Advances towards high-resolution pack-and-go displays: A survey

Aditi Majumder
Ezekiel S. Bhasker
Ray Juang

Abstract — Tiled displays provide high resolution and large scale simultaneously. Projectors can project on any available surface. Thus, it is possible to create a large high-resolution display by simply tiling multiple projectors on any available regular surface. The tremendous advancement in projection technology has made projectors portable and affordable. One can envision displays made of multiple such projectors that can be packed in one's car trunk, carried from one location to another, deployed at each location easily to create a seamless high-resolution display, and, finally, dismantled in minutes to be taken to the next location – essentially a pack-and-go display. Several challenges must be overcome in order to realize such pack-and-go displays. These include allowing for imperfect uncalibrated devices, uneven non-diffused display surfaces, and a layman user *via* complete automation in deployment that requires no user invention. We described the advances we have made in addressing these challenges for the most common case of planar display surfaces. First, we present a technique to allow imperfect projectors. Next, we present a technique to allow a photometrically uncalibrated camera . Finally, we present a novel distributed architecture that renders critical display capabilities such as self-calibration, scalability, and reconfigurability without any user intervention. These advances are important milestones towards the development of easy-to-use multi-projector displays that can be deployed anywhere and by anyone.

Keywords — *Tiled displays, projection-based displays, pack-and-go displays.*

1 Introduction

Projectors today are portable and lightweight, so much so that they can even fit on the palm of one's hand (Fig. 1). Therefore, projection technology is ready to realize portable seamless high-resolution displays by tiling multiple portable projectors. Such a pack-and-go display has been the coveted vision of a large number of research communities such as display, computer graphics, visualization, and human-computer interaction.[11,13,17,18,27,28,31,35,37] Until recently, the importability of projectors was considered to be the most serious shortcoming to make this possible. However, since such a display is not possible today even after the advent of the portable projectors, we realize that several technical advances are still critical to make pack-and-go display a reality.

FIGURE 1 — A commercial pocket projector.

1.1 State of the art

Tiling multiple projectors to create one high-resolution display entails two serious *calibration* challenges: (a) *geometric calibration* for stitching the image content across multiple projectors; (b) *color calibration* for achieving a visually seamless brightness and chrominance response within and across the projectors and the overlap region between them. A sub-problem of this is *photometric calibration* that addresses only the issue of brightness seamlessness across the display (Fig. 2).

A decade ago, when heavyweight expensive projectors (~$75,000) driven by monolithic rendering engines were used to build high-resolution displays, only a few of them could be used to create a multi-projector display. Manual calibration was possible and most in-vogue in such small systems. For example, mounts with six degrees of freedom were manipulated manually to achieve geometric calibration, and manual manipulation of color controls of projectors achieved color balancing across multiple projectors. Further, since use of expensive optical elements was justified for such an expensive set up, several such optical elements were used for calibration. For example, expensive

Extended revised version of a paper presented at the 2007 SID Symposium, Seminar & Exhibition (SID '07) held in Long Beach, California, May 20–25, 2007.

The authors are with the University of California, Department of Computer Science, 3029 Bren Hall, Irvine, CA 92697; telephone 949/824-8877, fax –4056, e-mail: majumder@ics.uci.edu.

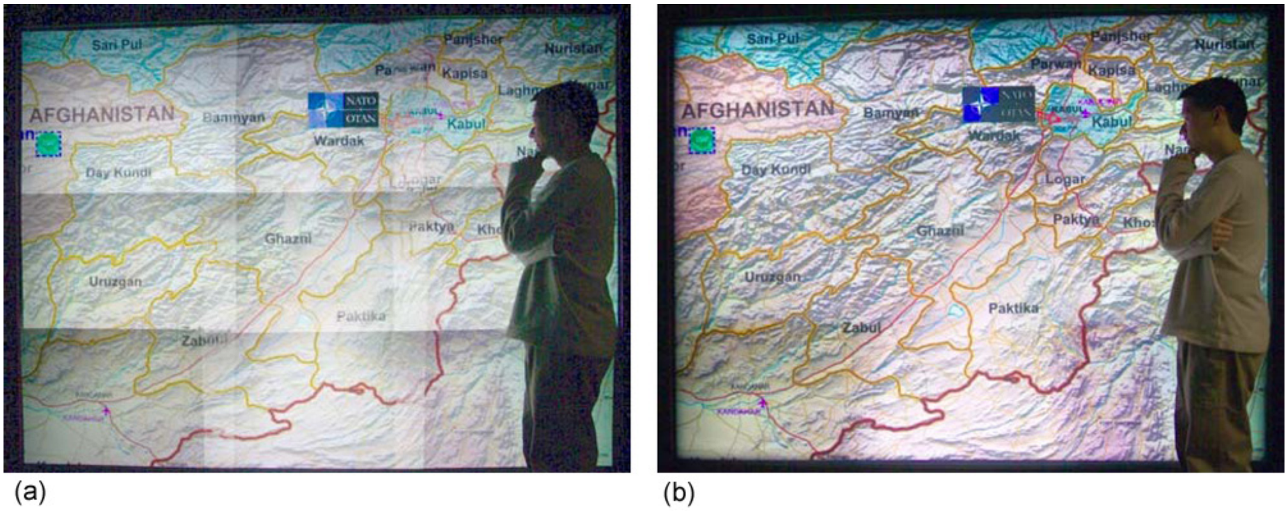

FIGURE 2 — Nine rear projectors arranged in a 3 × 3 rectangular array on a planar surface creating a large display of size 10 × 8 ft. and resolution of 3000 × 2500 offering 7 million pixels in our Large Area Display Laboratory (LADL) at UCI. The display is driven by a nine-PC cluster, each PC driving one projector. Left: The uncalibrated display showing geometric mismatches and photometric variations. Right: The seamless display created by centralized automated camera-based geometric and brightness calibration.

Fresnel lenses in projectors were common for achieving photometric uniformity in each projector.

Today, projectors are commodity products (~$1500). Further, the tremendous advancement in graphics hardware development has made a PC-cluster-based driving architecture possible at a cost which is an order of magnitude lower than the rendering engines of yesteryears. Hence, many organizations can afford a 10–12 projector display (~$50,000) that can provide a reasonably high resolution of about 10–12 Mpixels. Manual calibration is cost-prohibitive and infeasible in such a large system. Thus, the last decade has seen the development of a number of automated camera-based calibration techniques.[6,7,14,21,23,24,26,29,30,32,33,36,38,39] These use camera(s) to provide feedback on the geometric/photometric mismatches which is then used to correct the projected images digitally to achieve a seamless display.

1.2 Limitations

Development of automated camera-based techniques was driven by the promise of realizing multi-projector displays that can be easily deployed by the layman user, such as a doctor in a hospital or a historian in a museum.[31] However, we are still quite far from that vision. We believe this is primarily due to the rigidity in the associated architecture, devices, and algorithms.

