

Display Grid: Ad-hoc Clusters of Autonomous Projectors

Ramesh Raskar, Jeroen van Baar, Paul Beardsley

TR2004-033 May 2004

Abstract

Inspired by emerging trends in grid computing, we present techniques for creating multi-projector displays via self-configuring clusters of autonomous projectors. The ad-hoc clustering approach avoids large monolithic installations. We show a low-cost system that supports dynamic inclusion of new projectors, automatic geometric configuration and seamless blending of overlapping projectors.

SID 2004

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

To Appear in SID International Symposium 2004 Proceedings



59.1: Display Grid: Ad-hoc Clusters of Autonomous Projectors

Ramesh Raskar¹, Jeroen vanBaar, Paul Beardsley
Mitsubishi Electric Research Labs (MERL), Cambridge, MA 02139, USA

Abstract

Inspired by emerging trends in grid computing, we present techniques for creating multi-projector displays via self-configuring clusters of autonomous projectors. The ad-hoc clustering approach avoids large monolithic installations. We show a low-cost system that supports dynamic inclusion of new projectors, automatic geometric configuration and seamless blending of overlapping projectors.

1. Introduction

Projectors are getting smaller, brighter, and cheaper. In addition, they are becoming networked and are being built with on-board sensors and intelligence. The evolution of computers is suggestive of the ways in which future projectors might evolve. As computers evolved from mainframes to PCs, the application domain went from large scientific and business computations to small personal efficiency applications. Computing has also seen an evolution from well-organized configurations of mainframes to clusters of heterogeneous, self-sufficient computing units. In the projector world, we may see similar developments – towards portable devices for personal use; and a move from large monolithic installations towards ad-hoc, self-configuring displays made up of heterogeneous, self-sufficient projector units.

Emerging technologies such as sensor networks or cluster computing are blending multiple different devices into a single vast pool of resources. In simple terms, grid (or utility) computing refers to creating infrastructure so that access to computing power is as transparent as access to electricity. Can we similarly use intelligent projectors to create a seemingly pervasive infrastructure so that displays of any size can be realized with minimal planning? The research challenge is how to create **Plug-and-disPlay projectors**. We describe our effort in creating self-contained projector systems that are aware of their surroundings via built-in sensors. The projectors can respond to the display context. Single units can recover 3D geometric and photometric information about the environment. Multiple, possibly heterogeneous, units can work in clusters and discover the topology. Such projectors allow self-configuring, seamless and large-area support for interaction and display. In the future, such aware projectors will work flexibly in a variety of situations, communicate with other entities and harmoniously make use of available resources.

2. Grid Architecture

What does it mean to create a projector-based display grid? Let us compare with traditional grids such as compute or power grid.

Grids consist of **clusters** of self-contained units. The approach is a shift away from large monolithic systems. For computing, the aim is to replace expensive mainframe computers, which are difficult to maintain given their specialized architecture, operating system and software. For projectors, our goal is to avoid large fixed tiled

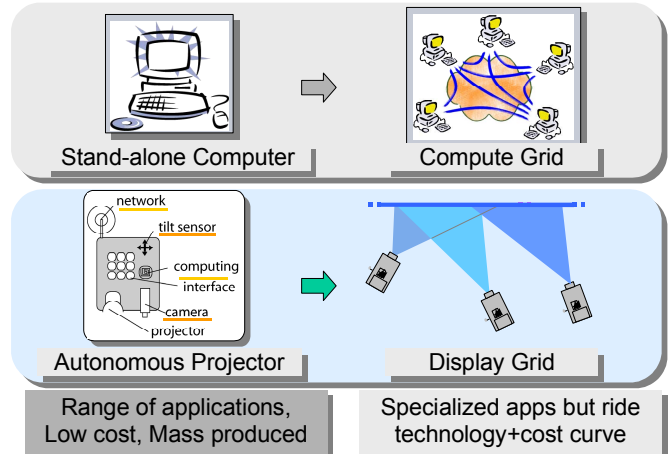


Figure 1. Projectors, like computers, are becoming commoditized. Compute grids allow blending heterogeneous devices into a single vast pool of resources for high-performance applications. In the future, can we do the same for Projector based Display Grids?

