumed. This meant that projected points and lines between projectors should be aligned, but were not because of slight mis-registrations. Using the collected data from the camera, a simulated annealing algorithm was used to compute each ${}_R\boldsymbol{P}_k$, such that errors between the corresponding projector points and angles between the corresponding lines where minimized. Y. Chen et al. reported that this approach could achieve near pixel accuracy in projector alignment [10]. While this approach proved to work, it suffered from being slow. Overall time to collect data from the PTU-mounted camera and to perform the simulated annealing was reported to be around one hour. Implementation improvements could undoubtedly reduce the overall time.

2) Using Multiple Cameras: More recently, H. Chen et al. [9] proposed a more scalable approach that uses multiple cameras. Several cameras observing the display are related to one another by camera-to-camera homographies. A root camera, ${}_R\boldsymbol{C}$, is established as the reference frame. Adjacent cameras, i and j, are related to one another by the ${}_i\boldsymbol{H}_j$ homographies. Point correspondences are established between adjacent cameras by observing projected points from which the ${}_i\boldsymbol{H}_j$s are computed. Next, each camera is registered to the root camera and thus to the reference frame, R. This can be done by computing a homography ${}_RH_j$, which is constructed by concatenating adjacent camera-to-camera homographies until the root camera is reached as follows (see Figure 3 (right)):

$$ {}_R\boldsymbol{H}_j = {}_R\boldsymbol{C} \times {}_R\boldsymbol{H}_i \times \cdots \times {}_i\boldsymbol{H}_j. \tag{2} $$

The homography ${}_R\boldsymbol{H}_j$ maps points in camera, j, to the reference frame R. To determine the path of this camera-to-reference frame concatenation, a minimum-spanning "homography tree" is built that minimizes registration errors in the camera-to-camera reference frame [9].

Each projector is now observed by one of the cameras in the system. A single camera in the system can typically only observe only 2-4 projectors from the entire display wall. The projectors can be related to their corresponding cameras via a homography, denoted AS, ${}_k\boldsymbol{C}_j$, where j is the camera index and k is the projector index. Using the homography tree computed between the cameras, the projector-to-reference homography, ${}_R\boldsymbol{P}_k$, for a given projector, k, can be computed as:

$$ {}_R\boldsymbol{P}_k = {}_R\boldsymbol{H}_j \times {}_j\boldsymbol{C}_k, \tag{3} $$

where ${}_R\mathbf{H}_j$ has been constructed using Eq. 2. Experiments in [9] showed that this approach can be very accurate in registering projectors. In examples using up to 32 projectors, sub-pixel accuracies could be achieved. In simulation, this technique was shown to be scalable to scores of projectors and cameras. In addition, this

approach took only a few minutes to reach a solution with large numbers of projectors.

B. Arbitrary Display Surfaces

While the previously mentioned approaches are able to achieve geometric registration via homographies, these approaches are applicable only if the display surface is planar. Here, we discuss approaches that address non-planar display surfaces. These include surround environments such as video domes and immersive environments. In addition, these techniques are geared for very flexible deployment in existing environments, like an office, where large empty planar display surfaces may be difficult to find.

The approaches we present address two modes of operation. One that assumes a stationary viewer and one that allows a moving user (i.e., a head-tracked viewer). These techniques can of course be applied to planar display surfaces as well.

1) Stationary Viewer: Raskar [41] and Surati [42] proposed a registration algorithm that uses a two-pass rendering technique to create seamless imagery on arbitrary display surfaces. In this approach, a single camera is placed at the location from where the viewer is supposed to observe the displayed imagery. A set of equally spaced features are projected from each projector, P_k, and registered in the camera image plane. The projected features $P_k(x, y)$ are typically used to form a tessellated grid in the projector space as well as the camera image space (see Figure 4). This establishes a non-linear mapping from the projector's features, $P_k(x, y)$, to their positions in the camera's image plane, $C(u, v)$, denoted as $C(u, v) \mapsto P_k(x, y)$.

To correct the displayed imagery, a two-pass rendering algorithm is used. In the first pass, the desired image to be seen by the viewer is rendered. This desired image from the first pass is then warped to the projected image based on the $C(u, v) \mapsto P_k(x, y)$ non-linear mapping. This warp is non-linear and can be realized by a piecewise texturing between tessellated meshes in the projector and the camera image space. This warping constitutes the second rendering pass. For clarity, Figure 4 shows this procedure using only one projector. This technique will, however, produce a seamless image even when multiple overlapping projectors are observed by the camera. Note that any camera distortion (such as radial distortion) will be encoded in the $C(u, v) \mapsto P_k(x, y)$ mapping. For cameras with severe radial distortion, e.g. a camera using a fish-eye lens, this distortion will be noticeable in the resulting image created by the projector mosaic. Care should be taken to first calibrate the camera to remove such distortion. Routines to perform this calibration are typically of computer vision software packages, such as Intel's OpenCV [19].

The warp specified from the $C(u, v) \mapsto P_k(x, y)$ mapping generates a geometrically correct view from where the camera is positioned. For this reason, the camera is positioned close to where the viewer will be while viewing the display. However, as the viewer moves away from this position, the imagery will begin to appear distorted. Thus, this technique is suitable for a stationary viewer only.

Yang et al. [50] incorporated this two-pass rendering algorithm into the University of North Carolina at Chapel Hill's *PixelFlex* display system. In the PixelFlex system, mirrors are mounted on pan-tilt units (PTU) positioned in front of the projectors. Software allows a user to dynamically modify the projectors' spatial alignment by moving the mirrors via the PTUs. New projector configurations are registered using the technique described above. Brown et al. [5] incorporated the same technique into the WireGL and Chromium [18] rendering architecture. This allows users to deploy PC-based tiled display systems that support unmodified OpenGL applications. Both Brown et al. [5] and Yang et al. [50] reported sub-pixel projector registration when using this two-pass approach. In addition, these approaches can register the displays in a matter of minutes. Further, the non-linear warp corrects for both non-linear projector lens distortion and display surface distortion. Thus, this approach allows for very flexible display configurations (non-rectangular projector arrangements on non-planar display surfaces). However, the need for two rendering passes can affect performance. Brown et al. [5] reported a drop in performance from 60 fps to 30 fps, when the second-pass warp was used on a 2×2 projector array using four PCs with nVidia GeForce3 cards. This overhead may be alleviated in the future by having the second-pass conducted directly on the projectors (see Section V). In addition, this technique uses a single camera, limiting its scalability to large projector arrays.