First, all camera-based calibration algorithms assume perfect projectors with no lens distortion or vignetting effect. Further, they assume a perfectly calibrated camera whose non-linear response, vignetting artifact, and lens distortion have been estimated by prior device calibration. Needless to say, such device calibration often is more complex than calibrating the display itself and cannot foster the confidence in a layman user to attempt to deploy such a display herself.

Second, all camera-based calibration algorithms devised up until now use a *centralized architecture* where a centralized server controls all the projectors and camera(s) as shown in Fig. 3. It dispatches imagery to a number of clients, each displaying the image on its display unit. To calibrate the display, the server (a) instructs the projector to project patterns; (b) instructs the camera to capture the patterns; (c) generates correction parameters; and (d) corrects the image and sends appropriate parts of it to the projectors to achieve a seamless display. Thus, this is an absolutely centralized and synchronized architecture where the server bears the overhead of managing the imagery while the clients are dumb display stations. Although simple to instrument, the entire responsibility of calibrating the display rests with the user. Hence, the user needs to be knowledgeable about the workings of the software and the camera. Further, this does not support easy scalability or reconfigurability to allow quick change in display scale, form-factor, and resolution. These are often important capabilities, especially in the context of pack-and-go displays where the display needs to adapt to the changing scale, form factor and resolution of the data.

1.3 Desiderata

Our vision is to build algorithms and architectures that can enable a common surface to be wall-papered by multiple projected images to create a single seamless display that is scalable (number of projectors can be increased easily), reconfigurable (projectors can be set up in different configurations based on application requirements), self-calibrating (does not need any input from the user to calibrate themselves), and fault-tolerant (responds to faults in an appropriate manner). Further, we would like to use regular commodity projectors/cameras that come with several opti-

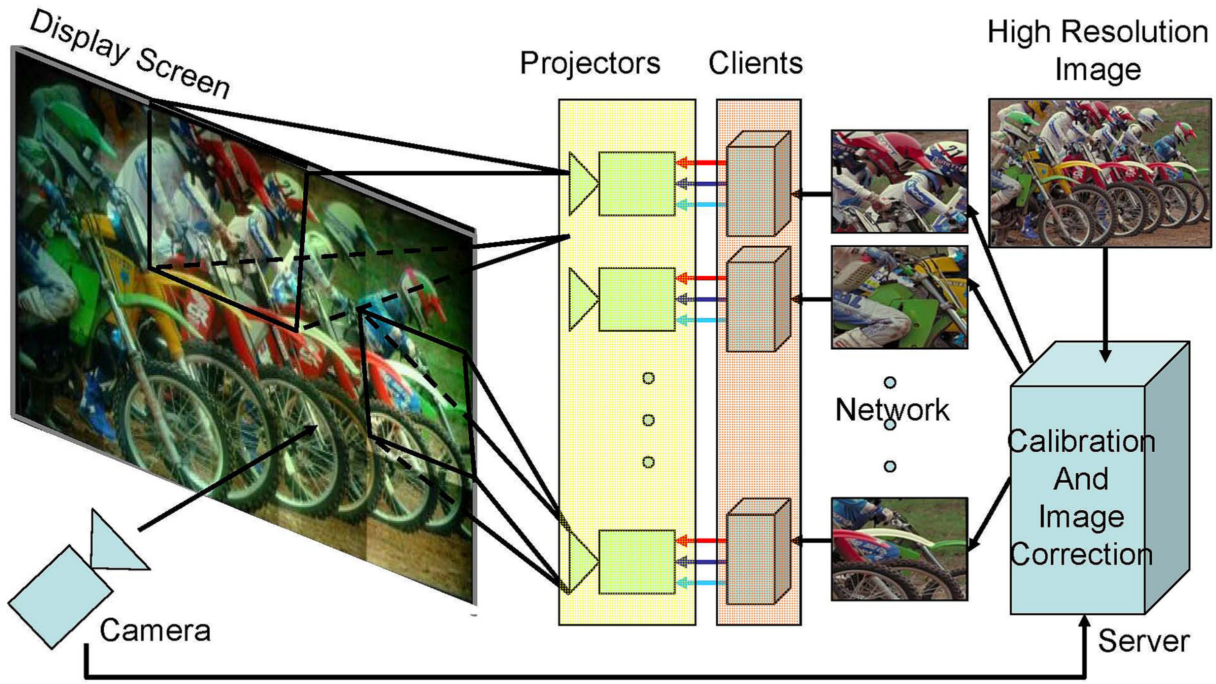

FIGURE 3 — A centralized architecture where a single server controls all the pixels [all projectors and camera(s)].

cal imperfections and are not pre-calibrated. Finally, all the above capabilities should be instrumented while maintaining an extreme simplicity in deployment so that they can be built almost anywhere, by almost anybody, at almost any scale and form factor. The ultimate goal is to trigger a paradigm shift in the seamless high-resolution display technology where people no longer perceive them to be mammoth structures quarantined in a huge room of a high-tech laboratory but can carry them around.

1.4 Our advances

Moving towards such a high-resolution pack-and-go display, we present a survey of the several advances we have made towards this goal in the past few years. However, all these advances are pertaining to the most common case of planar display surfaces and are as follows.

1. Staying within the simplicity of the centralized architecture, we have designed models and algorithms that can use geometrically and photometrically *imperfect* projectors when building multi-projector displays (Secs. 2 and 3).[19,20,22–24]
2. We have designed a new photometric self-calibration technique for a uncalibrated projector-camera pair. This, in turn, allows us to use a *photometrically uncalibrated* camera to calibrate a multi-projector display, when assuming a centralized architecture (Sec. 4).[3]
3. Finally, we break away from the centralized architecture and propose a novel *decentralized architecture* for a multi-projector display (Sec. 5).[4,10] Here, a distributed network of self-sufficient display units – made of a projector, a camera and embedded computation and communication hardware – is used to build the display. No central machine is responsible for managing all the pixels. Each display unit manages its own pixels autonomously yielding a completely scalable architecture. This enables several critical capabilities such as self-calibration with no user intervention, reconfigurability, and fault tolerance.

2 Geometrically imperfect projectors

Camera-based geometric calibration of a tiled planar display entails defining two functions: (a) a function relating each projector to the observing camera and (b) a function that relates the camera to the screen (the planar display surface). Geometric calibration techniques assume a camera that is pre-calibrated to exhibit no lens distortion. In such cases, the camera-to-screen function is linear, expressed by a 3×3 matrix. Geometric calibration methods can be classified in three following categories based on the function that relates the projector to the camera.