projector installations, which are governed by the rigid mechanical setup. Such installations freeze in time the invested resources. It is difficult to replace individual projectors with newer, smaller, brighter and more efficient projectors because the physical dimensions and optics of a given cube of the display are specific to a projector. The tiled display installation also requires frequent cumbersome geometric alignment. Even if the tiled display is modular, each module has a fixed opto-mechanical design locking the consumer into buying the same make of module. In addition, all the units need to have the same specifications in terms of resolution and brightness to create a uniform display, prohibiting heterogeneous mix and match of available units. We have already seen a movement toward mobile projection in single projector market. Since installing a projector is expensive, some medium-sized conference rooms now do not have a fixed installed projector. The presentations are conducted via a tabletop mobile projector. How can we take this further into multi-projector market and eliminate the need for large rigid installations?

The individual units in grids need to be **self-contained**. The unit, for example a computer or a smart projector, has to be useful even in isolation. This allows the device to have a range of applications. The mass appeal means the device is produced in quantity and the costs continue to decrease over time. Then such devices can be clustered for special applications. Over time the same cluster can be built via newer editions of the device. However, the self-contained unit needs to have computational and communication abilities to be able to sense and respond to its neighboring units.

Grids support **resource discovery** and management. This includes detection of network topology. In case of projector grids, in addition, the clusters need to determine the geometric neighborhood relationship. The topology is determined without involving a master controller or an entity outside the cluster. The communication takes

¹email: [raskar,jeroen,pab]@merl.com, <http://www.merl.com>

place via the network. In case of projector grids, the communication will also take place via light. The resource management includes allocation of image segments to corresponding projectors and efficient mechanism for data distribution.

Finally, grids allow a **dynamic entry** and exit of each unit, essentially a plug-and-play behavior. For projector grid, after the display is up and running, it should be possible to add (or remove) a projector to increase the effective size or brightness of the display.

3. Aware Projectors

What applications will be possible in a display grid? What problem do we need to solve to support those applications? What components do we need to include to make future projectors more intelligent?

We consider the following elements essential for geometric awareness – sensors such as camera and tilt-sensor, computing, storage, wireless communication and interface. The camera is used to observe (checkerboard) patterns projected by each projector and find 3D depth of the screen via triangulation. This data is used for alignment. The tilt sensor is used to make sure that the displayed image is level with the world horizontal and vertical. Note that the projector and these components can be combined in a single self-contained unit with just a single cable for power, or no cable at all with efficient batteries. Figure 2 illustrates a basic unit. This unit could be in a mobile form factor or could be a fixed projector.

Traditional tiled projector arrays create beautiful images at the cost of significant opto-mechanical infrastructure cost. While the projectors costs have come down to around US\$1000 to US\$3000 for XGA resolution, similar projectors modules in tiled arrays cost well over US\$10,000. Getting rid of the seams is still a big problem, as is the required frequent maintenance for alignment. Research efforts have produced prototypes that can create seamless displays. Existing work on projector clusters doing camera-based registration, such as [Raskar 2000; Yang et al. 2001; Chen et al. 2002], involves projection of patterns or texture onto the display plane, and measurement of homographies induced by the display plane. The homographies are used together with some Euclidean frame of reference to pre-warp images so that they appear geometrically registered and undistorted on the display. We improve on these techniques to allow the operation without environmental sensors and beyond the range of any one sensor.

The smart projector unit can communicate with other devices and objects to learn geometric relationships as required. The ability to learn these relationships on the fly is a major departure from most existing projector-based systems that involve a pre-configured geometric setup or, when used in flexible environments, involve detailed calibration, communication and human aid. As discussed later, even existing systems that use a simple planar homography and avoid complete calibration require some Euclidean information on the screen (e.g., screen edges or markers) or assume the camera is in the ideal sweet-spot position.