2) A Moving (Head-Tracked) Viewer: For a moving viewer in an arbitrary display environment, the necessary warping function between each projector and the desired image must be dynamically computed as the view changes. Raskar et al. [40] presented an elegant two-pass rendering algorithm to address this situation. Figure 5(a) illustrates this two-pass rendering approach. The desired image from the viewer's position is rendered in the first pass. This image is then projected from the viewer's point of view onto a 3D model of the display surface using projective textures. This textured 3D model is then rendered from the view point of the projector as the

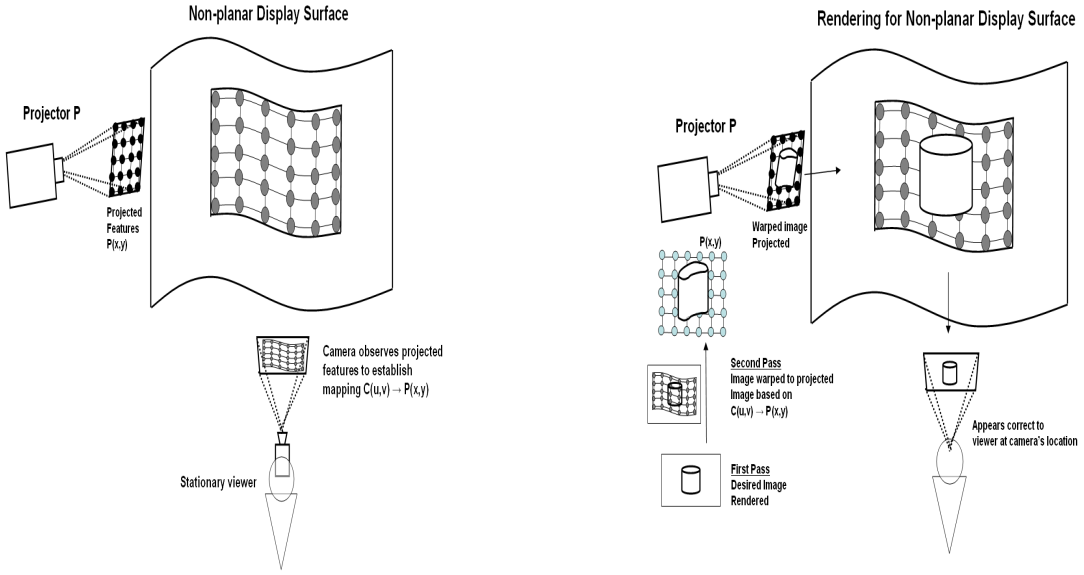

Fig. 4. (Left) Projectors display features that are observed by a camera placed near the desired viewing location. (Right) The desired image is (1) rendered and then (2) warped to the projected imagery based on its mapping to the camera.

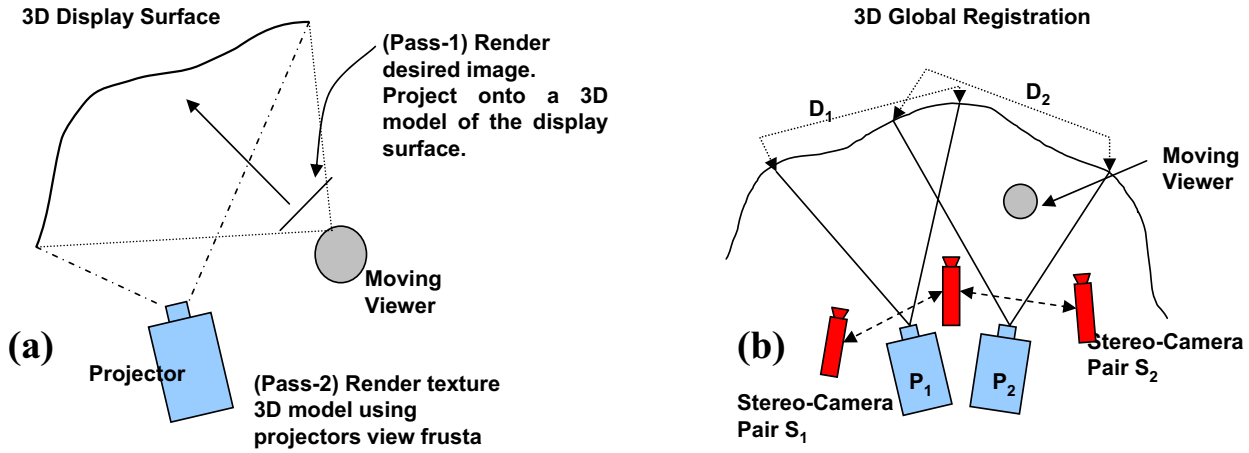

Fig. 5. a) A two-pass rendering algorithm for a moving viewer and an arbitrary display surface. The first pass renders the desired image to be observed by the user. This is used as a projective texture and projected from the viewer's point of view onto the display surface. The textured display surface is then rendered from the projector's point of view constituting the second pass render. When projected, second pass rendered image will look correct to the viewer. (b) Stereo-camera pairs are used to determine the 3D display surfaces D_1 and D_2 and projector locations P_1 and P_2. These are then registered into a common coordinate system along with the head tracker.

second rendering pass. When projected by the projector, this second pass image will appear geometrically correct to viewer.

In this algorithm, three components must be known: (1) a 3D model of the display surface, (2) the projectors' locations (in the form of a view frustum with respect to the display surface) and (3) the viewer's location (with respect to the display surface). These three components need to be registered in a common coordinate frame for the algorithm to work. Raskar et al. [37] presented a system that uses several cameras to determine automatically the 3D geometry of the display surface and the location of the projectors within the display environment. The data is then integrated with a head tracker to provide the third necessary component of viewer location. Figure 5(b) shows an overview of the approach.

In this system, multiple cameras are first used to form several stereo-camera pairs, S_i, to observe the projected imagery. Typically, one stereo pair is established for each projector. Each stereo pair is calibrated using a large calibration pattern. Note that a particular camera may be a member of more than one stereo pair. Using the stereo pair S_i, the display surface, D_i, seen by the

projector, P_i, can be determined using a structured-light technique. Each recovered 3D display surface, D_i, is represented as a 3D mesh. From D_i, the projector's P_i' view frustum (i.e., 3D location) with respect to the display surface can be computed. This completes the computation of the initial unknowns of the 3-D display surface and the projector location for every projector in the display. However, each D_i and P_i pair is still registered to different coordinate frames. The next step is to unify them within a common coordinate frame.

A stereo pair, S_i, can see its corresponding projector, P_i, and any overlapping portion from an adjacent projector, P_j. Using 3-D point correspondences acquired in the overlapping region between two display surfaces, D_i and D_j, a rigid transformation consisting of a rotation, ${}_iR_j$, and a translation, ${}_iT_j$, can be computed to bring D_i and D_j into alignment. Once the display surfaces are in alignment view frustums, P_i and P_j, can also be computed in the same common coordinate frame. Finally, the head tracker is registered to the global coordinate frame. This can be done by registering tracker positions to 3D points on the displays surface. A rotation and translation to bring the tracker's coordinates into alignment with the display surface can be computed. With all three necessary components registered to a common coordinate frame, the two-pass rendering algorithm for a moving user can be used to generate seamless imagery.

