iLamps: Geometrically Aware and Self-Configuring Projectors

Ramesh Raskar, Jeroen van Baar, Paul Beardsley, Thomas Willwacher, Srinivas Rao, Clifton Forlines

TR2003-23 July 2003

Abstract

Projectors are currently undergoing a transformation as they evolve from static output devicesand applications.

SIGGRAPH 2003

iLamps: Geometrically Aware and Self-Configuring Projectors

Ramesh Raskar Jeroen van Baar Paul Beardsley Thomas Willwacher Srinivas Rao Clifton Forlines
Mitsubishi Electric Research Labs (MERL), Cambridge MA, USA [*]

Abstract

Projectors are currently undergoing a transformation as they evolve from static output devices to portable, environment-aware, communicating systems. An enhanced projector can determine and respond to the geometry of the display surface, and can be used in an ad-hoc cluster to create a self-configuring display. Information display is such a prevailing part of everyday life that new and more flexible ways to present data are likely to have significant impact. This paper examines geometrical issues for enhanced projectors, relating to customized projection for different shapes of display surface, object augmentation, and co-operation between multiple units.

We introduce a new technique for adaptive projection on non-planar surfaces using conformal texture mapping. We describe object augmentation with a hand-held projector, including interaction techniques. We describe the concept of a display created by an ad-hoc cluster of heterogeneous enhanced projectors, with a new global alignment scheme, and new parametric image transfer methods for quadric surfaces, to make a seamless projection. The work is illustrated by several prototypes and applications.

CR Categories: B.4.2 [Input/output and Data Communications]: Input/Output Devices—Image display ; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Imaging geometry

Keywords: projector, calibration, seamless display, augmented reality, ad-hoc clusters, quadric transfer.

1 Introduction

Traditional projectors have been static devices, and typical use has been presentation of content to a passive audience. But ideas have developed significantly over the past decade, and projectors are now being used as part of systems which sense the environment. The capabilities of these systems range from simple keystone correction to augmentation overlay on recognized objects, including various types of user interaction.

Most such systems have continued to use static projectors in a semi-permanent setup, one in which there may be a significant calibration process prior to using the system. Often too the systems are specifically designed for a particular configuration. But the increasing compactness and cheapness of projectors is enabling much more flexibility in their use than is found currently. For example, portability and cheapness open the way for clusters of projectors which are put into different environments for temporary deployment, rather than a more permanent setup. As for hand-held use, projectors look like a natural fit with cellphones and PDAs. Cellphones provide access to the large amounts of wireless data which surround us, but their size dictates a small display area. An attached projector can maintain compactness while still providing a reasonably-sized display. A hand-held cellphone-projector becomes a portable and easily-deployed information portal.

These new uses will be characterized by opportunistic use of portable projectors in arbitrary environments. The research challenge is how to create *Plug-and-disPlay* projectors which work flexibly in a variety of situations. This requires generic application-independent components, in place of monolithic and specific solutions.

This paper addresses some of these new problems. Our basic unit is a projector with attached camera and tilt-sensor. Single units can recover 3D information about the surrounding environment, including the world vertical, allowing projection appropriate to the display surface. Multiple, possibly heterogeneous, units are deployed in clusters, in which case the systems not only sense their external environment but also the cluster configuration, allowing self-configuring seamless large-area displays without the need for additional sensors in the environment. We use the term *iLamps* to indicate intelligent, locale-aware, mobile projectors.

1.1 Overview

The focus of this paper is geometry. Successive sections address issues about the geometry of display surfaces, 3D motion of a hand-held projector, and geometry of a projector cluster. Specifically, we make the following contributions –

Shape-adaptive display: We present a new display method in which images projected on a planar or non-planar surface appear with minimum local deformation by utilization of conformal projection. We present variations to handle horizontal and vertical constraints on the projected content.

Object-adaptive display: We demonstrate augmentation of objects using a hand-held projector, including interaction techniques.

Planar display using a cluster of projectors: We present algorithms to create a self-configuring ad-hoc display network, able to create a seamless display using self-contained projector units and without environmental sensors. We present a modified global alignment scheme, replacing existing techniques that require the notion of a master camera and Euclidean information for the scene.

Curved display using a cluster of projectors: We extend planar surface algorithms to handle a subset of curved surfaces, specifically quadric surfaces. We introduce a simplified parameterized transfer equation. While several approaches have been proposed for seamless multi-projector planar displays, as far as we know, literature on seamless displays is lacking in techniques for parameterized warping and registration for curved screens.

We omit discussion of photometric issues, such as the interaction between color characteristics of the projected light [Majumder et al.

[*] email:[raskar, jeroen, pab, willwach, raos, forlines]@merl.com
http://www.merl.com/projects/Projector/

2000] and environmental characteristics like surface reflectance, orientation, and ambient lighting. We also omit discussion about non-centralized cluster-based systems and issues such as communication, resource management and security [Humphreys et al. 2001; Samanta et al. 1999]. Finally a full discussion of applications is outside the scope of the paper, though we believe the ideas here will be useful in traditional as well as new types of projection systems.

1.2 Evolution of Projectors

Projectors are getting smaller, brighter, and cheaper. The evolution of computers is suggestive of the ways in which projectors might evolve. As computers evolved from mainframes to PCs to handheld PDAs, 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 systems towards ad-hoc, self-configuring displays made up of heterogeneous, self-sufficient projector units.

The most exploited characteristic of projectors has been their ability to generate images that are larger in size than their CRT and LCD counterparts. But the potential of other characteristics unique to projector-based displays is less well investigated. Because the projector is decoupled from the display (i) the size of the projector can be much smaller than the size of the image it produces, (ii) overlapping images from multiple projectors can be effectively superimposed on the display surface, (iii) images from projectors with quite different specifications and form factors can be easily blended together, and (iv) the display surface does not need to be planar or rigid, allowing us to augment many types of surfaces and merge projected images with the real world.

1.3 Relevant Work

The projector's traditional roles have been in the movie, flight simulator, and presentation markets, but it is now breaking into many new areas.

Projector-based environments There are many devices that provide displays in the environment, and they are becoming more common. Some examples are large monitors, projected screens, and LCD or plasma screens for fixed installations, and hand-held PDAs for mobile applications. Immersion is not a necessary goal of most of these displays. Due to shrinking size and cost, projectors are increasingly replacing traditional display mediums. We are inspired by projector-based display systems that go beyond the traditional presentation or multi-projector tiled displays: Office of the Future [Raskar et al. 1998], Emancipated Pixels [Underkoffler et al. 1999], Everywhere Display [Pinhanez 2001] and Smart Presentations [Sukthankar et al. 2001]. Many new types of projector-based augmented reality have also been proposed [Raskar et al. 2001; Bimber et al. 2002]. From a geometric point of view, these 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 (or ubiquitous computing). In the last four years, many authors have proposed automatic registration for seamless displays using a cluster of projectors [Yang et al. 2001; Raskar et al. 2002; Chen et al. 2002; Brown and Seales 2002]. We improve on these techniques to allow the operation without environmental sensors and beyond the range of any one sensor. We also extend the cluster based approach to second-order display surfaces.