1. *Linear methods* assume a geometrically perfect projector and use a linear 3×3 matrix to relate the projector to the camera.[6,36] This produces reasonable results in high-end expensive projectors where setting the projectors in a "sweet" zoom level assures no non-linear distortion.
2. *Piecewise linear methods* use a dense piecewise linear triangulation to relate the projector and camera[38] and can handle non-linear distortions in the projector. However, since this method needs a

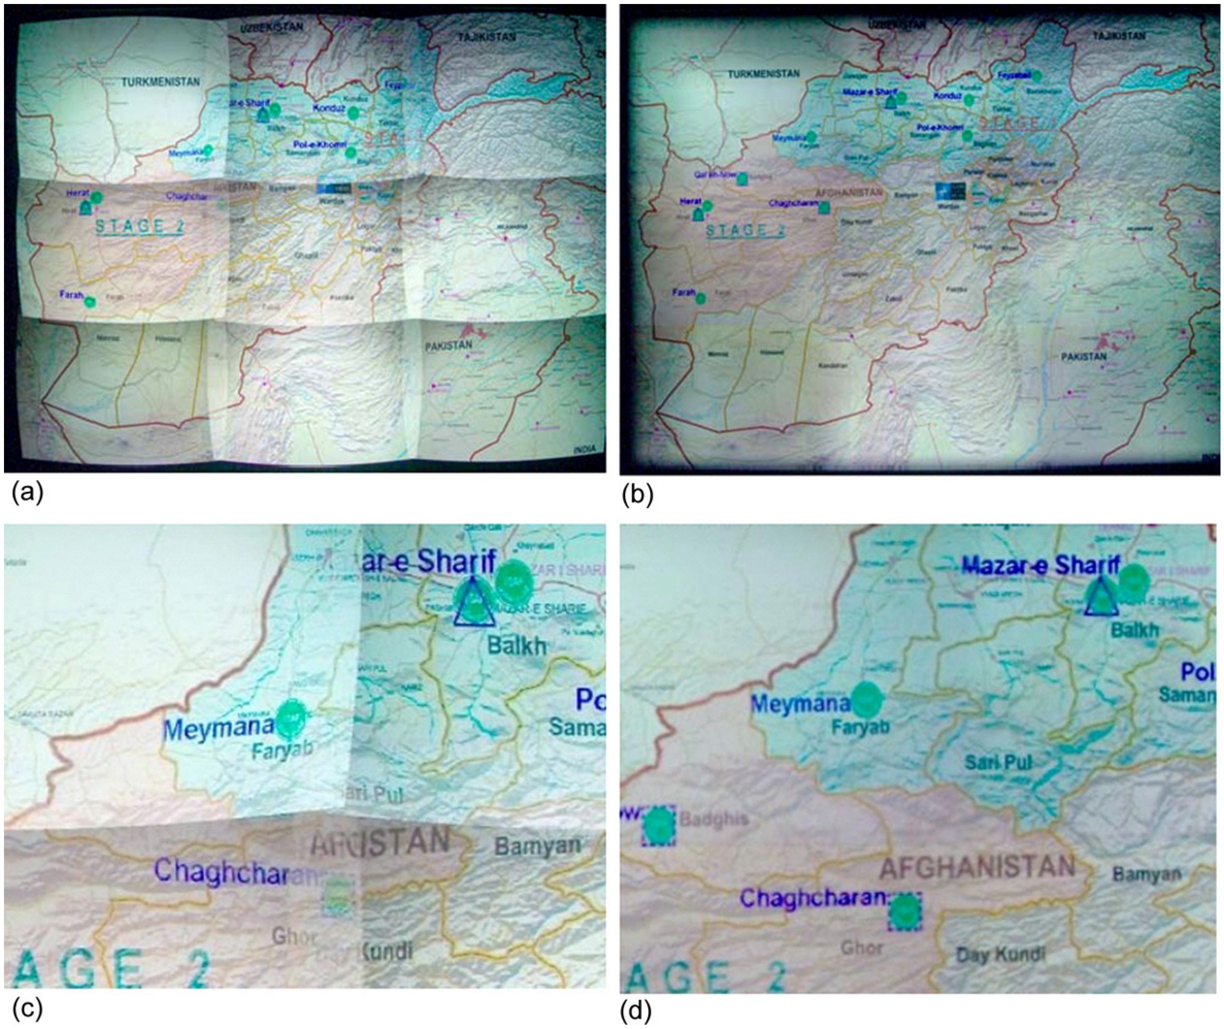

FIGURE 4 — Top row – Left: Our nine-projector 3000 × 2500 display with severely distorted projectors. Right: The same display geometrically calibrated using a VGA 640 × 480 camera. Bottom row: zoomed-in view of the respective top pictures.

dense sampling of this function, especially in the presence of adverse non-linearities, a very high-resolution expensive camera is essential to instrument it.

3. *Polynomial methods* use a cubic polynomial to approximate the projector non-linearities.[14] However, the projector-to-camera relationship involves a perspective projection which is not adequately modeled by a simple cubic polynomial. Thus, these methods assume a near rectangular array of projectors that is difficult to setup, especially in front-projection systems.

As is evident, all these methods address geometric imperfections in projectors only in a limited manner. But, current commodity projectors have several geometric imperfections – non-linearities that cause straight lines to become curved. The most common cause of these non-linearities is lens distortion that has been modeled extensively before.[5,12] However, optical folding, common in today's compact projectors, can also cause higher-order geometric non-linearities that cannot be modeled by these traditional models.

In Ref. 3, we present a new geometric model that can unify both the non-linearities of the projector and the perspective projection between the projector and the camera. We show that a Bezier patch can adequately model the different non-linearities of the projector, including the lens distortions, in a unified manner. The richer cross terms of this new model allows for accurate modeling of these effects even with a relatively lower order patch. A generalized form of the Bezier, called the rational Bezier, retains all the properties of the Bezier and is additionally perspective invariant. We show that this rational Bezier is well-suited to unify the non-linear distortions of the projector and the perspective projection between the projector and the camera. Thus, a *rational Bezier patch* can accurately and compactly model the relationship between the projector and the camera. Using this model, we present a new geometric calibration method that can allow the use of severely distorted projectors when building a tiled display. We also show that the

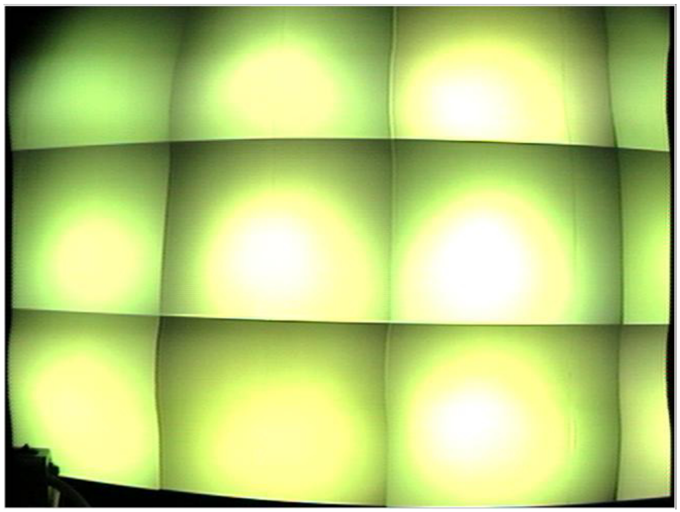

FIGURE 5 — Severe projector vignetting.

parameters of this model can be easily estimated from a sparse set of correspondences between the projector and camera, and hence the geometric calibration can be achieved using an inexpensive low-resolution camera. For example, our 3000×2500 display made of nine projectors was calibrated using a VGA camera of 640×480 resolution (Fig. 4).