This approach allows flexible projector alignment and display surfaces and, hence, recovers the underlying display surface and projector locations automatically. In addition, this technique is scalable, allowing immersive displays to be deployed in a wide-range of environments. However, due to the large number of parameters that need to be estimated (e.g. camera parameters,3D surface, projector parameters) the accuracy of this system is roughly 1-2 pixels. In addition, this technique is fairly slow, requiring around 30 minutes to register a six-projector system.

More recently, Raskar et al. [38] presented a more efficient method to solve for a smaller set of 3D surfaces, specifically quadratic surfaces. Examples of quadric surfaces are domes, cylindrical screens, ellipsoids, and paraboloids. Such specialized surfaces are used in systems built for training and simulation purposes. For quadratic surfaces, the warping to register the images on quadric surfaces can be expressed by a parameterized transfer equation. In comparison with the full 3D reconstruction approach in [37], this parameterized approach has substantially fewer unknown parameters to estimate. A registration accuracy of one pixel was reported for this method [38]. However, note that at rendering time the parameterized transfer equation is discretized into a densely tesselated mesh. Some of the accuracy may be lost during this process.

III. Photometric Correction

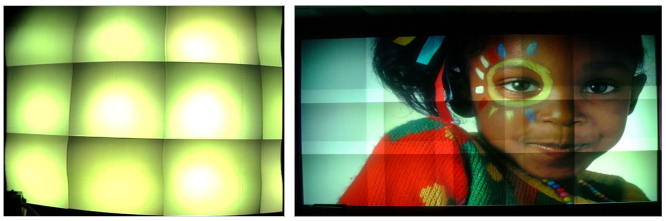

Fig. 6. Digital photographs of tiled displays showing the color variation problem. (a): Example of severe photometric variation across a display made of abutting projectors. Though difficult to believe, it is true that every pixel of this display is projecting the identical input of the maximum intensity for green. (b): A tiled display made of a 3×5 array of fifteen projectors ($10' \times 8'$ in size) with perfect geometric registration, but with color variation.

In this section, we address the color variation problem. Current commodity projectors, our target products for building large-area displays inexpensively, do not have sophisticated lens systems to assure color uniformity across the projector's field-of-view. Thus, the color variation in multi-projector displays made of commodity projectors, can be significant. Figure 6(a) shows the severe color variation of the 40 projector display used by NCSA of UIUC. In addition, even after perfect geometric alignment, the color variation problem can be the sole factor in causing a 'break' in creating the illusion of the single large display, as shown in Figure 6(b). Thus, color variation problems need to be addressed to achieve truly seamless displays.

Color is a three-dimensional quantity defined by one-dimensional luminance (defining brightness) and two-dimensional chrominance (defining hue and saturation). The entire range of luminance and chrominance that can be reproduced by a display is represented by a 3D volume called the *color gamut* of the display. Since color is defined by luminance and chrominance, the color variation problem involves spatial variation in both luminance and chrominance. It has been shown that most current tiled displays composed of projectors of the *same manufacturer model* show large spatial variation in luminance while the chrominance is almost constant spatially [29], [24]. Also, humans are at least an order of magnitude more sensitive to luminance variation than to chrominance variation. For example, humans have higher spatial and temporal frequency acuity for luminance than for chrominance; or, humans can resolve a higher luminance resolution than chrominance resolution. Detailed

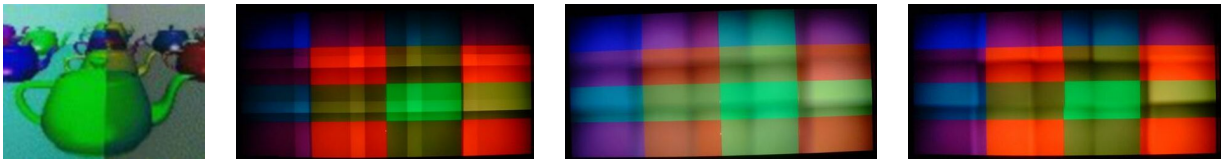

Fig. 7. From left: (1) Correction done by luminance matching for a display made of two abutting projectors; (2), (3), and (4) respectively: fifteen-projector tiled display, before blending, after software blending, and after optical blending using physical mask, respectively.

discussion of such perceptual capabilities can be found in the psychophysics literature [11], [15], [48]. These show that perceptually, the subproblem of *photometric variation* (luminance variation) is the most significant contributor to the color variation problem.

The color variation in multi-projector displays has been classified in three different categories [29].

1) *Intra-Projector Variation:* Luminance varies significantly across the field-of-view of a single projector. Luminance fall-off of about $40 - 80\%$ from the center to the fringe is common in most commodity projectors. This is caused by several reasons like the distance attenuation of light and the angle at which the light from the projector falls on the screen. This also results in an asymmetric fall-off, especially with off-axis projection. The non-Lambertian nature of the screen further pronounces the problem. There are many front projection screens available that are close to Lambertian in nature. However, this is a rare property among the rear projection screens making intra-projector variation more pronounced for rear-projection systems. However, the chrominance remains almost spatially constant within a single projector.
2) *Inter-Projector Variation:* Luminance can vary significantly across different projectors. This is caused by differences in the properties of the projector lamps and their ages, by differences in the position and orientation of the projector with respect to the screen, and also by differences in the projector settings like brightness, contrast and zoom. However, chrominance variation across projectors is relatively much smaller and is almost negligible for same model projectors.
3) *Overlap Variation:* The luminance in the region where multiple projectors, overlap is multiplied by the number of overlapping projectors creating a very high brightness region. If the chrominance properties of the overlapping projectors are not close to each other as in video walls made of different model projectors, this can also lead to visible chrominance variations in the overlapped region. However, these variations are at least an order of magnitude smaller than the luminance variation.

A related problem is that of the *black offset*. Any ideal display device should project no light for the red, green, and blue (RGB) input $(0, 0, 0)$. This is true for most cathode ray tube (CRT) projectors since the electron beam inside can be switched off completely at zero. However, most commodity projectors use light blocking technology like lithium crystal display (LCD) or digital light projector (DLP), through which some light is always projected. This is called the *black offset*. This reduces the contrast of projectors and current technology is driven towards reducing this black offset as much as possible.

In abutting projector displays, traditionally, the color compensation was done by manipulating the projector controls (like brightness, contrast and zoom) manually using feedback from a human user on the quality of the color uniformity achieved. Unfortunately, this is a very labor-intensive process and ideal color uniformity is not always achievable given the limited control allowed to the user. Therefore, automatic color calibration methods were devised to create scalable displays.