Previously, creating wide aspect ratios using camera-based registration of projector clusters has been a problem. We are able to overcome this problem because a single master camera sensor is not required and we use a new global alignment strategy that relies on pair-wise homographies between a projector of one unit and the camera of the neighboring unit. Figure 4 shows a heterogeneous cluster of five units, displaying seamless images after accurate geometric registration. The pair-wise homographies are used to compute a globally consistent set of homographies by solving a

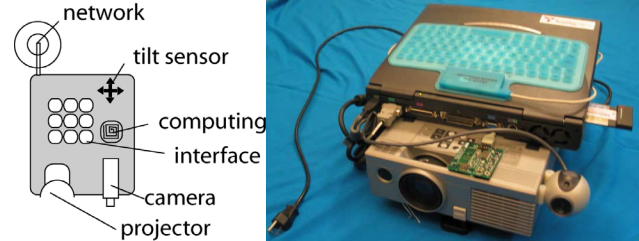


Figure 2. Self-contained intelligent projector unit, which can communicate with other units via both, light and RF, to learn geometric relationships and topology. (a) Components of a single unit and (b) our prototype with a single external cable for power.

linear system of equations. Figure 6 is a close-in view demonstrating the good quality of the resulting registration.

4. Self-organizing Projector Clusters

From a geometric point of view, camera-assisted multi-projector display systems are based on the notion of one or more environmental sensors assisting a central intelligent device. This central hub computes the Euclidean or affine relationships between projector(s) and displays. In contrast, our system is based on autonomous units, similar to self-contained computing units in cluster computing. In the proposed approach, each individual unit senses its geometric context within the cluster. This can be useful in many applications. For example, the geometric context can allow each projector to determine its contribution when creating a large area seamless display. Multiple units can also be used in the shape and object-adaptive projection systems for projector-based augmented reality.

This approach to display allows very wide aspect ratios, short throw distance between projectors and the display surfaces and hence higher pixel resolution and brightness, and the ability to use heterogeneous units. An ad-hoc cluster also has the advantages that it (a) operates without a central commanding unit, so individual units can join in and drop out dynamically, (b) does not require environmental sensors, (c) displays images beyond the range of any single unit, and (d) provides a mechanism for bypassing the limits on illumination from a single unit by having multiple overlapping projections.

Similar to computing units in a computing cluster, projector units can dynamically enter and leave a display cluster, and the alignment operations are performed without requiring significant pre-planning or programming. This is possible because (a) every unit acts

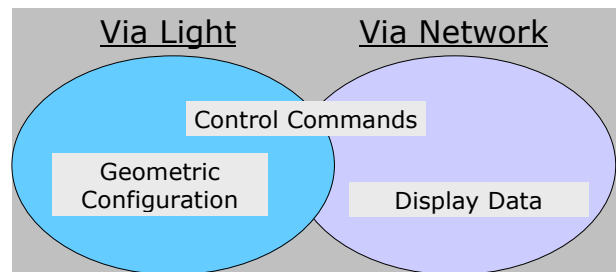


Figure 3. The communication of smart projectors occurs not just via wireless network, but also via light.

independently and performs its own observations and calculations, in a symmetric fashion (b) no Euclidean information needs to be fed to the system (such as corners of the screen or alignment of the master camera), because tilt-sensors and cameras allow each projector to be geometrically aware. In contrast to our approach, systems with centralized operation for multi-projector display quickly become difficult to manage.

5. Ad-hoc Operation

The ideas of a display grid provides geometric underpinnings for a new generation of projectors – autonomous devices, easily adapting to operation within a cluster, and adaptive to their surroundings. The approach is described below in the context of creating a large planar display. Imagine a collection of projectors distributed in a building. The proximity of projectors that may work together to create a large display is detected via light and via wireless network (typically short range radio frequency or infrared). A *group* of projectors display a seamless image, but there may be more than one group in the *vicinity*. For example, projectors in the same room maybe in the same vicinity within RF range. But if they are pointing in opposite directions, they cannot form a group.