The biggest impact of this method can be in compact multi-projector display design. Short-throw lenses are inevitable for such compact designs, but they usually introduce severe non-linearities to the projected imagery. To avoid such non-linearities, current lenses use optical corrective measures which make them 4–5 times more expensive than the projectors themselves. Our method allows software solutions so that inexpensive lens can be mounted on projectors, just as they are used for cameras today.

3 Photometrically imperfect projectors

The most common photometric imperfection in projectors is that of vignetting – a spatial fall-off in brightness from center to fringe.[23] This can be aggravated by high-gain screens causing distracting brightness variations that can completely break the illusion of single display (Fig. 5). It is often beneficial to overlap adjacent projectors to alleviate the brightness fall-off at the fringes. However, the brighter overlap regions now result in a display than show a very large brightness variation – almost 25–30% of the maximum brightness at some regions.

There was no existing camera-based calibration technique to compensate for this photometric variation. In Refs. 22 and 23, we presented the first method to measure the spatial brightness variation of the display using a camera. We also presented a method to achieve photometric uniformity by matching the photometric response of the brighter pixels to the dimmer ones which are most limited in their capabilities. However, this results in all the pixels sacrificing their higher capabilities to operate at a very limited capability. The resulting display, though seamless, showed severe contrast degradation – to the extent of being almost useless. To alleviate this problem, we presented a perception-based smoothing technique that smoothens the brightness response across the display (instead of matching) leading to a photometrically non-uniform display. But the controlled smoothing, catered specifically for the particular display, removes the photometric seams without any perceivable contrast degradation,[20,24] resulting in a seamless high-contrast display (Fig. 6).

4 Photometrically uncalibrated camera

The photometric calibration results presented in Sec. 3 assumes a photometrically calibrated camera, *i.e.*, a camera whose gamma function is known *apriori*. This requires complex device pre-calibration mechanisms.[9,25]

In our recent work,[3] we present a photometric self-calibration method for an uncalibrated projector–camera pair that can estimate the gamma function of the camera and projector, and the vignetting effect of the projector *simultaneously.* The results of this method are illustrated in Fig. 7. By using this method, we can now photometrically calibrate a display using a photometrically uncalibrated camera (Fig. 2).

5 A decentralized approach

As mentioned in Sec. 1, the difficulty in deployment of multi-projector displays primarily stems from a centralized architecture that burdens the user with responsibilities to

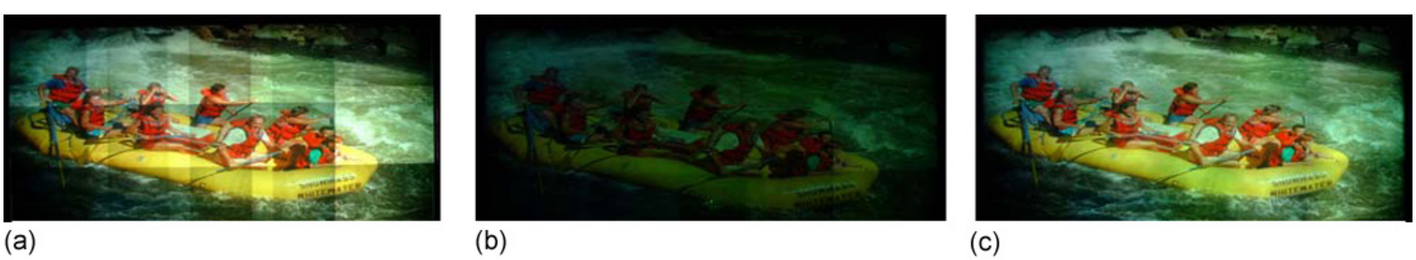

FIGURE 6 — A 15-projector display at Argonne National Laboratory before any photometric calibration. Middle: The display after photometric uniformity showing severe contrast reduction. Right: The same display after photometric smoothing showing no perceptual seams and retaining most of the original contrast.