A. Gamut Matching

This approach was the first to automate the process of using manual feedback and manipulation of controls. A point light measuring instrument (like a spectroradiometer) [44], [45] is used to measure the *color gamut* of each projector at one spatial location. The spatial variation of color within a single projector is assumed to be negligible and the colors between the different projectors are matched by a two-step process. First, a *common color gamut* is identified that is the intersection of the gamuts of different projectors. This represents the range of colors that all the projectors in the display are capable of producing. Second, by assuming projectors to be linear devices, 3D linear transformations are used to convert the color gamut of each display to the common color gamut.

This method is applicable to devices that use three primary colors (most devices use red, green and blue as color primaries). Since three primary colors form a basis

for describing all colors, a color in the three-primary-system can be represented by a *unique* combination of the three primaries. However, some DLP projectors use a clear filter to project the grays instead of projecting the superposition of light from the red, green and blue filters. This makes these DLP projectors behave like four primary devices (like printers that use cyan, magenta, yellow and black as primaries). Adding the fourth primary brings in linear dependencies in such systems. As a result, a color cannot be represented using unique combinations of the four primaries. The gamut-matching method depends on the linear independence of the primaries and becomes inapplicable in such four-primary systems. [49] presents a solution by which a gamut-matching method can be extended to be applied to the DLP projectors.

The theoretical disadvantage of the gamut-matching method lies in the fact that there is no practical method to find the common color gamut. [4] presents an optimal method (where n is the number of color gamuts whose intersection is being sought), which finds this intersection in $O(n^6)$ time. This is clearly not a scalable solution, especially for large-scale displays of over 10 projectors. [26] tries to address this problem by matching only luminance across different projectors. Since most display walls are made of the same model projectors, which differ negligibly in chrominance, achieving luminance matching across different projectors can suffice. The result of this method is shown in Figure 7. However, since spatial color variation is ignored, these methods cannot produce entirely seamless displays. Further, expensive instrumentation makes these methods cost prohibitive. A relatively inexpensive radiometer costs at least four times more than a projector. Expensive radiometers can cost as much as a dozen projectors.

B. Using a Common Lamp

Using a common lamp is a wonderful engineering feat [33]. In this method, the lamps of the multiple projectors are taken off and replaced by a common lamp of much higher power. Light is distributed from this common lamp to all the different projectors using optical fibres. However, this method is cost intensive because it requires skilled labor. Further, power and thermal issues (heat generated by the high-power lamp) make this approach unscalable. A maximum of nine projectors can be illuminated by a common lamp using this method. Also, this approach addresses only the color variation caused by the differences in the lamp properties. All the other kinds of variations still exist.

C. Blending

Blending or feathering techniques, adopted from image-mosiacing techniques, address overlapped regions and try to smooth color transitions across these regions. The smooth transitions can be achieved by using a linear or cosine ramp which attenuate pixel intensities in the overlapped region. For example, considering a pixel x in the overlap region of projectors P_1 and P_2, as illustrated in Figure 8(Left). Let the contributions of these projectors at x be given by $P_1(x)$ and $P_2(x)$ respectively. When using linear ramping the intensity at x is computed by a linear combination of the intensities $P_1(x)$ and $P_2(x)$, i.e,

$$\alpha_1(x)P_1(x) + \alpha_2(x)P_2(x)$$

where $\alpha_1 + \alpha_2 = 1$. These weights, α_1 and α_2, are chosen based on the distance of x from the boundaries of the overlapped region. For example, when using a linear ramp, these functions can be chosen as,

$$\alpha_1(x) = \frac{d1}{d1 + d2}; \quad \alpha_2(x) = \frac{d2}{d1 + d2}.$$

.

This two-projector example can be extended to an arbitrary number of projectors [37]. To do so, the hull, H_i, in the camera's image plane of observed projector P_i's pixels is computed. The alpha-weight, $A_m(x)$, associated with projector, P_m's pixel, x, is evaluated as follows:

$$A_m(x) = \frac{\alpha_m(m,x)}{\sum_i \alpha_i(m,x)}, \tag{4}$$

where $\alpha_i(m,x) = w_i(m,x) * d_i(m,x)$ and i is the index of the projectors observed by the camera (including projector m).

In the above equation, $w_i(m,x) = 1$ if the camera's observed pixel of projector P_m's pixel, x, is inside the convex hull, H_i; otherwise, $w_i(m,x) = 0$. The term $d_i(m,x)$ is the distance of the camera's observed pixel of projector P_m's pixel, x, to the nearest edge of H_i. Figure 8 (right) shows the alphamasks created for four overlapping projectors. The alphamasks are applied after the image has been warped. This can be performed efficiently as a single alpha-channel textured quad the size of the framebuffer.

Blending can be achieved in three ways. First, it can be done in software [40] where the distances from the projector boundaries and the number of projectors contributing to every pixel in the overlap region can be accurately calculated using geometric calibration information. Thus, the ramps can be precisely controlled by software. However, this cannot attenuate the black offset, which is especially important with scientific or astronomy data,

Fig. 8. Blending Techniques: (Left) The intensity at any pixel x in the overlapped region of two projectors, P_1 and P_2, is the combination of the intensity of P_1 and P_2 at x. (Right) The resulting alphamasks computed for four projectors.

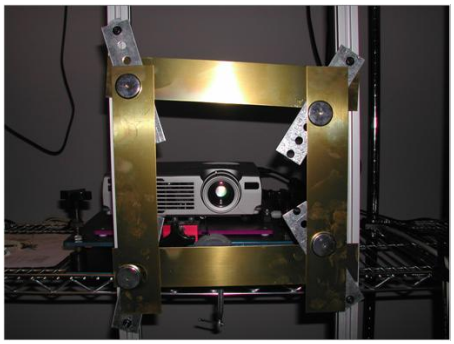

Fig. 9. Aperture blending by mounting metal masks on the optical path of the projector that attenuates the light physically.

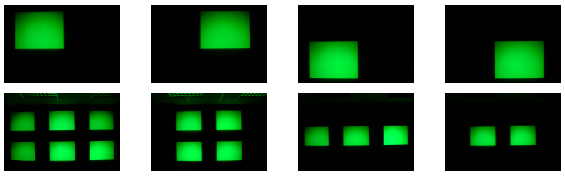

Fig. 10. To compute the display luminance surface of each projector, we need only four pictures per channel. Top: Pictures taken for a display made of a 2×2 array of 4 projectors. Bottom: The pictures taken for a display made of a 3×5 array of 15 projectors (both for green channel).

which often have black backgrounds. Alternate optical methods thus try to achieve this blending by physical attenuation of lights so that it can also affect the black offset. In one method, physical masks mounted at the projector boundaries on the optical path attenuate the light in the overlapped region [23], as shown in Figure 9. In another method, optical masks are inserted in front of the projection lens to achieve the attenuation [8]. The results of blending are shown in Figure 7. Though blending methods are automated and scalable, they ignore the inter- and intra-projector spatial color variation. Also, the variation in the overlapped region is not accurately estimated. Thus, blending works well if the brightness of the projectors whose overlap is being blended have similar luminance ranges which is often assured by an initial manual brightness adjustment using the projector controls. However, for displays where the luminance has a large spatial variation (like for most rear projection systems), blending results in softening of the seams in the overlapping region, rather then removing them.