Joining a group When a unit, U_k , containing a projector, P_k , plus a camera, C_k , wants to join a group, it informs the group in two ways. Over the proximity network (such as wireless Ethernet, RF or infrared) it sends a ‘request to join’ message with its own unique id, which is received by all the m units, U_i for $i = 1..m$, in the vicinity. This puts the cameras, C_i for $i = 1..m$, of all the units in attention mode and the units respond with ‘ready’ message to U_k . The second form of communication occurs via light. Unit U_k projects a structured pattern, which may interrupt the display and is observed by all the m cameras embedded in the units. If any one camera from the existing group views the projected pattern, the whole groups moves onto a quick calibration step to include P_k in their display. Otherwise, the group assumes that U_k is in the vicinity but does not overlap with its own extent of the display. Without a loss of generality let us assume that the first n units now form a group.

Pairwise Geometric affine relationship A well-known method to register overlapping projectors is to express the relationship using a homography matrix. The mapping between two arbitrary perspective views of an opaque planar surface in 3D can be expressed using a planar projective transfer, expressed as a 3×3 matrix defined up to a scale. The 8 degrees of freedom can be

computed from four or more correspondences in the two views. In addition, due to the linear relationship, homography matrices can be cascaded to propagate the image transfer.

Unit U_k directs, using wireless communication, each projector, P_i for $i = 1..n$, in the group to project a structured pattern (a uniform checkerboard), one at a time. Projection is simultaneously viewed by the camera of each unit in the group. This creates pairwise homographies for transferring the image of projector P_i into image in camera C_j . We calculate homography for each pairwise projector pair, cascading two such homographies.

Global Alignment After pairwise homographies are exchanged, each unit performs the global adjustment separately (in parallel). In the absence of environmental sensors, we compute the relative 3D pose between the screen and all the projectors to allow a seamless display. Without a known pose, the computed solution is correct up to a transformation by a homography and will look skewed on screen (Figure 5). We use the tilt sensor readings in each unit so that, if the screens are vertical planes, our approach automatically aligns the projected image with the world horizontal and vertical.

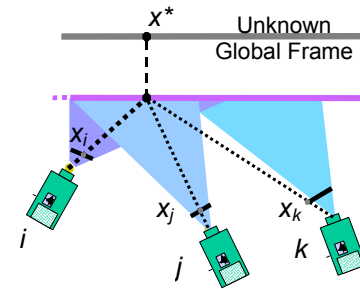


Figure 5. Dynamic entry of Unit k in a display of two units, i and j . The image of a fixed point on screen in each projector, e.g. X_i and X_j are related by a homography.

In the work here, the full size image is multi-cast to each projector unit, of which each unit displays an appropriate part. If bandwidth were a limiting constraint for high frame rate applications, one would want to decompose content and transmit to an individual projector only the pixels that it requires. We texture map the input




Figure 4. Application of self-configuring projectors in building wide aspect ratio displays. Every unit is **self-contained**. We use a modified global alignment scheme based on local pairwise relationships. This replaces existing techniques that require the notion of a master camera and Euclidean information for the scene. Left-top: Uncorrected projection from each of the five projectors; Left-bottom: Registered images; Right: Setup of self-contained units and seamless display.




Figure 6. (Left) Visual indication of registration accuracy (Right) Intensity blending map.

image onto a unit rectangle (of correct aspect ratio) and render with a projection matrix derived from the homography [Raskar 2000]. The intensity blending is implemented using the alpha-channel in graphics hardware. The blending weights are assigned proportional to the distance to the image boundary.