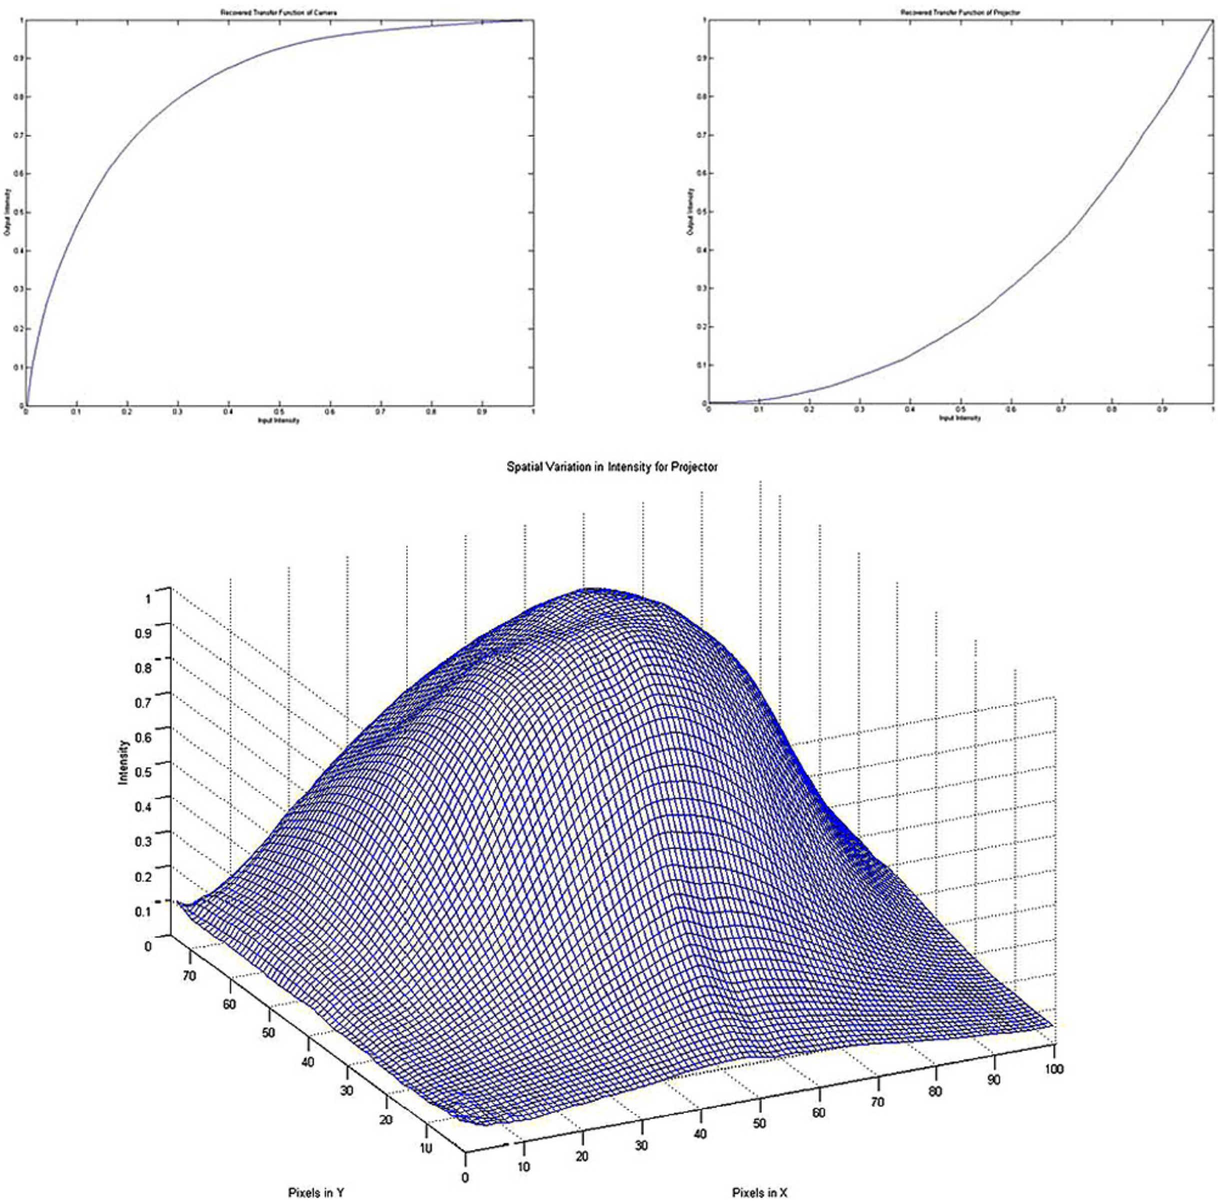

FIGURE 7 — Estimated parameters from the photometric self-calibration of the uncalibrated projector–camera pair. Left: the camera gamma function; middle: the projector gamma function; right: the projector vignetting effect.

calibrate the display, pre-calibrate the devices, and provide new input to trigger a complete recalibration with a change in the scale and form factor of the display. Instead, we would like to build a multi-projector display that would be able to calibrate itself, with no input from the user. It should detect additions, removals, and faults, and recalibrate in response to these events. Essentially, all the user should do is to arrange the projectors physically and the rest will be taken care of by the system. Thus, the projector will become almost like a flashlight, light from which can be moved around wherever one desires. And a cluster of these projectors should be able to create a giant high-resolution display without the user worrying about manual set-up or maintenance.

In our recent work, we presented a completely decentralized approach towards this end.[4,10] First, we presented a self-sufficient display unit that becomes the building block for such self-calibrating displays. Next, we present a decentralized architecture to build a multi-projector display *via* a distributed network of these display units. Finally, we present an asynchronous decentralized calibration process that provides several novel capabilities such as self-calibration, reconfigurability, and fault tolerance.

5.1 The display unit

Traditionally, pixels have been the most important commodity of any workspace being extensively used for visualization, collaboration, and interface. As opposed to pixels provided by traditional displays such an LCD panel or a CRT monitor, pixels from a projector have greater flexibility and mobility

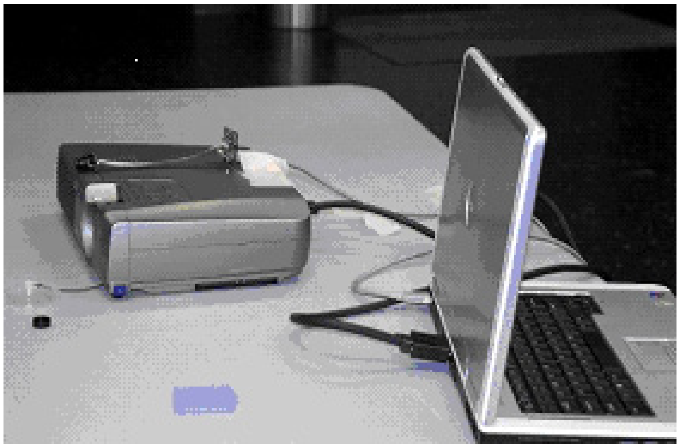

FIGURE 8 — A prototype self-sufficient display unit.

and can be used to illuminate any arbitrary surface.[31,38] However, when used alone, projectors act like passive digital illumination devices, where pixels act as dumb entities while what we desire are intelligent entities that can adapt/react to environmental changes. So, a combination of projectors and cameras is inevitable to provide 'intelligent' pixels and has been envisioned by many other researchers.[1,2,8,27,33] Our self-sufficient display unit is largely inspired by such previous works.

We propose an augmented display unit consisting of a projector, a camera, and embedded computation and communication hardware. Each display unit is thus self-sufficient with the capacity to sense environmental changes (using the camera), adapt or react to those changes (using the computation unit), and communicate those changes to other display units (using the communication unit). We use a standard projector, attach a standard web-cam to it, and use a laptop to simulate the computation and communication unit (Fig. 8).

5.2 The architecture

In Refs. 4 and 10 we propose to build our display by constructing a distributed network of the self-sufficient display units (Fig. 9). This architecture removes the dependency on a single central machine and every display unit takes complete control of the part of the display it is responsible for. They act like a client and request the appropriate part of the data from a traditional data server. In fact, this server can be oblivious of the fact that the clients requesting data are in reality display units. So, each display unit behaves just like any other independent data client and is entirely responsible for managing its own pixels.

Existing distributed rendering architectures such as Chromium and SAGE[15,16,34] use distributed methodologies only for rendering the pixels and use centralized architecture for calibration and data handling. Thus, the user defines, in the central server, the total number of display units and the relationship of the image projected by each unit with respect to the large image they are creating together. This centralized unit then streams the appropriate data to the dumb display units or projectors. Unlike such a centralized architecture, our architecture uses distributed methodologies in all aspects of multi-projector display, including calibration, data handling, and rendering.