D. Camera-based Photometric Uniformity

All the methods mentioned so far address only the inter-projector or overlapped variation. None addresses the intra-projector variation that can be significant. Also, only the gamut matching method makes an effort to estimate the color response of the projectors. However, since the spatial variation in color is significant, a high-resolution estimation of the color response is the only means towards an accurate solution. Thus, the use of a camera is inevitable. However, since a camera has a limited color gamut (as opposed to a spectroradiometer), estimating the color gamut of the display at an high resolution is difficult. However, different exposure settings of a camera can be used to measure the luminance accurately and faithfully. Exploiting this fact, [28], [29] use a camera to correct for the photometric variation (variation in luminance) across a multi-projector display. Since most current displays use the same model projectors that have similar chrominance properties, this method achieves reasonable seamlessness.

This camera-based method aims at achieving *identical* photometric response at every display pixel. This is called *photometric uniformity*. We describe the method for a single channel. All three channels are treated similarly and independently. The method comprises two steps. The first step is a one-time *calibration* step that uses the camera to estimate the luminance response of the multi-projector display. At the end of this step, a per-projector per-pixel map called the *luminance attenuation map (LAM)* is generated. In the *image correction* step,

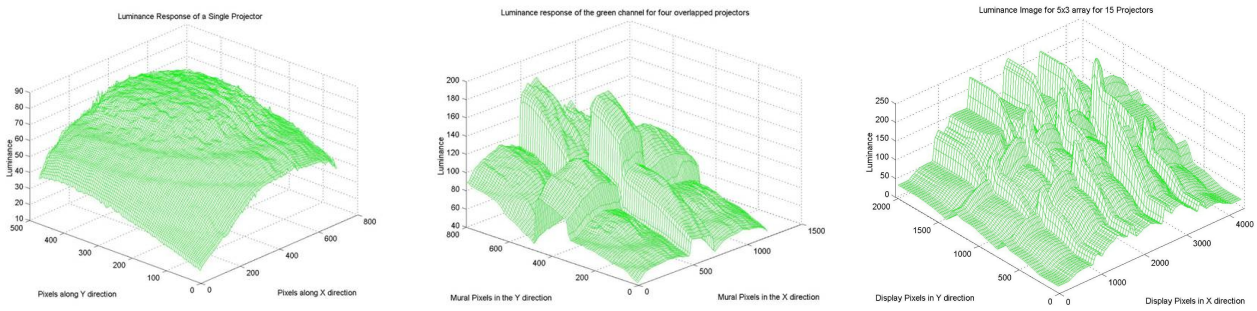

Fig. 11. Left: The luminance surface for one projector. Middle and Right: The display luminance surface for a 2×2 array of a four projector and a 3×5 array of fifteen projectors, respectively. (all for the green channel)

the LAM is used to correct any image to be displayed.

1) Calibration: Let the display, D, be made of N projectors, each denoted by P_j. Let the camera used for calibration be C. First, geometric calibration is performed to find the geometric warps, $T_{P_j \to C}$, relating the projector coordinates (x_j, y_j) with the camera coordinates (x_c, y_c), and $T_{P_j \to D}$, relating the projector coordinates with the global display coordinate (x_d, y_d). Any of the geometric calibration methods described in Section II can be used for this purpose. Photometric calibration has three steps.

Capturing the Display Luminance Response: Note that the camera should be in the same location as the geometric calibration method throughout this process of photometric calibration. Using the digital camera, two functions are acquired to perform photometric calibration.

The variation of the projected intensity from a channel of a projector with the variation in the input is defined by the *intensity transfer function (ITF)*. This is commonly called gamma function. In projectors, this function cannot be expressed by a power function and hence we prefer to call it the intensity transfer function. [24] shows this function to be spatially invariant i.e varies only with input and does not change from one pixel to another within the projector. Hence, the ITF for each projector is first estimated using a point light measuring instrument like a photometer at one location for each projector. However, since such instruments can be cost prohibitive, [35] presents a method in which the high dynamic range (HDR) imaging method developed by Debevec and Malik [13] is applied to measure the ITF of all the projectors at a time using a inexpensive video camera.

Next, the display luminance surface is captured. Images with maximum luminance are projected from each projector and captured using the camera. More than one non-overlapping projector can be captured in the same image. The images taken for this purpose for a four and fifteen projector display are shown are in Figure 10. From these images the luminance surface for each projector, L_{P_j}, is generated by using the warp $T_{P_j \to C}$. Standard RGB to YC_rC_b conversion is used for this purpose. The luminance surface from these projectors are then added up spatially using the warp, $T_{P_j \to D}$, to create the display luminance surface, L_D. The luminance surfaces generated for a projector and the whole display are shown in Figure 11.

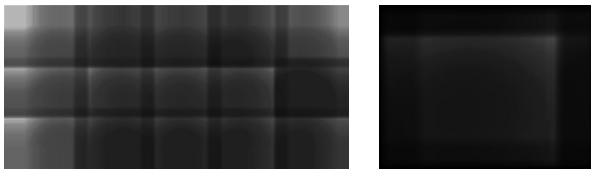

Fig. 12. Left: The display luminance attenuation map for a 3×5 array of fifteen projectors. Right: The LAM for a single projector cut out from the display LAM on the left. (both examples are for the green channel)

Finding the Common Achievable Response: Next, the luminance response that can be achieved at every pixel of the display is identified. Since the dimmer pixels cannot match the brighter pixel, the common achievable response is given by

$$L_{min} = \min_{\forall (x_d, y_d)} L_D.$$

Generating the Attenuation Maps: The luminance attenuation map (LAM), A_D, for the whole display is first generated by

$$A_D(x_d, y_d) = \frac{L_{min}}{L_D(x_d, y_d)}.$$

From this display LAM, a luminance attenuation map for each projector is generated using the inverse of the warp $T_{P_j \to D}$. The display and projector LAMs thus generated are shown in Figure 12. This concludes the calibration.

Fig. 13. The top row shows the image before correction and the bottom row shows the image after luminance matching. Left and middle: Digital photograph of a 2×2 array of projectors. Right: Digital photograph of a 5×3 array of projectors. In this case, the image after correction was taken at a higher exposure.

2) Image Correction: Once the per-projector LAMs are generated, the per-projector image correction is done in two steps. These correction steps are applied to any image that is projected from the display. First, the per pixel multiplication of the image with the LAM is performed. This multiplication assumes linear ITF. In practice, however, the ITF is non-linear. To compensate for that, an inverse of the ITF is applied to the image after the LAM is applied. The results of this method are shown in Figure 13.