Enhanced projectors Another related area of research is the enhancement of projectors using sensors and computation. Underkoffler et al. [1999] described an I/O bulb (co-located projector and camera). Hereld et al. [2000] presented a smart projector with an attached camera. Raskar and Beardsley [2001] described a geometrically calibrated device with a rigid camera and a tilt sensor to allow automatic keystone correction. Many have demonstrated user interactions at whiteboards and tabletop surfaces including the use of gestural input and recognition of labeled objects [Rekimoto 1999; Rekimoto and Saitoh 1999; Crowley et al. 2000; Kjeldsen et al. 2002]. We go further, by adding network capability, and by making the units self-configuring and geometrically-aware. This allows greater portability and we investigate techniques which anticipate the arrival of hand-held projectors.

Grids Our approach to ad-hoc projector clusters is inspired by work on grids such as ad-hoc sensor networks, traditional dynamic network grids, and ad-hoc computing grids or network of workstations (NoW, an emerging technology to join computers into a single vast pool of processing power and storage capacity). Research on such ad-hoc networks for communication, computing and datasharing has generated many techniques which could also be used for context-aware 'display grids'.

2 Geometrically Aware Projector

What components will 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. 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 1 illustrates a basic unit. This unit could be in a mobile form factor or could be a fixed projector. Because we do not wish to rely on any Euclidean information external to the device (e.g., markers in the room, boundaries on screens, or human aid), we use a completely calibrated projector-camera system.

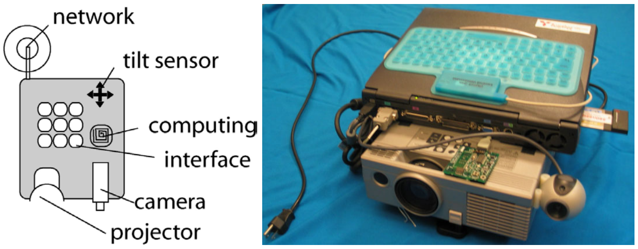

Figure 1: *Our approach is based on self-contained iLamps. Left: components of enhanced projector; Right: our prototype, with a single power cable.*

In isolation, the unit can be used for several applications including (i) smart keystone correction (ii) orientation compensated intensities (iii) auto brightness, zoom, focus (iv) 3D scanner for geometry and texture capture (with auto zippering of piecewise reconstructions by exploiting the camera with accelerometer) (v) smart flash for cameras, with the projector playing a secondary role to provide intelligent patterns or region specific lighting.

The 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 preconfigured geometric setup or, when used in flexible environments, involve detailed calibration, communication and human aid. 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) [Sukthankar et al. 2001] or assume the camera is in the ideal sweet-spot position [Yang et al. 2001; Raskar et al. 2002; Brown and Seales 2002].

3 Shape-adaptive Display

When using projectors casually and portably, an ideal planar display surface is not always available, and one must take advantage of other surfaces such as room corners, columns, or oddly shaped ceilings. The shape-adaptive display in this section emulates existing examples of texture on curved surfaces, such as large advertisements and news tickers on curved displays, and product labels on curved containers. The issue is how to generate images that appear 'correctly' to multiple simultaneous viewers. This is a different problem to pre-warping an input image so that it appears perspectively correct from a *single* sweet-spot location [Raskar et al. 1999]. Human vision interprets surface texture in the context of all three-dimensional cues – when viewing a poster on the wall from one side, or reading the label of a cylindrical object such as a wine bottle, or viewing a mural on a curved building. The goal therefore is to create projected texture which is customized to the shape of the surface, to be consistent with our usual viewing experience.

3.1 Conformal Projection

This section describes how to display an image that has minimum stretch or distortion over the illuminated region. Consider first a planar surface like a movie screen – the solution is to project images as if the audience is viewing the movie in a fronto-parallel fashion, and this is achieved by keystone correction when the projector is skewed. Now consider a curved surface or any non-planar surface in general. Intuitively, we wish to 'wallpaper' the image onto the display surface, so that locally each point on the display surface is undistorted when viewed along the surface normal.

Since the normal may vary, we need to compute a map that minimizes distortion in some sense. We chose to use conformality as a measure of distortion. A conformal map between the input image and the corresponding areas on the display surface is angle preserving. A scroll of the input image will then appear as a smooth scroll on the illuminated surface, with translation of the texture but no change in size or shape.

A zero-stretch solution is possible only if the surface is developable. Example developable surfaces are two planar walls meeting at a corner, or a segment of a right cylinder (a planar curve extruded perpendicular to the plane). In other cases, such as three planar walls meeting in a corner, or a partial sphere, we solve the minimum stretch problem in the least squares sense. We compute the desired map between the input image and the 3D display surface using the least squares conformal map (LSCM) proposed in Levy et al. [2002]. An example of texture projection using this approach is shown in Figure 2.

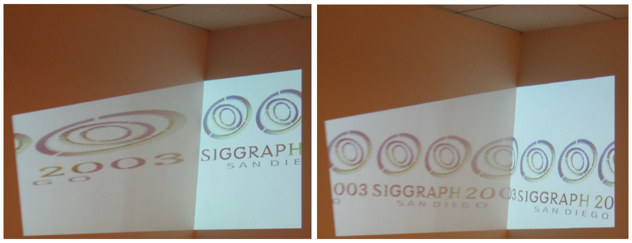

Figure 2: *Shape-adaptive projection. Left: the projector is skew relative to the left wall so direct, uncorrected projection of texture gives a distorted result; Right: the projector still in the same position but use of LSCM removes the distortion in the projected image.*

LSCM minimizes angle deformation and non-uniform scaling between corresponding regions on a 3D surface and its 2D parameterization space, the solution being fully conformal for a developable surface. For a given point X on the 3D mesh, if the 2D texture coordinates (u,v) are represented by a complex number $(u+iv)$ and the display surface uses coordinates in a local orthonormal basis $(x+iy)$, then the goal of conformal mapping is to ensure that tangent vectors to the iso-u and iso-v curves passing through X are orthogonal and have the same norm, i.e.,

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y} \qquad\qquad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}$$

In Levy et al. [2002], the authors solve this problem on a per triangle basis and minimize the distortion by mapping any surface homeomorphic to a disk to a (u,v) parameterization. The steps of our algorithm are as follows.

1. Project structured light from the projector, capture images with a rigidly attached calibrated camera, and create a 3D mesh D of the surface.
2. Use LSCM to compute texture coordinates U of D, thereby finding a mapping D^{Π} of D in the (u,v) plane.
3. Find the displayable region in D^{Π} that (a) is as large as possible and (b) has the vertical axis of the input image aligned with the world vertical. The method for determining the rotation between the input image and the world vertical is described below.
4. Update U into U' to correspond to the displayable region.
5. Texture-map the input image onto the original mesh D, using U' as texture coordinates, and render D from the viewpoint of the projector.