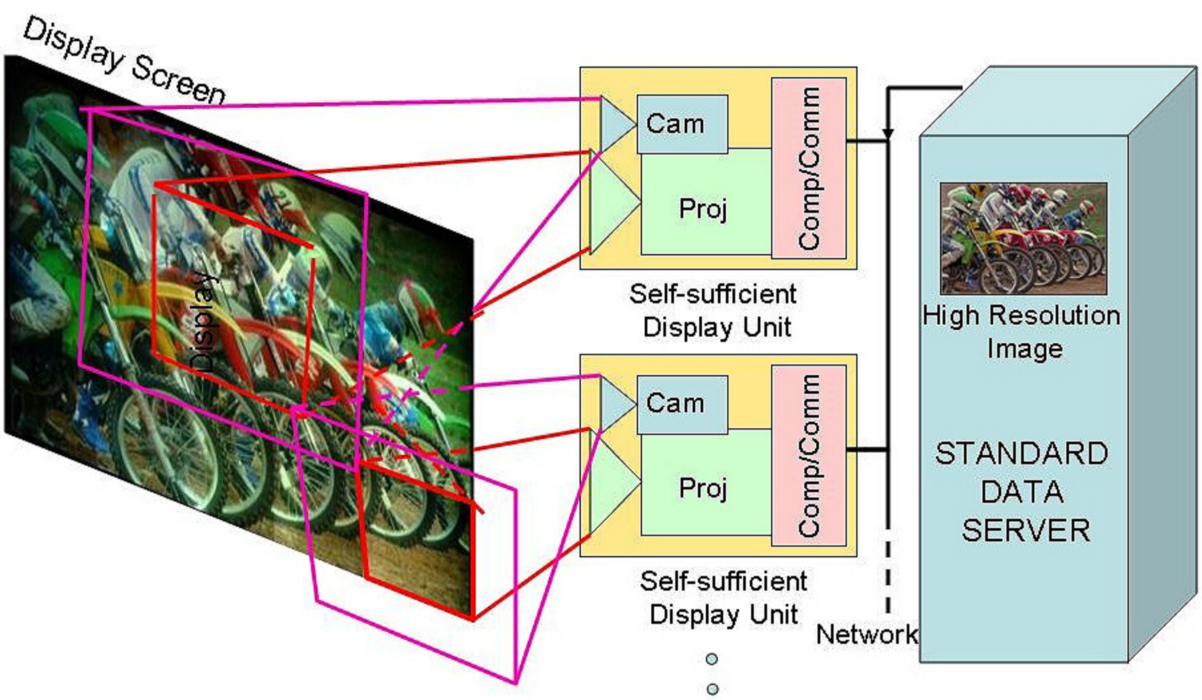

FIGURE 9 — A tiled multi-projector display made of a distributed network of self-sufficient display units.

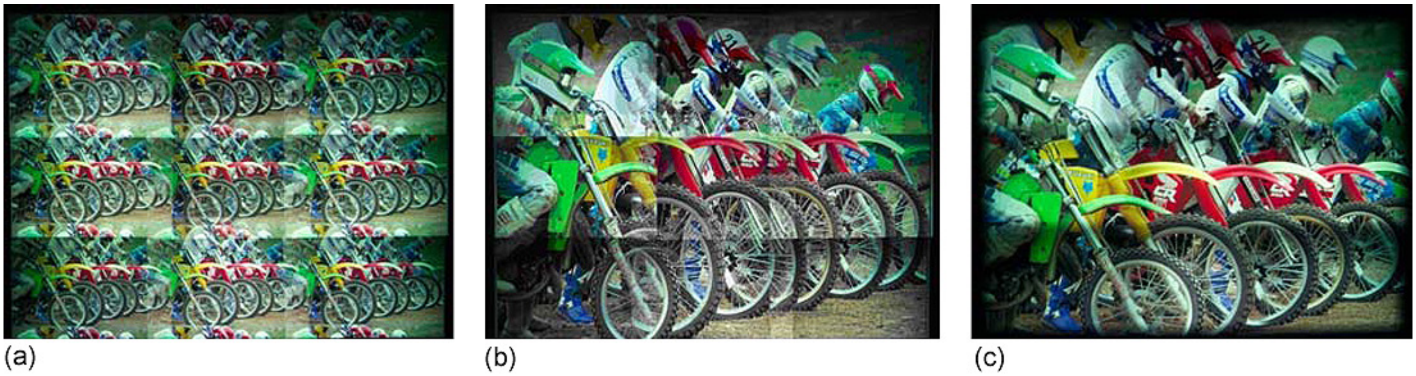

FIGURE 10 — Left: Initially, every display unit thinks that it is only display unit present and is therefore solely responsible for displaying the whole image. Middle: After configuration identification, each display unit knows the display configuration – total number of projectors, and total display dimensions – and their own coordinates in the array. Thus, they know which parts of the display they are responsible for, but still do not know the relative orientations of their neighbors. Thus, the image is not seamless. Right: after alignment each display unit matches geometrically and photometrically with its neighbors to create a seamless display.

The critical capabilities of a pack-and-go display are self-calibration and easy reconfigurability leading to an unforeseen ease in deployment. Our distributed architecture enables several such critical capabilities listed as follows.

5.3 Camera-based communication

The camera on each display unit sees only a part of the display. However, the camera field of view needs to be a little larger than that of the projector, so that it can observe parts of its neighbors' projections and react accordingly. This ability of each unit to sense changes in its neighbor via its own camera enables us to use an alternate modality of camera-based communication. Patterns observed by a neighboring projector are analyzed to find the relative orientation between the neighboring projectors, and locations in the display array. Thus, the overlapping field of view of sensors in adjacent display units provides an underlying mesh interconnection network at no additional cost.

5.4 Self-calibration

The camera-based communication enables the design of an asynchronous decentralized self-calibration algorithm. This is a SPMD (single program multiple data) algorithm that runs on each display unit to achieve calibration with no user intervention. Figure 10 provides an illustration of the state of the display at different steps of this decentralized self-calibration process. Initially, every unit believes it is alone in the environment and has the sole responsibility of displaying the data. Then each display unit runs the identical SPMD algorithm consisting of the two following steps.

First, each display unit identifies its immediate neighbors, the display configuration, and its own coordinates in the display. Each display unit uses asynchronous camera-based and wireless communication to find its own (row, column) in the 2-D array, and left, right, top, and bottom neighbors, whenever they exist. This takes $O(m + n)$ parallel steps for an $m \times n$ array of display units.

Finally, it projects an image that is geometrically matched with its neighbors and edge-blended to achieve a relatively seamless display. Because no single camera sees the entire display, the geometric calibration cannot assume any global coordinates. We use a decentralized method to perform geometric correction without considering a global coordinate system, where one projector is dynamically chosen to be the leader against whom everyone else is calibrated. This runs in $O(m + n)$ for an $m \times n$ array of display units. To avoid a change of coordinate in every frame (which can lead to visual flicker), the leader does not change unless it crashes. Photometric calibration is achieved by a rudimentary local edge blending that can be computed from the relative geometric positions of only the adjacent display units. Since this does not need communications with non-adjacent units, it is performed in constant time.

5.5 Flexibility

Since the camera and projector is in a feedback loop in every display unit, addition or removal (can be due to a fault) of a unit is easily detected by neighboring units. This initiates the recalculation of display configuration and individual positions of the display units, followed by the reshaping of the imagery to fit into the largest rectangle contained within the non-rectangular projected area of all the active display units. This is illustrated in Fig. 11. Thus, we achieve easy reconfigurability, complete scalability, and fault tolerance which allows the display to easily adapt itself to different sizes, aspect ratio, and resolution as demanded by the user/data/situation.