The corrections required to achieved photometric uniformity or to compensate for the surface reflectance are encoded as per-pixel linear operations and a 1D color look-up-table (LUT). These form a very efficient way of representing the non-linear correction because these operations can be applied in real time using commodity graphics hardware. Recent advances in programmable graphics hardware make it possible to implement complex per-pixel operations that can run natively on the graphics processor without taking a toll on the main CPU [31], [2]. Details of how these can be used to create interactive displays are available at [27].

However, since this method aims at photometric uniformity, the photometric response of every pixel is matched to the 'worst' pixel on the display, ignoring all the 'good' pixels that are very much in the majority. This results in a compression in dynamic range making the method unscalable. Ongoing research [25] is trying to address this issue by achieving a perceptual uniformity rather than a strict phorometric uniformity.

E. Camera-Based Compensation for Non-White Surfaces

The methods described so far can be used to compensate for color variation in a multi-projector display when projecting on a white screen. Recent work addresses the issue of using projectors to project on displays that are not necessarily white but have colors and textures, like brick walls or a poster boards [30], for scenarios where it may not be possible to find white display surfaces. In this approach, the camera and projectors are assumed to be linear devices and the color transformation between them is expressed by a 3×3 matrix, V. The RGB color, C, measured by a camera for a projector input, P, is related by the matrix multiplication $C = VP$.

The camera is first used to measure the response of several images projected from the projector. Each image is made of identical input at every projector pixel. With the projector pixel inputs and the corresponding measured outputs from the camera established, V can be estimated *for each pixel* by solving a set of over-determined linear equations. Once V is estimated, V^{-1} is applied to the input image to generate the desired response that would look seamless on an imperfect surface. The estimated V is further refined by a continuous feedback and an estimation loop between the projector and the camera. The non-linearities of the projector and the camera are also considered to validate the assumption of linear devices. Greater details of this method are available in [30]. This method has not yet been scaled to displays made of multiple projectors.

IV. Discussion

All of the approaches discussed in the previous sections have been used and tested in deploying various projector-based display systems. Table I provides a list of representative systems and supporting publication references. Table I lists these systems in chronological order based on publication date and itemizes key aspects of these systems, including the types of display surface, the number of cameras and projectors in the system, geometric and photometric registration approach used, the number of rendering passes required, and targeted viewer mode (stationary vs. moving). Since a variety of approaches are available for different applications and display configurations, we use this section to discuss the positive and negative aspects of the various approaches for geometric and photometric registration.

On the geometric front, restricting the display surface to be planar has many benefits. First, there are more scalable techniques to register very large arrays with sub-pixel accuracy, such as the homography tree approach [9]. In addition, the alignment procedure using a 2D linear homography can be performed in the graphics pipeline, allowing for efficient rendering [36], [50]. Planar homography-based approaches, however, can correct for only linear geometric distortions. For instance, non-linear radial distortion introduced by a projector's optical system cannot be corrected by this method. Yang et al. [50] showed that the zoom setting of some projectors effected the radial distortion. This limited the projectors usable zoom range to the positions that minimized radial distortion.

The parameterized transfer equation introduced by Raskar et al [38] extends planar surface algorithms to quadric surfaces. While some screens (i.e., dome and cylindrical screens) can be modelled as quadric surfaces, this requires precise manufacturing. For applications that used cheaper constructed surfaces that do not require head-tracking, it may still be better to use the direct mapping technique (see Section II-B.1) that can compensate for the imperfections in the display surface geometry.

For arbitrary display surfaces, the direct mapping from the camera space to the projector space is a very efficient way to generate seamless images from one fixed view location. The resulting two-pass rendering algorithm compensates for display surface distortion as well as projector lens distortion. For small arrays (4-5 projectors), this approach is very flexible and can allow quick deployment of projector-based displays in a wide range of environments. However, because this technique requires the camera to see the entire display and it is not scalable to large projector arrays.

The technique presented by Raskar et. al [37] for a moving user and arbitrary display surfaces involves a full 3D modeling of the display environment including the projector positions and display surface geometry. While this approach is the most general solution to large scale display deployment, it is non-trivial to implement a robust and practical system. Due to its complexity, the best registration error reported so far is about 1-2 pixels.

For correcting the color variation problem, solutions like blending (Section III-C) do not estimate the spatial variation and hence cannot achieve entirely seamless display, especially for large displays. However, for small systems of $2-4$ projectors, blending can achieve effective results if it is preceded by color balancing across different projectors. This color balancing can be manual or can be automated using gamut matching techniques (Section III-A). The camera based technique (Section III-D) can achieve reasonable photometric seamlessness across the display, which is sufficient for displays made of same brand projectors. The advantage of this method lies in its complete automation and scalability. However, the limitation of both gamut matching or photometric uniformity lies in degrading the color quality of the display in terms of dynamic range and color resolution. Thus, achieving perceptual uniformity (rather than strict uniformity) while maintaining high display quality is the current area of research. Finally, current camera-based correction do not address chrominance variation, arbitrary display geometry or a moving user, which are still active areas of research.

V. Hardware Support for Image Correction

To correct geometric and photometric distortions in a projector-based display requires changes to be made to the desired image, which causes overhead during rendering time. This shortcoming has been ameliorated by recent advances in computer hardware. Modern graphics hardware provides a tremendous amount of image processing power. Thus, many of the correction operations can be off-loaded to the graphics board. The overhead to warp and blend a screen resolution image becomes negligible. In addition, certain correction operations can be integrated into the graphics rendering pipeline, such as the one-pass rendering algorithm for off-axis projection on planar surfaces [36]. These approaches completely eliminate the image correction overhead when rendering 3D contents. With increasing programmability in graphics hardware, we expect that new techniques that leverage the power of programmable graphics hardware will emerge to reduce the rendering overhead in a wide range of configurations.

TABLE I
CHARACTERISTICS OF REPRESENTATIVE LARGE-FORMAT DISPLAYS USING CAMERA-BASED CALIBRATION

System	Display surfaces	number of projectors	number of cameras	Resolution (mega pixels)	Geometric registration	Photometric correction	Rendering passes
Surati [42]	arbitrary$^{\heartsuit}$	4	one	1.9	fixed warping	color attenuation	two
Raskar et al. [37]	arbitrary$^{\diamondsuit}$	5	multiple	3.8	full 3D model	software blending	three
Y. Chen et al. [10]	planar	8	one on PTU	5.7	simulated annealing	optical blending	one
PixelFlex [50]	arbitrary$^{\heartsuit}$	8	one	6.3	fixed warping	software blending	two
H. Chen et al. [9]	planar	24	multiple	18	homography tree	optical blending	one
Metaverse [20]	multiple walls$^{\diamondsuit}$	14	one	11	homography	software blending	one
iLamp [38]	quadric surfaces	4	one/projector	3.1	full 3D model	software blending	two

$^{\diamondsuit}$ head-tracked moving viewer. $^{\heartsuit}$ static viewer (image is correct for a fixed location).