3.2 Vertical Alignment

The goal of vertical alignment is to ensure the projected image has its vertical direction aligned with the world vertical. There are two cases – (a) if the display surface is non-horizontal, the desired texture vertical is given by the intersection of the display surface and the plane defined by the world vertical and the surface normal, (b) if the display surface is horizontal, then the texture vertical is undefined. Regarding condition (b), the texture orientation is undefined for a single horizontal plane, but given any non-horizontal part on the display surface this will serve to define a texture vertical which also applies to horizontal parts of the surface.

The update of U into U' involves a rotation, R, for vertical alignment, in addition to a scale and shift in the (u,v) plane. If the computed 3D mesh were perfect, the computation of R could use a single triangle from the mesh. But the 3D data is subject to error, so we employ all triangles in a least-squares computation. The approach is as follows.

1. For each non-horizontal triangle, t_j, in the 3D mesh D, (i) Compute the desired texture vertical as the 3D vector $p_j = n \times (v \times n)$, where n is the surface normal of the triangle and v is the world vertical (obtained from the tilt-sensor), and $\times$ is the cross-product operator, (ii) Use the computed LSCM to transform p_j into normalized vectors $q_j = (u_j, v_j)$ in the (u,v) plane.
2. Find the rotation which maximizes the alignment of each q_j with direction (0,1) – compute the covariance matrix $M = \sum_j [u_j \quad v_j]^T [0 \quad 1]$, find the singular value decomposition of M as TSV^T, and compute the desired 2D rotation as $R = TV^T$.

3.3 Shape Constraints

It is sometimes desirable to constrain the shape of projected features in one direction at the cost of distortion in other directions. For example, banner text projected on a near-vertical but non-developable surface such as a sphere-segment should appear with all the text characters having the same height, even if there is distortion in the horizontal direction. Additional constraints on the basic four partial derivatives in LSCM are obtained by introducing equations of the form $\lambda_{vert} \cdot (\frac{\partial v}{\partial y} - const) = 0$. Typically, only one such equation will be used. The equation above, for example, keeps stretch along the vertical direction to a minimum, i.e., it penalizes and minimizes the variance in $\partial v / \partial y$ over all triangles. This modification also requires that the local orthonormal x, y-basis on the triangles is chosen appropriately – in this case, the x-axis must point along the horizontal everywhere on the surface. Figure 3 shows an example.

Results Surfaces used for conformal display, shown here and in the accompanying materials, include a two-wall corner, a concertina-shaped display, and (as an example of a non-developable surface) a concave dome.

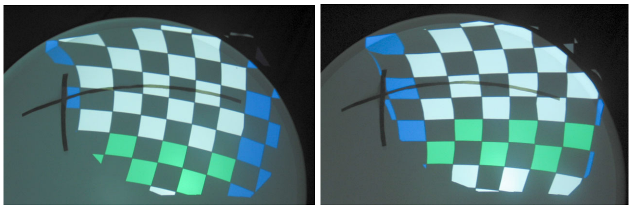

Figure 3: *Left: uncorrected projection from a skew projector; Right: correction of the texture using constrained LSCM. Observe the change in the area at upper-left. Image is world horizontal aligned. Vertical stretch is minimized (at the cost of horizontal distortion) so that horizontal lines in the input texture remain in horizontal planes.*

4 Object-adaptive Display

This section describes object augmentation using a hand-held projector, including a technique for doing mouse-style interaction with the projected data. Common to some previous approaches, we do object recognition by means of fiducials attached to the object of interest. Our fiducials are ’piecodes’, colored segmented circles like the ones in Figure 4, which allow thousands of distinct color-codings. As well as providing identity, these fiducials are used to compute camera pose (location and orientation) and hence projector pose since the system is fully calibrated[1]. With projector pose known relative to a known object, content can be overlaid on the object as required.

Advantages of doing object augmentation with a projector rather than by annotated images on a PDA include (a) the physical size of a PDA puts a hard limit on presented information; (b) a PDA does augmentation in the coordinate frame of the camera, not the user’s frame, and requires the user to context-switch between the display and physical environment; (c) a PDA must be on the user’s person while a projector can be remote; (d) projection allows a shared experience between users. Eye-worn displays are another important augmentation technique but they can cause fatigue, and there are stringent computational requirements because of the tight coupling between user motion and the presented image (e.g., a user head rotation must be matched precisely by a complementary rotation in the displayed image). Projection has its own disadvantages – it is poor on dark or shiny surfaces, and can be adversely affected by ambient light; it does not allow private display. But a key point is that projector-based augmentation naturally presents to the user’s own viewpoint, while decoupling the user’s coordinate frame from the processing. This helps in ergonomics and is easier computationally.

A hand-held projector can use various aspects of its context when projecting content onto a recognized object. We use proximity to the object to determine level-of-detail for the content. Other examples of context for content control would be gestural motion, history of use in a particular spot, or the presence of other devices for cooperative projection. The main uses of object augmentation are (a) information displays on objects, either passive display, or training applications in which instructions are displayed as part of a sequence (Figure 4(top)); (b) physical indexing in which a user is guided through an environment or storage bins to a requested object (Figure 4(bottom)); (c) indicating electronic data items which have been attached to the environment. Related work includes the Magic Lens [Bier et al. 1993], Digital Desk [Wellner 1993], computer augmented interaction with real-world environments [Rekimoto and Nagao 1995], and Hyper mask [Yotsukura et al. 2002].

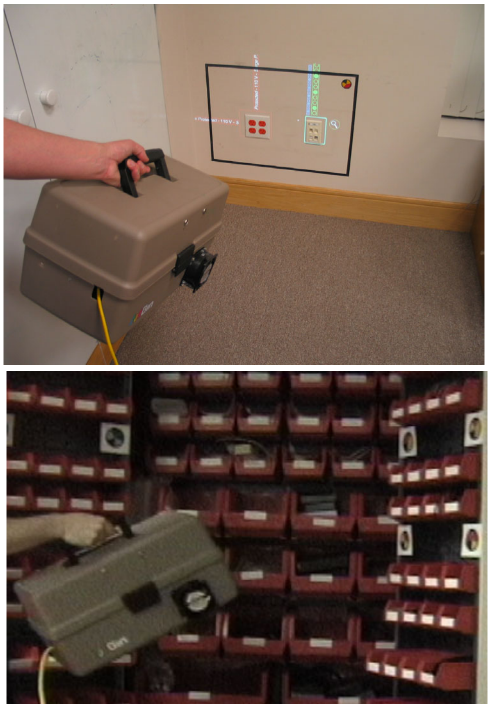

Figure 4: *Context-aware displays. Top: augmentation of an identified surface; Bottom: guidance to a user-requested object in storage bins.*

Mouse-Style Interactions with Augmentation Data. The most common use of projector-based augmentation in previous work has been straightforward information display to the user. A hand-held projector has the additional requirement over a more static setup that there is fast computation of projector pose, so that the augmentation can be kept stable in the scene under user motion. But a hand-held projector also provides a means for doing mouse-style interactions – using a moving cursor to interact with

[1] We use four coplanar points in known position in a homography-based computation for the pose of the calibrated camera. The points are obtained from the segments of a single piecode, or from multiple piecodes, or from one piecode plus a rectangular frame. For good results, augmentation should lie within or close to the utilized points.

the projected augmentation, or with the scene itself.