6 Conclusion

In this paper, we have presented a survey of the advances we have made towards the vision of pack-and-go displays. Such

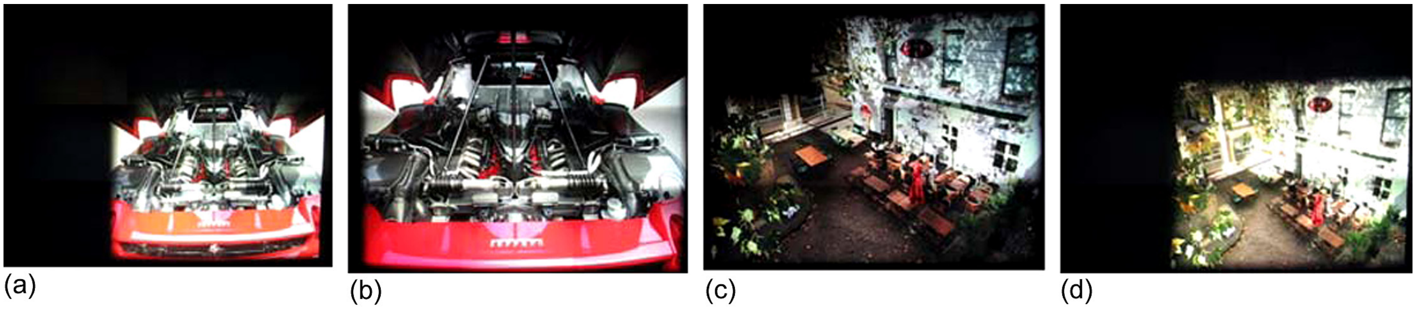

FIGURE 11 — Left two: This illustrates recalibration when new display units are added to a 2 × 2 array to create a large display of 3 × 3 array. Right two: here, the top left display unit in the 3 × 3 display fails, following which the display deactivates the appropriate units to create a 2 × 2 array with limited resolution but the same aspect ratio.

pack-and-go displays can have broad applications in many areas such as archeology, education, medicine, and entertainment. In particular, it can enable mobile science laboratories and mobile command-and-control units (in war grounds or emergency-incident sites).

We believe that the methods outlined in this paper, especially the decentralized architecture and algorithms will be critical to instrument pack-and-go displays that can be set-up easily, dismantled, and moved to a different location when required. However, more interestingly, the same devices, architectures, and algorithms have the potential to be used in many other domains where it can change the way we interact with our everyday environment.

First, a display unit that can both display and sense its environment, realized by combining a projector and a camera, could spark and foster novel paradigms of collaboration where each person carries his own display unit and when more than one person meet for collaboration, their respective displays are put together to create a seamless high-resolution display. This display can easily scale as the number of collaborators increase. Further, such a shared display space that has access to data from multiple machines might foster new directions of research in user interfaces for data sharing.

Second, the algorithms and the architecture developed can be instrumental to realize mammoth seamless visualization systems. Large visualization systems today are built primarily using very large arrays of LCD panels. However, the bezels around the panels result in seams that can be distracting and even detrimental in executing certain tasks due to mangling of text and patterns. But we still tend to use LCD panels due to the relative ease in setting them up on a single substrate in a rigid fashion, even though this initial set-up demands huge financial, engineering, and personnel resources. People do not consider a seamless projection-based tiled display since installing and maintaining them is very difficult. And, of course, the necessity of advanced capabilities such as scalability and reconfigurability are not even evaluated. The decentralized algorithms and architectures will enable next-generation super-high-resolution entirely seamless visualization, training and simulation systems where the number of pixels can be scaled easily to even billions. Thus, this will enable displays that can match the size and resolution of the exponentially growing size of today's data.

6.1 Future work

Our advances, however, are just the first steps towards moving the frontier of technology to realize pack-and-go displays. A large number of challenges still need to be handled to make pack-and-go displays the reality of the future. Some of these challenges are listed below.

1. The distributed architecture and algorithms (Sec. 5) use geometrically and photometrically pre-calibrated projectors and cameras. It will be important to investigate ways to extend our centralized techniques to handle imperfect uncalibrated devices (Secs. 2–4) in such a distributed architecture.
2. We have used only a local edge blending for the distributed photometric calibration. Edge blending can only yield reasonable results when the projectors have similar brightness.[20,24] We have to investigate ways to extend our centralized algorithm of brightness smoothness to the distributed architecture.
3. Currently, only photometric calibration can be done using a camera which addresses only the brightness of the display. No camera-based automatic color calibration technique exists that can handle both brightness and chrominance, even for the relatively simple centralized architecture. Hence, the projectors in most displays are color balanced *apriori* via manual or semi-automatic methods that use a radiometer as a feedback device. Camera-based color calibration requires addressing the fundamental issue of measuring color using a camera. It is critical to perform extensive analysis in this direction to evaluate the feasibility of the solution space and then explore the feasible solutions, if any.
4. For both centralized and distributed architecture, there exists no geometric calibration method that can allow the use of uncalibrated cameras that show significant geometric non-linearities.

5. We have only addressed the case of planar multi-projector display. However, to instrument tiled displays on regularly available surfaces, that are neither planar nor diffused, is a big challenge and very little has been done so far in this direction.

References

1 L. Childers, T. Disz, R. Olson, M. E. Papka, R. Stevens, and T. Udeshi, "Immersive group-to-group collaborative visualization," *Proc. Immersive Projection Technol.* (2000).

2 M. Ashdown, M. Flagg, R. Sukthankar, and J. M. Rehg, "A flexible projector-camera system for multi-planar displays," *Proc. Computer Vision Pattern Recognition (CVPR)* (2004).

3 E. Bhasker, R. Juang, and A. Majumder, "Registration techniques for using imperfect and partially calibrated devices in planar multi-projector displays," *IEEE Trans. Visualization Computer Graphics (Visualization)* (2007).

4 S. B. Ezekiel, P. Sinha, and A. Majumder, "Asynchronous distributed calibration for scalable reconfigurable multi-projector displays," *IEEE Trans. Visualization Computer Graphics (TVCG)* **12**(5), 1101–1108 (2006).

5 D. C. Brown, "Decentering distortion of lenses," *Photometric Eng.* **32**(3), 444–462 (1966).

6 H. Chen, R. Sukthankar, G. Wallace, and K. Li, "Scalable alignment of large-format multiprojector displays using camera homography trees," *Proc. IEEE Visualization* (2002).

7 Y. Chen, D. W. Clark, A. Finkelstein, T. Housel, and K. Li, "Automatic alignment of high-resolution multi-projector displays using an uncalibrated camera," *Proc. IEEE Visualization* (2000).

8 D. Cotting, R. Ziegler, M. Gross, and H. Fuchs, "Adaptive instant displays: Continuously calibrated projections using per-pixel light control," *Proc. Eurographics,* 705–714 (2005).