The need for more flexibility in projector-based displays is also being addressed by projector manufacturers. Recent projectors are equipped with more options to adjust the projected images. For example, projectors from EPSON provide automatic keystone correction using a built-in tilt sensor [43]. 3D-Perception CompactView was one of the first companies to offer a projector that performed real-time corrective warping to the incoming video stream [1]. This feature is used to help compensate for projection on smooth, curved surfaces, such as those in video domes. Recently, other projector manufacturers have provided similar options. The Barco Galaxy-WARP projector is also capable of performing real-time corrective warping to the incoming video stream [3]. Both products allow for non-linear image mapping. Thus, a wide range of configurations can be accommodated without incurring any rendering overhead.

Currently, these products allow control of the non-linear warping via user interfaces. However, it is a matter of time before an interface between the projector and camera-based registration techniques will allow this warping to be specified automatically. With the projectors performing the warping in real-time, performance overhead will not be an issue. This should be a great benefit to the current two-pass rendering algorithms. Merging this technology with camera-based registration will truly allow a new generation of flexible and highly configurable projector-based display environments.

VI. CONCLUSION AND FUTURE WORK

Camera-based calibration techniques have enabled a much wider range of configurations for projector-based displays. The capability of automatic geometric alignment and photometric correction of multiple projected images eases the setup and reduces the cost of large-format displays. Coupled with advances in distributed rendering software and graphics hardware, the possibility of creating inexpensive and versatile large format displays using off-the-shelf components becomes a reality. It is our hope that the information provided in this survey will provide projector-based display users a better a guide to the currently available techniques and their associated advantages and disadvantages.

Looking forward, there are a number of research topics that can further advance the state of the art.

a) Geometric Registration Quality: Registration quality is often reported as pixel registration accuracy in local overlapped regions and not in the context of the global display coordinate frame. Moreover, using a pixel as a unit of measure is ill-defined when imagery is projecting on arbitrary display surfaces or contributing projector pixels are not uniform in size. Better metrics and analysis approaches are needed to fully evaluate overall registration accuracy.

b) Color Correction: The shortcoming of the automated color correction method presented here is the severe degradation in image quality. Methods should be devised that optimize the available resources in terms of brightness and contrast of the display and achieve perceptual uniformity, which may not require strict photometric uniformity. Also, only the problem of spatial photometric variation has been addressed while assuming that most current displays have negligible chrominance variation. However, when using different model projectors, the chrominance variation cannot be ignored. Finally, arbitrary 3D display surfaces with arbitrary reflectance properties for moving users is still to be addressed.

c) Image Resampling: Most of the geometric correction techniques involved a resampling of an original rendered image. How this resampling affects the overall resolution of the display and ways to avoid fidelity loss need to be addressed.

d) Continuous Calibration: Almost all camera-based techniques treat the calibration procedure as a pre-processing routine. The correction function derived from the calibration information remains fixed until the next calibration. But during the normal operation of a display system, there are many factors that affect the validity of the calibration, such as vibrations, electronic drift, aging of projector light bulbs, or even transient events such as temporary occlusion of projector light. To deal with these problems, techniques could be developed to continuously monitor the projected imagery and correct any undesired distortions online. Promising work, such as continuous monitoring of display surfaces [51] and shadow removal [21], [46], [7] have demonstrated the potential of this research area.

e) Display and User Interaction: The real-time feedback of cameras in the display environment make it possible to develop interaction techniques between the user and the display. For example, laser pointer interaction inside a camera-registered displays can be easily realized [47], [6]. Significantly more ambitious goals have been set forth in UNC's Office of the Future project [40] and Gross et al's [16] *blue-c* system. These systems aim to provide immersive 3D telecommunication environments where cameras capture real-time 3D information of users inside the display. The tightly coupled relationship between the camera and display environment offers great potential for novel user interaction metaphors within such environments.

REFERENCES

[1] 3D Perception AS, Norway. CompactView X10, 2001. http://www.3d-perception.com/.

[2] ATI Technologies Inc. ATI Radeon 9800, 2003. http://www.ati.com/products/radeon9800.

[3] Barco, Kortrijk Belgium. Barco Galaxy Warp. http://www.barco.com/.

[4] M. Bern and D. Eppstein. Optimized color gamuts for tiled displays. *ACM Computing Research Repository, cs.CG/0212007, 19th ACM Symposium on Computational Geometry, San Diego*, 2003.

[5] M. S. Brown and W. B. Seales. A Practical and Flexible Tiled Display System. In *Proc of IEEE Pacific Graphics*, pages 194–203, 2002.

[6] M. S. Brown and W. Wong. Laser Pointer Interaction For Camera-Registered Multi-Projector Displays. In *Proceedings of the Proceedings of International on Image Processing (ICIP)*, Barcelona, Sept 2003.

[7] T.J. Cham, J. Rehg, R. Sukthankar, and G. Sukthankar. Shadow elimination and occluder light suppression for multi-projector displays. *Proceedings of Computer Vision and Pattern Recognition*, 2003.

[8] C. J. Chen and Mike Johnson. Fundamentals of Scalable High Resolution Seamlessly Tiled Projection System. *Proceedings of SPIE Projection Displays VII*, 4294:67–74, 2001.

[9] H. Chen, R. Sukthankar, and G. Wallace. Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. In *Proceeding of IEEE Visualization 2002*, pages 339–346, 2002.

[10] Y. Chen, D. Clark, A. Finkelstein, T. Housel, and K. Li. Automatic Alignment Of High-Resolution Multi-Projector Displays Using An Un-Calibrated Camera. In *Proceeding of IEEE Visualization 2000*, pages 125–130, 2000.

[11] R.A. Chorley and J. Laylock. Human Factor Consideration for the Interface between Electro-Optical Display and the Human Visual System. In *Displays*, volume 4, 1981.

[12] C. Cruz-Neira, D. Sandin, and T. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of SIGGRAPH 1993*, pages 135–142, 1993.

[13] P. E. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. *Proceedings of ACM Siggraph*, pages 369–378, 1997.

[14] Fakespace Systems Inc. PowerWall, 2000. http://www.fakespace.com.

[15] E. Bruce Goldstein. *Sensation and Perception*. Wadsworth Publishing Company, 2001.

[16] Markus Gross, Stephan Wuermlin, Martin Naef, Edouard Lamboray, Christian Spagno, Andreas Kunz, Esther Koller-Meier, Thomas Svoboda, Luc Van Gool, Silke Lang, Kai Strehlke, Andrew Vande Moere, and Oliver Staadt. blue-c: A Spatially Immersive Display and 3D Video Portal for Telepresence. In *Proceedings of SIGGRAPH 2003*, pages 819–827, San Diego, July 2003.