Consider first the normal projected augmentation data – as the projector moves, the content is updated on the projector's image plane, so that the projected content remains stable on the physical object. Now assume we display a cursor at some fixed point on the projector image plane, say at the center pixel. This cursor will move in the physical scene in accordance with the projector motion. By simultaneously projecting the motion-stabilized content and the cursor, we can emulate mouse-style interactions in the scene. For example, we can project a menu to a fixed location on the object, track the cursor to a menu item (by a natural pointing motion with the hand-held projector), and then press a button to select the menu item. Alternatively the cursor can be used to interact with the physical scene itself, for example doing cut-and-paste operations with the projector indicating the outline of the selected area and the camera capturing the image data for that area. In fact all the usual screen-based mouse operations have analogs in the projected domain.

5 Cluster of Projectors

The work so far has been on individual projector units. This section deals with ad-hoc clusters of projector units. 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 described above.

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.

These concepts are shown working in the context of planar display, and also for higher order surfaces, such as quadric surfaces. For the latter, we present a new image transfer approach. In the work here, each projector unit has access to the same full-size image, of which it displays an appropriate part. If bandwidth were an important constraint, one would want to decompose content and transmit to an individual projector only the pixels which it requires, but that topic is not discussed.

5.1 Planar Display using Ad-Hoc Clusters

This section deals with a cluster projecting on the most common type of display surface, a plane. Existing work on projector clusters doing camera-based registration, such as [Raskar et al. 2002; Brown and Seales 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.

However, creating wide aspect ratios 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 5 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 linear system of equations. Figure 6(left) is a close-in view demonstrating the good quality of the resulting registration.

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 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.

The approach is described below in the context of creating a large planar display . A *group* of projectors display a seamless image, but there may be more than one group in the *vicinity*.

Joining a group When a unit, U_k, containing a projector, P_k, and 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. 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 3x3 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 H_{PiCj} for transferring the image of projector P_i into image in camera C_j.

We calculate pairwise projector homography, H_{PiPj}, indirectly as $H_{PiCi}H_{PjCi}{}^{-1}$. For simplicity, we write H_{PiPj} as H_{ij}. In addition, we store a confidence value, h_{ij}, related to the percentage of overlap in image area and it is used during global alignment later. Since we use a uniform checkerboard pattern, a good approximation for the overlap percentage is the ratio r_{ij} of (the number of features of projector P_j seen by camera C_i) to (the total number of features projected by P_j). We found confidence $h_{ij} = r_{ij}{}^4$ to be a good metric. The value is automatically zero if the cameras i did not see the projector j.

Global Alignment 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 distorted on screen. Further, if the screens are vertical planes, our approach automatically aligns the projected image with the world horizontal and vertical.

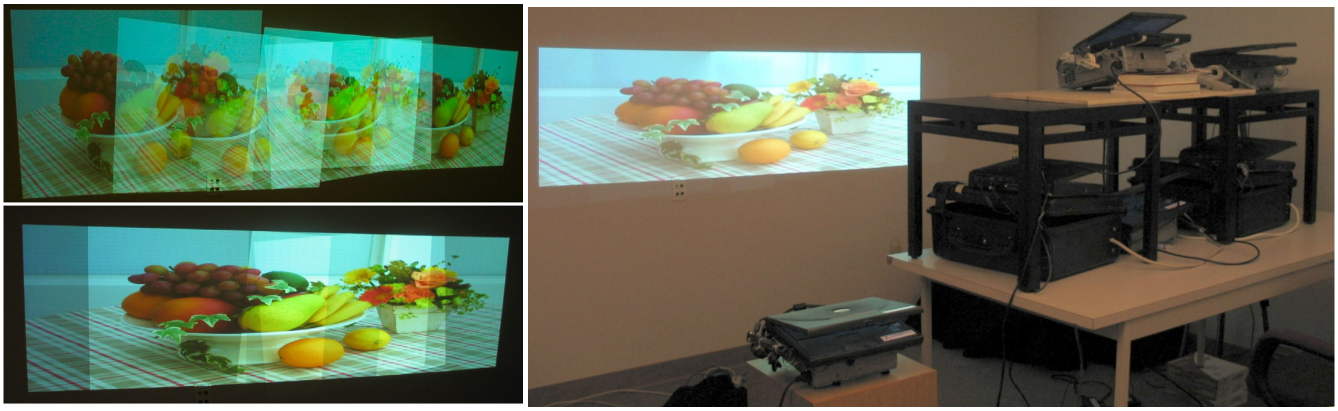

Figure 5: *Application of self-configuring projectors in building wide aspect ratio displays. Left-top: Uncorrected projection from each of the five projectors; Left-bottom: Registered images; Right: Setup of self-contained units and seamless display.*

However, to compute the 3D pose for global alignment, we avoid low confidence information such as stereo calibration and tilt-sensor parameters of units, and exploit relatively robust image space calculations such as homographies. Thus, we first compute globally consistent homographies. A set of homography transforms cascaded along a cycle $\prod_{i=0}^{k-1} H_{(i+1\%k)(i)}$ should equal identity. This is seldom the case due to feature location uncertainty. Thus the goal is to create a meaningful mosaic in the presence of minor errors. A simple solution is to use a single camera viewing the display region [Yang et al. 2001; Chen et al. 2000] and perform computation in that space. When a single camera cannot view the whole display, Chen et al. [2002] use a minimum spanning tree-based approach. The graph nodes are projectors which are connected to overlapping projectors by edges. The solution keeps only a minimal set of homographies. By using a projector near the center of the display as the root node of the tree, this approach reduces the number of cascaded homographies and thereby lowers the cumulative error.

We instead present a simple and efficient method for finding a globally consistent registration by using information from all the available homographies *simultaneously*. The basic approach is based on a scheme proposed by [Davis 1998]. Cumulative errors due to homography cascading are concurrently reduced with respect to a chosen reference frame (or plane) by solving a linear system of equations. The linear equations are derived from independently computed pairwise homography matrices. We modify the method to consider the confidence h_{ij} in measured pairwise projector homography H_{ij}.