9 P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," *Proc. ACM Siggraph,* 369–378 (1997).

10 E. Bhasker and A. Majumder, "Self-calibrating tiled displays," *SID Symposium Digest Tech. Papers* **38**, 23–26 (2007).

11 T. L. Disz, M. E. Papka, and R. Stevens, "Ubiworld: An environment integrating virtual reality," *Heterogeneous Computing Workshop,* Geneva, Switzerland (1997).

12 J. G. Fryer and D. C. Brown, "Lens distortion for close-range photogrammetry," *Photogrammetric Engineering and Remote Sensing* **52**(1), 51–58 (1986).

13 M. Hereld, I. Judson, and R. Stevens, "Introduction to building projection-based tiled display systems," *IEEE Computer Graphics and Applications* (2000).

14 M. Hereld, I. R. Judson, and R. Stevens, "Dottytoto: A measurement engine for aligning multiprojector display systems," Argonne National Laboratory Preprint ANL/MCS-P958-0502 (2002).

15 G. Humphreys, M. Eldridge, I. Buck, G. Stoll, M. Everett, and P. Hanrahan, "Wiregl: A scalable graphics system for clusters," *Proc. ACM SIGGRAPH* (2001).

16 G. Humphreys, M. Houston, R. Ng, R. Frank, S. Ahem, P. Kirchner, and J. Klosowski," Chromium, "A stream processing framework for interactive rendering on clusters," *ACM Trans. Graphics* (2002).

17 J. Lai, A. Levas, P. Chou, C. Pinhanez, and M. Viveros, "Bluespace: Personalizing workspace through awareness and adaptability," *Intl. J. Human Computer Studies* **57**(5), 415–428 (2002).

18 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 Applications* **20**(4), 671–680 (2000).

19 A. Majumder, "Contrast enhancement of multi-displays using human contrast sensitivity," *Proc. IEEE Intl. Conf. Computer Vision Pattern Recognition (CVPR)* (2005).

20 A. Majumder, "Luminance management for seamless multi-projector displays," *SID Symposium Digest Tech. Papers* **36**, 1506–1509 (2005).

21 A. Majumder, Z. He, H. Towles, and G. Welch, "Achieving color uniformity across multiprojector displays," *Proc. IEEE Visualization* (2000).

22 A. Majumder and R. Stevens, "LAM: Luminance attenuation map for photometric uniformity in projection based displays," *Proc. ACM Virtual Reality Software Technology* (2002).

23 A. Majumder and R. Stevens, "Color nonuniformity in projection-based displays: Analysis and solutions," *IEEE Trans. Visualization Computer Graphics* **10**(2) (March–April, 2003).

24 A. Majumder and R. Stevens, "Perceptual photometric seamlessness in tiled projection-based displays," *ACM Trans. Graphics* **24**(1) (January, 2005).

25 T. Mitsunaga and S. Nayar, "Radiometric self calibration," *Proc. IEEE Conf. Computer Vision Pattern Recognition* (1999).

26 S. K. Nayar, H. Peri, M. D. Grossberg, and P. N. Belhumeur, "A projection system with radiometric compensation for screen imperfections," *Proc. IEEE Intl. Workshop on Projector-Camera Systems* (2003).

27 C. Pinhanez, "The everywhere displays projector: A device to create ubiquitous graphical interfaces," *Proc. Ubiquitous Computing* (2001).

28 C. S. Pinhanez, F. C. Kjeldsen, A. Levas, G. S. Pingali, M. E. Podlaseck, and P. B. Chou, "Ubiquitous interactive graphics," Technical Report RC22495 (W0205-143), IBM Research (2002).

29 A. Raij, G. Gill, A. Majumder, H. Towles, and H. Fuchs, "PixelFlex2: A comprehensive, automatic, casually-aligned multi-projector display," *IEEE Intl. Workshop on Projector-Camera Systems* (2003).

30 R. Raskar, M. S. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs, "Multi-projector displays using camera based registration," *Proc. IEEE Visualization* (1999).

31 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 display," *Proc. ACM Siggraph,* 168–176 (1998).

32 R. Raskar, "Immersive planar displays using roughly aligned projectors," *Proc. IEEE Virtual Reality* (2000).

33 R. Raskar, J. van Baar, P. Beardsley, T. Willwacher, S. Rao, and C. Forlines, "ilamps: Geometrically aware and self-configuring projectors," *ACM Trans. Graphics* **22**(3) (2003).

34 L. Renambot, A. Rao, R. Singh, B. Jeong, N. Krishnaprasad, V. Vishwanath, V. Chandrasekhar, N. Schwarz, A. Spale, C. Zhang, G. Goldman, J. Leigh, and A. Johnson, "Sage: The scalable adaptive graphics environment," *Proc. WACE* (2004).

35 P. Steurer and M. B. Srivastava, "System design of smart table," *Proc. First IEEE Intl. Conf. on Pervasive Computing and Communications* (2003).

36 R. Sukthankar, R. Stockton, and M. Mullin, "Smarter presentations: Exploiting homography in cameraprojector systems," *Proc. Intl. Conf. on Computer Vision (ICCV)* (2001).

37 J. Underkoffler, B. Ullmer, and H. Ishii, "Emancipated pixels: Real-world graphics in the luminous room," *Proc. ACM Siggraph* (1999).

38 R. Yang, D. Gotz, J. Hensley, H. Towles, and M. S. Brown, "Pixelflex: A reconfigurable multi-projector display system," *Proc. IEEE Visualization* (2001).

39 R. Yang, A. Majumder, and M. Brown, "Camera based calibration techniques for seamless multi-projector displays," *IEEE Trans. Visualization Computer Graphics* **11**(2) (March–April, 2005).

Aditi Majumder is an Assistant Professor in the Department of Computer Science at University of California, Irvine. She graduated with a Ph.D. from the University of North Carolina in Chapel Hill in 2003. Her research spans the domains of visualization, computer vision, and graphics and focuses primarily on high-resolution tiled displays. She has published more than 35 conference and journal publications that include premium venues such as SID, *IEEE Visualization, IEEE Transactions on Visualization and Computer Graphic,* and *ACM Transaction on Graphics,* and holds multiple patents. She has offered several tutorials on the topic of tiled projection-based displays at premium computer graphics and vision venues. She is the author of the book *Practical Multi-Projector Display Design,* published by A.K. Peters, released in August 2007.

Ezekiel S. Bhasker is a Ph.D. student at the Department of Computer Science, University of California, Irvine, from Fall 2005 and works on tiled projection-based displays. He received his B.S. degree from University of California, San Diego, in 2005.

Ray Juang graduated with a M.S. degree from the Department of Electrical Engineering and Computer Science, University of California, Irvine, in June 2007. He is currently a research associate at the Johns Hopkins University Applied Physics Laboratory. His research interests are in computer vision, machine learning, and tracking.