[17] G. Humphreys, I. Buck, M. Eldrige, and P. Hanrahan. Chromium: A Stream Processing Framework for Interactive Rendering on Clusters. In *Proceedings of SIGGRAPH*, July 2002.

[18] G. Humphreys, M Eldridge, Ian B., G Stoll, M Everett, and P Hanrahan. WireGL: A Scalable Graphics System for Clusters. In *Proceedings of SIGGRAPH 2001*, August 2001.

[19] Intel. Open Source Computer Vision Library (OpenCV). http://www.intel.com/research/mrl/research/opencv/.

[20] C. Jaynes, B. Seales, K. Calvert, Z. Fei, and J. Griffioen. The Metaverse - A Collection of Inexpensive, Self-configuring, Immersive Environments. In *Proceeding of 7th International Workshop on Immersive Projection Technology*, 2003.

[21] C. Jaynes, S. Webb, M. Steele, M. S. Brown, and B. Seales. Dynamic Shadow Removal from Front Projection Displays. In *Proceeding of IEEE Visualization 2001*, pages 174–181, San Diego, CA, 2001.

[22] K. Li, H. Chen, Y. Chen, D.W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J.P. Singh, G. Tzanetakis, and J. Zheng. Early Experiences and Challenges in Building and Using A Scalable Display Wall System. *IEEE Computer Graphics and Applications*, 20(4):671–680, 2000.

[23] K. Li and Y. Chen. Optical Blending for Multi-Projector Display Wall System. In *Proceedings of the 12 th Lasers and Electro-Optics Society 1999 Annual Meeting*, 1999.

[24] A. Majumder. Properties of Color Variation Across Multi-Projector Displays. *Proceedings of SID Eurodisplay*, 2002.

[25] A. Majumder. A Practical Framework to Achieve Perceptually Seamless Multi-Projector Displays, PhD Thesis. Technical report, University of North Carolina at Chapel Hill, 2003.

[26] A. Majumder, Z. He, H. Towles, and G. Welch. Achieving Color Uniformity Across Multi-Projector Displays. *Proceedings of IEEE Visualization*, 2000.

[27] A. Majumder, D. Jones, M. McCrory, M. E. Papka, and R. Stevens. Using a Camera to Capture and Correct Spatial

Photometric Variation in Multi-Projector Displays. *IEEE International Workshop on Projector-Camera Systems*, 2003.

[28] A. Majumder and R. Stevens. LAM: Luminance Attenuation Map for Photometric Uniformity in Projection Based Displays. *Proceedings of ACM Virtual Reality and Software Technology*, 2002.

[29] A. Majumder and R. Stevens. Color Nonuniformity in Projection-Based Displays: Analysis and Solutions. *IEEE Transactions on Visualization and Computer Graphics*, 10(2), 2003.

[30] S. K. Nayar, H. Peri, M. D. Grossberg, and P. N. Belhumeur. A Projection System with Radiometric Compensation for Screen Imperfections. *IEEE International Workshop on Projector-Camera Systems*, 2003.

[31] NVIDIA Corporation. GeForce FX, 2003. http://www.nvidia.com/page/fx desktop.html.

[32] T. Okatani and K. Deguchi. Autocalibration of a Projector-Screen-Camera System: Theory and Algorithm for Screen-toCamera Homography Estimation. In *Proceedings of Internation Conference on Computer Vision (ICCV)*, volume 2, pages 125–131, 2002.

[33] B. Pailthorpe, N. Bordes, W.P. Bleha, S. Reinsch, and J. Moreland. High-Resolution Display with Uniform Illumination. *Proceedings Asia Display IDW*, pages 1295–1298, 2001.

[34] Panoram Technologies Inc. PanoWalls, 1999. http://www.panoramtech.com/.

[35] A. Raij, G. Gill, A. Majumder, H. Towles, and H. Fuchs. PixelFlex2: A Comprehensive, Automatic, Casually-Aligned Multi-Projector Display. *IEEE International Workshop on Projector-Camera Systems*, 2003.

[36] R. Raskar. Immersive Planar Display using Roughly Aligned Projectors. In *IEEE VR 2000*, pages 109–116, 2000.

[37] R. Raskar, M. S. Brown, R. Yang, W. Chen, G. Welch, H. Towles, B. Seales, and H. Fuchs. Multi-projector displays using camera-based registration. In *Proceeding of IEEE Visualization 1999*, pages 161–168, 1999.

[38] R. Raskar, J. van Baar, P. Beardsley, T. Willwacher, S. Rao, and C. Forlines. iLamps: Geometrically Aware and Self-configuring Projectors. *ACM Transactions on Graphics (SIGGRAPH 2003)*, 22(3):809–818, 2003.

[39] R. Raskar, J. vanBaar, and J. Chai. A Low Cost Projector Mosaic with Fast Registration. In *Fifth International Conference on Computer Vision (ACCV.02)*, 2002.

[40] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. *Computer Graphics*, 32(Annual Conference Series):179–188, 1998.

[41] R. Raskar, G. Welch, and H. Fuchs. Seamless Projection Overlaps using Image Warping and Intensity Blending. In *Proc. of 4th International Conference on Virtual Systems and Multimedia*, 1998.

[42] R.Surati. *Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays*. PhD thesis, Department of Computer Science, Massachusetts Institute of Technology, 1998.

[43] SEIKO EPSON Corp, Japan. Epson PowerLite 730p. http://www.epson.com/.

[44] M. C. Stone. Color Balancing Experimental Projection Displays. *9th IS&T/SID Color Imaging Conference*, 2001a.

[45] M. C. Stone. Color and Brightness Appearance Issues in Tiled Displays. *IEEE Computer Graphics and Applications*, 2001b.

[46] R. Sukthankar, T.J. Cham, and G. Sukthankar. Dynamic shadow elimination for multi-projector displays. *Proceedings of Computer Vision and Pattern Recognition*, 2001.

[47] R. Sukthankar, R. Stockton, and M. Mullin. Smarter Presentations: Exploiting Homography in Camera-Projector Systems. In *Proceedings of the Proceedings of International Conference on Computer Vision (ICCV)*, Vancouver, July 2001.

[48] R. L. De Valois and K. K. De Valois. *Spatial Vision*. Oxford University Press, 1990.

[49] G. Wallace, H. Chen, and K. Li. Color Gamut Matching for Tiled Display Walls. *Immersive Projection Technology Workshop*, 2003.

[50] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. Brown. PixelFlex: A Reconfigurable Multi-Projector Display System. In *Proceeding of IEEE Visualization*, pages 167–174, 2001.

[51] R. Yang and G. Welch. Automatic Projector Display Surface Estimation Using Every-Day Imagery. In *9th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, 2001.