A global alignment homography G_i transfers image I_i in P_i into the global reference frame. But it can also be written as a transfer of I_i to I_j with H_{ij} followed by transfer to global frame with G_j.

$$G_j H_{ij} \cong G_i$$

Thus, we can build a sparse linear system of equations, where each H_{ij} is known, to compute the unknown global transfers. Since, each homography is defined up to a scale factor, it is important that we normalize so we can directly minimize $|G_j H_{ij} - G_i|$. The determinant of normalized homography, $\hat{H}$ is 1 i.e., $\hat{H} = H/(det|H|)^{1/3}$. We omit the hat for simplicity. We further scale individual linear equations by the confidence in H_{ij} computed as h_{ij} above. Thus, the set of equations are, for each projector pair (i, j),

$$h_{ij}(G_j H_{ij} - G_i) = 0$$

If the number of independent pairwise homographies are larger than the number of projector units, the system of linear equations is over constrained and we solve it in a least squares sense to produce a set of global registration matrices. After pairwise homographies are exchanged, each unit performs the global adjustment separately (in parallel) by treating its own image as the reference plane. Since the projector image plane is not aligned with the screen, at this stage the solution for global image transfer is known up to an unknown homography, H_{i0}, between each projector P_i and the screen.

Euclidean reconstruction In the absence of Euclidean information in the environment, we find H_{i0} by computing the 3D pose of P_j with respect to the screen. Although pose determination in 3D is noisy, our image transfer among projectors remains the same as before, i.e., based on global transfer ensuring a seamless image. Our approach to compute H_{i0} is based on [Raskar and Beardsley 2001]. Note that this computation is performed in parallel by each unit U_i. The steps are as follows.

1. Triangulate corresponding points in P_i and C_i image space to generated 3D points $X = \{x1, x2..\}$.
2. Find a best fit plane, Π, for X. Assign a local coordinate system with the plane Π as the *x*-*y* coordinate plane Find rotation R^{π}, between local and projector coordinate system.
3. Apply the rotation due to tilt, $X_0 = R^{tilt} R^{\pi} X$, so that X_0 is now world and screen aligned i.e., z coordinate is 0, points in a world vertical plane have the same x coordinate and points in a world horizontal plane have the same y coordinate.
4. Compute homography, H_{i0} , mapping the image of X in P_i framebuffer to the corresponding 2D points in *x*-*y* plane for X_0 (ignoring *z*-value)
5. Finally for all other projectors, $H_{j0} = H_{i0} H_{ji}$, where H_{ji} is the global homography G_j computed by unit *i*.

Display Finding a suitable projection matrix for rendering for a projector unit involves first computing the region of contribution of the projector on the screen and then re-projecting that complete display extent into the projector image.

In the first step, we find the extent of the illuminated quadrilateral in screen-aligned (X_0) coordinates by transforming the four corners of the normalized projector image plane $[\pm 1 \pm 1]$ using $H_{j0}[\pm 1 \pm 1\ 1]^T$. The union of all projected quadrilaterals in screen X_0 space is a (possibly concave) planar polygon, *L*. In the Appendix we describe a method to compute the largest inscribed rectangle, *S* in *L*.

Each of *n* units calculates a slightly different rectangle S_i depending on errors in Euclidean reconstruction of X_0. For a global agreement for *S*, we take the weighted average for each of the four vertices of S_i. The weight is the distance of the projector P_i from that vertex.

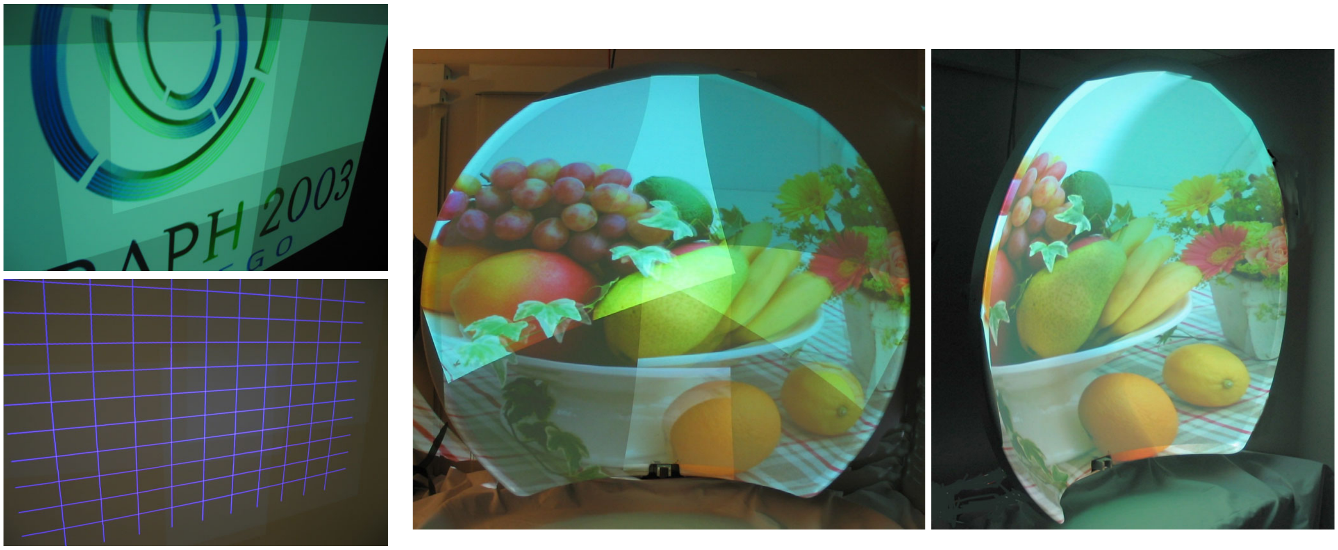

Figure 6: *Left : closeup of projector overlap for a four-unit cluster in a 2x2 arrangement. Left-top: projection of ordinary images (without normalization of illuminatioin the overlap area); Left-bottom: a projected grid for the same configuration provides a clear visual indication that the registration after global alignment is good. Middle: four self-contained units displaying overlapping images on a spherical surface with discernible mis-registration after linear estimate of quadric transfer; Right: seamless geometrically corrected display after non-linear refinement of quadric transfer parameters and intensity normalization.*

In the second step, we find the projection of the corners of S in the projector image space using H_{j0}^{-1}. Note that the re-projected corners of S will most likely extend beyond the physical image dimensions of the projector. Since S represents the displayable region, it indicates the extents of the input image, T, to be displayed. Thus, we can find the homography between input image and the projector image, H_{Tj}.

We texture map the input image onto a unit rectangle (of correct aspect ratio) and render with a projection matrix derived from H_{Tj} [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.

All the computations are performed symmetrically. After casual installation, it takes about 5 seconds per projector to 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 after casual installation is 30 seconds.

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

5.2 Curved Display using Ad-Hoc Clusters

The second type of display surface we consider for ad-hoc cluster projection is the quadric. Examples of quadric surfaces are domes, cylindrical screens, ellipsoids, or paraboloids. A solution for quadric surfaces may inspire future work in ad-hoc cluster projection on higher-order and non-planar display surfaces.