6. Results

Our prototype is a box that includes an XGA (1024x768) Mitsubishi X80 projector, a Dell Inspiron laptop with ATI Radeon graphics board, a Logitech USB camera (640x480), a tilt-sensor by Analog Devices ADXL202A (angular resolution about 1 degree), and a wireless LAN card. The box is closed with two circular holes in the front, one for the projector and one for the camera. Other components inside the box are a power supply, cooling fan, and a numerical keypad. The total cost per box is about US\$5000, \$3200 for the projector, \$1200 for the laptop, \$50 for camera, \$100 for tilt sensor evaluation board. The only cable coming out of the box is for power. We also built prototypes that have the same components but not put in a box. It should be noted that our goal is to demonstrate the concept of an ad-hoc grid. Many of the bandwidth, resolution and computation constraints will become less significant in the future. e.g. with new 802.11 protocols, LCOS imagers and shrinking faster processors. Consumers may accept decreased resolution due to image warping for mechanical simplicity and cost.



Figure 7. A cluster of enhanced projectors.

All the computations described above are performed symmetrically, in parallel on each unit. After casual installation or when a new projector joins the group, each unit projects a checkerboard that is observed by all the cameras. It takes about 5 seconds per projector to project the pattern and find pairwise homographies. Global alignment, inscribed rectangle and blending weights computations take an additional 3 seconds. For a six projector unit setup, the total time to display a seamless image after casual installation is no more than 35 seconds.

We show several demonstrations in the video available on our website, (a) various projector configurations, (b) very wide aspect ratio, (c) high-brightness display by superimposition of projected texture.

We also show a close-up to demonstrate the accuracy of geometric registration. In this paper, we are not discussing the photometric correction aspects or issues of image quality due to warping. The demonstration shows a front projection setup, but the ideas can be used in a rear-projection setup with no modifications.

7. Impact

After projectors become cheap and commoditized, consumers will start using multi-projector displays in casual applications at home to watch movies, at work for low-cost ad-hoc visualization, in game arcades and in retail stores for advertising. These environments do not permit large rigid installations with significant maintenance costs. We believe, in the future, the software intelligence and embedded sensing will become critical to support such casual use. A cluster-based approach, compared to building massive displays, allows consumers to ride the cost and technology curve. Projector manufacturers and display integrators will try to take advantage of parallel trends in cluster computing and sensor networks. They will build modular pieces that will support inter-device communication and allow the same device (a projector) to be used in a multi-purpose fashion.

8. Conclusion

We have presented a vision of a display grid built with autonomous networked and intelligent projector units. We have demonstrated working prototypes and techniques to support the main idea - like the transition from mainframes to compute clusters, in the future, large projector installations will be replaced by self-configuring projector clusters.

Please see the video demonstrating our system at <http://www.merl.com/people/raskar/SID04/>

9. Acknowledgements

We thank Thomas Willwacher, Srinivas Rao, Cliff Forlines, Bill Yerazunis, Darren Leigh and Paul Dietz, Joe Marks as well as Yoshihiro Ashizaki, Masatoshi Kameyama, and Keiichi Shiotani.

References

1. Barco Display Systems, <http://www.barco.com>
2. Chen, H., Sukthankar, R., Wallace, G., Li, K. Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. IEEE Visualization, 2002.
3. Raskar, R., van Baar, J., Beardsley, P., Willwacher, T., Rao, S., Forlines, C., iLamps: Geometrically Aware and Self-Configuring Projectors, Proceedings of SIGGRAPH 2003
4. Raskar, R.; van Baar, J.; Chai, J.X., A Low-Cost Projector Mosaic with Fast Registration, ACCV 2002
5. Raskar, R. 2000. Immersive Planar Display using Roughly Aligned Projectors. IEEE VR 2000
6. Seos Reality Rooms <http://www.seos.com/reality.htm>
7. Yang, R., Gotz, D., Hensley, J., Towles, H., Brown, M. PixelFlex: A Reconfigurable Multi-Projector Display System. In IEEE Visualization 2001