In computer vision literature, some relevant work has used quadrics for image transfer [Shashua and Toelg 1997]. In multi-projector systems however, although several approaches have been proposed for seamless multi-projector planar displays based on planar transfer (planar homography) relationships [Raskar 2000; Chen et al. 2000; Yang et al. 2001], there has been little or no work on techniques for parameterized warping and automatic registration of higher order surfaces. This is an omission because quadrics do appear in many shapes and forms in projector-based displays. Large format flight simulators have traditionally been cylindrical or dome shaped, planetariums and OmniMax theaters use hemispherical screens, and many virtual reality setups [Trimension Systems Ltd 2002] use a cylindrical shaped screen.

Alignment is currently done manually. Sometimes this manual process is aided by projecting a 'navigator' pattern [Trimension Systems Ltd 2002; Jarvis 1997]. We propose a completely automatic approach similar to the approach for planar clusters. Parametric approaches lead to reduced constraints on camera resolution, better tolerance to pixel localization errors, faster calibration and finally a simpler parameterized warping process.

Our main contribution here is the re-formulation of the quadric transfer problem and its application for seamless display. See Figure 6.

Simplification of Quadric Transfer Mapping between two arbitrary perspective views of an opaque quadric surface in 3D can be expressed using a quadric transfer function, Ψ. While a planar transfer can be computed from 4 or more pixel correspondences, quadric transfer requires 9 or more correspondences. If a homogeneous point in 3D, X (expressed as a 4×1 vector) lies on the quadric surface Q (expressed as a symmetric 4 by 4 matrix), then $X^T QX = 0$ and the homogeneous coordinates of the corresponding pixels x in the first view and x' in the second view are related by

$$x' \cong Bx - \left(q^T x \pm \sqrt{(q^T x)^2 - x^T Q_{33} x}\right)e$$

Given pixel correspondences (x, x'), this equation is traditionally used to compute the 21 unknowns: the unknown 3D quadric $Q = [\, Q_{33} \quad q;\ q^T \ 1\,]$, a 3x3 homography matrix B and the epipole in homogeneous coordinates, e. The epipole is the image of the center of projection of the first view in second view. This form used in Shashua and Toelg [1997] and even in later papers such as Wexler and Shashua [1999] contains 21 variables, 4 more than needed. Our method is based on a simple observation that we can remove part of this ambiguity by defining

$$A = B - eq^T \qquad\qquad E = qq^T - Q_{33}$$

and obtain the form we use,

$$x' \cong Ax \pm \left(\sqrt{x^T Ex}\right) e$$

Here $x^T Ex = 0$ defines the outline conic of the quadric in the first view and A is the homography via the polar plane between the second and the first view. Note that this equation contains (apart from the overall scale) only one ambiguous degree of freedom resulting from relative scaling of E and e. This can be removed by introducing an additional normalization constraint, such as $E(3,3) = 1$. Further, the sign in front of the square root is fixed within the outline conic in the image.

The suggested method to calculate parameters of quadric transfer, Ψ, i.e., $\{A, E, e\}$, directly from point correspondences involves estimating the quadric, Q, in 3D [Shashua and Toelg 1997], [Cross and Zisserman 1998] using a triangulation of corresponding pixels and a linear method. If the internal parameters of the two views are not known, all the calculations are done in projective space after computing the fundamental matrix. However, we noticed that when projectors rather than cameras are involved, the linear method produces very large reprojection errors, in the order of 20 or 30 pixels for XGA projectors. The computation of the fundamental matrix is inherently ill-conditioned given that the points on the quadric illuminated by a single projector do not have significant depth variation in most cases. We instead use known internal parameters and estimated Euclidean rigid transformations. Hence, unlike the planar case, computation of accurate image transfer in this case, involves three-dimensional quantities early in the calculation.

Our algorithm The steps are (a) Compute correspondences between cameras and projectors (b) triangulate and find the equation of the quadric in 3D (c) compute the quadric transfer parameters and (d) use quadric transfer to prewarp the input images.

As in the planar cluster case, each projector P_k, for $k = 1..n$ projects a structured pattern on the quadric, one at a time and is viewed by cameras C_i, for $i = 1..n$, of the n units in the group. Let us consider the problem of computing the quadric transfer, T_{CkPi}, mapping image in camera C_k to projector P_i. However, we cannot directly find quadric transfer, Ψ_{CkPi}, without first finding a rigid transformation , Γ^{Ck}_{Pi}, between P_i and C_k.

The steps for a camera C_k are as follows.

1. Since $\Gamma^{C_i}_{P_i}$ is known, we triangulate corresponding points in P_i and C_i to get 3D points and a mesh, D_i, of the display surface, and store them in C_i coordinate system.
2. Given 3D points in D_i (projected by P_i) and corresponding pixels observed by camera, C_k, of neighboring units, we find Γ^{Ck}_{Ci}. Then,
$$\Gamma^{Ck}_{Pi} = \Gamma^{Ck}_{Ci}\Gamma^{Ci}_{Pi}$$
3. Fit a 3D quadric Q_i to points in D_i transformed into C_k coordinate system
4. Find parameters of Ψ_{CkPi} from (Q_i, projection matrix M_{Ck} and M_{Pi} and pixel correspondences) using the new formulation of simplified quadric transfer
5. Perform nonlinear refinement of Ψ_{CkPi} , i.e., A_{CkPi},E_{CkPi} and e_{CkPi} to minimize pixel reprojection error i.e., distance between pixels in projector P_i and transferred corresponding pixels from camera C_k

Note that finding the pose of a camera from known 3D points on a quadric is error-prone because the 3D points are usually quite close to a plane. Since we know the camera internal parameters, we first find an initial guess for the external parameters based on a homography and then use an iterative algorithm described in Lu et al. [2000]. We use Powell's method for nonlinear refinement of the reprojection error.

Rendering For rendering, we treat the quadric transfer as a homography via the polar plane (i.e., A) plus a per-pixel shift defined by E and e. Similar to the cluster for planar display, without the aid of any environmental sensor or Euclidean markers in the scene, we exploit this homography along with tilt-sensor reading at U_k to align the display with the world horizontal and vertical. Given the relationship between the input image and the image in C_k, as well as quadric transfer between C_k and all P_i, $i = 1..n$, each unit warps the input image into its own image space via Ψ_{CkPi}. We defer the discussion on computing relationship between input image and image in C_k, plus intensity blending to an upcoming paper. Note that warping an image using a quadric transfer is different than rendering quadric surfaces [Watson and Hodges 1989].

We have implemented the rendering using a simple vertex shader program. For each projector unit, we map the input image as a texture onto a densely tessellated rectangular mesh, and compute the projection of each vertex of the mesh using the quadric transfer.

Results There are two demonstrations in the accompanying video. We show registration of three and four units on a concave spherical segment (Figure 6)and on a convex spherical segment. Our calibration process is relatively slow (about one minute per projector) compared to the planar case. The two time-consuming steps are computing the camera pose from near-planar 3D points using an iterative scheme in Lu et al. [2000], and a non-linear refinement of A, E, e to minimize the pixel reprojection error.

The techniques are ideal for creating single or multi-projector seamless displays without expensive infrastructure using the self-contained projector units. New possible applications are low-cost and flexible dome displays, shopping arcades, cylindrical columns or pillars. The approach and proposed ideas can also be treated as an intermediate step between planar to arbitrary free form shaped displays.

6 Implementation

We have built several prototypes. The oldest is a box which includes an XGA (1024x768) Mitsubishi X80 projector, a Dell Inspiron 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 only cable coming out of the box is for power. Our newest prototype uses an XGA Plus V1080 projector with a Sony Vaio.

The Euclidean calibration of a single display unit uses an auxiliary camera. First the camera of the display unit and the auxiliary camera undergo full stereo calibration using a checkerboard. The projector then projects onto a blank plane, for two or more orientations of the plane, and the projected points are reconstructed in 3D. Finally the projection matrix of the projector is determined using the correspondences between projector image points and the 3D data.

The display surfaces are mostly everyday surroundings like walls and corners, but we used an Elumens VisionStation dome for the quadric projection to ensure a truly quadric surface.

For the object augmentation in Section 4, the projected augmentation shows wobbles of a few mm for a projector which is 1-2m away, but this seems not particularly detrimental to the viewing experience (there is currently no smoothing on the camera motion).

For the planar transfer in Section 5, the reprojection error is about 2 pixels with pairwise homographies and about 0.3 pixels with global alignment. As seen in Figure 6(left), the images appear registered to sub-pixel accuracy.

For the quadric transfer in Section 5, reprojection error after linear estimation of the 3D quadric directly from point correspondences is about 20 pixels. After estimating the external parameters using method in Lu et al. [2000], the error is about 4 pixels. Finally, non-linear minimization of reprojection using Powell's method reduces the error to about 1.0 pixel. As seen in the video, when displayed on the half-dome, registration is about 1 pixel accurate.

7 Future Directions

Several new modifications and applications are possible with the proposed intelligent projector: (a) a steady-projection projector – a handheld projector that creates a stable image by responding to the geometric relationship between the projector and display surface and the instantaneous acceleration (measured by a tilt sensor); (b) intelligent flash – the projector can provide direction-dependent illumination for images captured by the camera in the unit, (c) shadow elimination – cooperation between overlapping projectors to fill in shadows during augmentation [Jaynes et al. 2001].

A geometrically aware projector can be further improved by considering photometric aspects. The projection can adapt for variations in surface reflectance or surface orientation and an intelligent augmentation may look for suitable surfaces to project on, avoiding low reflectance areas.

The system we have built is still fairly bulky, but the trend is for miniaturization of projectors. The techniques we have presented are ideal for a mobile form factor or temporary deployments of projectors. LEDs are replacing lamps and reflective instead of transmissive displays (DLPs, LCOS) are becoming popular. Both lead to improved efficiency requiring less power and less cooling. Several efforts are already on-going and show a great promise. For example Symbol Technologies [Symbol 2002] has demonstrated a small laser projector (two tiny steering mirrors for vertical and horizontal deflection) and has even built a handheld 3D scanner based on such a projector. Siemens has a built a 'mini-beamer' attachment for mobile-phones [Siemens 2002]. Cam3D has built a 'Wedge' display where a projector can be converted into a 'flat panel' display by projecting images at the bottom of a wedge shaped glass [Travis et al. 2002]. A future mobile projector may double up as 'flat panel' when there is no appropriate surface to illuminate, or ambient light is problematic. Super bright, sharp infinite focus laser projectors are also becoming widespread [Jenoptik 2002] which may allow shape-adaptive projection without focus and ambient lighting problems. In addition suitable input devices are also appearing e.g., Canesta [2002] has built a projected laser pattern on which you can type. The finger movement is detected by IR sensing. Finally novel lamp designs, especially those based on LEDs or lasers are creating smaller, lighter, efficient and long-life solutions.

8 Conclusion

Projectors are showing the potential to create new ways of interacting with information in everyday life. Desktop screens, laptops and TVs have a basic constraint on their size – they can never be smaller than the display area. Hand-helds such as PDAs are compact but the display size is too limited for many uses. In contrast, projectors of the near future will be compact, portable, and with the built-in awareness which will enable them to automatically create satisfactory displays on many of the surfaces in the everyday environment. Alongside the advantages, there are limitations, but we anticipate projectors being complementary to other modes of display for everyday personal use in the future, and to have new application areas to which they are especially suited.

This paper has investigated how to use projectors in a flexible way in everyday settings. The basic unit is a projector with sensors, computation, and networking capability. Singly or in a cluster, it can create a display that adapts to the surfaces or objects being projected on. As a hand-held, it allows projection of augmentation data onto a recognized object, plus mouse-style interaction with the projected data. It works with other units in an ad hoc network to create a seamless display on planar and curved surfaces. The ideas presented provide geometric underpinnings for a new generation of projectors – autonomous devices, easily adapting to operation within a cluster, and adaptive to their surroundings.

Appendix: Largest inscribed rectangle

Here we describe how to compute the largest world axis-aligned rectangle with given aspect ratio, a, inside a (possibly concave) polygon L. L is formed by the union of projected quadrilaterals. The solution for the three unknowns, two for position and one for scale of the rectangle, can be solved by stating a set of linear inequalities for convex polygons but not for concave polygons. We provide a simple re-parameterization of the problem.

Imagine L is drawn in $z = 1$ plane and a rectangle, R, of aspect ratio a is drawn in $z = 0$ plane. A center of projection $W{=}(x,y,z)$, z in [0,1], to map R into a rectangle S in z=1 plane, is considered valid if S remains completely inside L. We search for a center of projection (CoP) with minimum z (blue circle) because it will create the largest inscribed rectangle.

Consider the forbidden zones for W. Any CoP, that is inside the set of pyramids, created by a point on L with R as base, is invalid (yellow). Since the faces of the pyramid connecting z=0 and z=1 are all triangles, our algorithm computes the intersection of each triangle triple and keeps the one with smallest z value. We only need to consider two types of triangles, those connecting a vertex of L with edge of R and those connecting edge of L with a vertex or R. For n-sided polygon, L, we have 8n triangles with an O(n^4) algorithm. This is clearly suboptimal in comparison to the O(n^2) algorithms [Agarwal et al. 1996]. Yet this algorithm is very easy to implement in a few lines of code. Since n is O(number of projectors), runtime of the algorithm is still negligible.

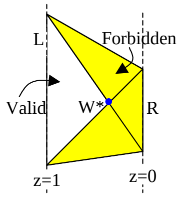

Figure 7: *Search for inscribed rectangle reparameterized as search for a COP*

Acknowledgements We would like to thank Mr Yoshihiro Ashizaki, Mr Masatoshi Kameyama, and Dr Keiichi Shiotani for helping us by providing industrial applications to motivate the work; Paul Dietz, Darren Leigh, and Bill Yerazunis who advised about and built the devices; Shane Booth for sketch in Figure 1; Joe Marks and Rebecca Xiong for many helpful comments; Debbi VanBaar for proof-reading, and Karen Dickie for timely and unrelenting administrative support.

References

Agarwal, P., Amenta, N., Aronov, B., and Sharir, M. 1996. Largest Placements and Motion Planning of a Convex Polygon. In *2nd International Workshop on Algorithmic Foundation of Robotics, 1996*, 28–37.

Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. 1993. Toolglass and Magic Lenses: The See-Through Interface. In *Proceedings of ACM SIGGRAPH 1993*, 73–80.

Bimber, O., Gatesy, S., Witmer, L., Raskar, R., and Encarnao, E. 2002. Merging Fossil Specimens with Computer-Generated Information. In *IEEE Computer*, 32–39.

Brown, M. S., and Seales, W. B. 2002. A Practical and Flexible Large Format Display System. In *The Tenth Pacific Conference on Computer Graphics and Applications*, 178–183.

Canesta, 2002. Miniature Laser Projector Keyboard. http://www.canesta.com, Cited December 2002.

Chen, Y., Chen, H., Clark, D. W., Liu, Z., Wallace, G., and Li, K. 2000. Automatic Alignment of High-Resolution Multi-Projector Displays Using An Un-Calibrated Camera. In *IEEE Visualization 2000*, 125–130.

Chen, H., Sukthankar, R., Wallace, G., and Li, K. 2002. Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. In *Proceedings of Visualization, 2002*, 135–142.

Cross, G., and Zisserman, A. 1998. Quadric Surface Reconstruction from Dual-Space Geometry. In *Proceedings of 6th International Conference on Computer Vision(Bombay, India)*, 25–31.

Crowley, J., Coutaz, J., and Berard, F. 2000. Things That See. *Communications of the ACM* (Mar.), 54–64.

Davis, J. 1998. Mosaics of Scenes with Moving Objects. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 354–360.

Hereld, M., Judson, I. R., and Stevens, R. L. 2000. Introduction to Building Projection-based Tiled Display Systems. *IEEE Computer Graphics and Applications 20*, 4, 22–28.

Humphreys, G., Eldridge, M., B., I., Stoll, G., Everett, M., and Hanrahan, P. 2001. WireGL: A Scalable Graphics System for Clusters. In *Proceedings of SIGGRAPH 2001*, 129–140.

Jarvis, K. 1997. Real Time 60Hz Distortion Correction on a Silicon Graphics IG. *Real Time Graphics 5*, 7 (Feb.), 6–7.

Jaynes, C., Webb, S., Steele, R., Brown, M., and Seales, B. 2001. Dynamic Shadow Removal from Front Projection Displays, . In *IEEE Visualization 2001*, 152–157.

Jenoptik, 2002. Laser Projector. http://www.jenoptik.com, Cited December 2002.

Kjeldsen, R., Pinhanez, C., Pingali, G., Hartman, J., Levas, T., and Podlaseck, M. 2002. Interacting with Steerable Projected Displays. In *Proc. of the 5th International Conference on Automatic Face and Gesture Recognition*, 12–17.

Levy, B., Petitjean, S., Ray, N., and Maillot, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *ACM Transactions on Graphics*, vol. 21, 3, 162–170.

Lu, C., Hager, G., and Mjolsness, E. 2000. Fast and Globally Convergent Pose Estimation from Video Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22*, 6, 610–622.

Majumder, A., He, Z., Towles, H., and Welch, G. 2000. Color Calibration of Projectors for Large Tiled Displays. In *IEEE Visualization 2000*, 102–108.

Pinhanez, C. 2001. The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces. In *Ubiquitous Computing 2001 (Ubicomp"01)*, 12–17.

Raskar, R., and Beardsley, P. 2001. A Self Correcting Projector. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 626–631.

Raskar, R., Welch, G., Cutts, M., Lake, A., Stesin, L., and Fuchs, H. 1998. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. In *Proceedings of ACM SIGGRAPH 1998*, 179–188.

Raskar, R., Brown, M., Ruigang, Y., Chen, W., Welch, G., Towles, H., Seales, B., and Fuchs, H. 1999. Multiprojector Displays using Camera-based Registration. In *IEEE Visualization*, 161–168.

Raskar, R., Welch, G., Low, K., and Bandyopadhyay, B. 2001. Shader Lamps: Animating Real Objects With Image-Based Illumination. In *Rendering Techniques 2001, The Eurographics Workshop on Rendering*, 89–102.

Raskar, R., VanBaar, J., and Chai, X. 2002. A Low Cost Projector Mosaic with Fast Registration. In *Fifth Asian Conference on Computer Vision*, 114–119.

Raskar, R. 2000. Immersive Planar Display using Roughly Aligned Projectors. In *IEEE VR 2000*, 27–34.

Rekimoto, J., and Nagao, K. 1995. The World Through the Computer: Computer Augmented Interaction with Real World Environments. In *Proceedings of UIST'95*, 29–36.

Rekimoto, J., and Saitoh, M. 1999. Augmented Surfaces: A Spatially Continuous Workspace for Hybrid Computing Environments. In *Proceedings of CHI'99*, 378–385.

Rekimoto, J. 1999. A Multiple-device Approach for Supporting Whiteboard-based Interactions. In *Proceedings of CHI'98*, 344–351.

Samanta, R., Zheng, J., Funkhouser, T., Li, K., and Singh, J. P. 1999. Load Balancing for Multi-Projector Rendering Systems. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, 12–19.

Shashua, A., and Toelg, S. 1997. The Quadric Reference Surface: Theory and Applications. In *IJCV*, vol. 23(2), 185–189.

Siemens, 2002. Siemens Mini Beamer. http://w4.siemens.de/en2/html/press//newsdesk_archive/2002/-foe02121b.html, Cited December 2002.

Sukthankar, R., Stockton, R., and Mullin, P. 2001. Smarter Presentations: Exploiting Homography in Camera-Projector Systems . In *International Conference on Computer Vision*, 82–87.

Symbol, 2002. Laser Projection Display. http://www.symbol.com/products/oem/lpd.html, Cited December 2002.

Travis, A., Payne, F., Zhong, J., and Moore, J. 2002. Flat panel display using projection within a wedge-shaped waveguide. http://ds.dial.pipex.com/cam3d/technology/technology01.html, Cited December 2002.

Trimension Systems Ltd, 2002. . http://www.trimension-inc.com/ , Cited Dec 2002.

Underkoffler, J., Ullmer, B., and Ishii, H. 1999. Emancipated Pixels: Real-world Graphics in the Luminous Room. In *Proceedings of ACM SIGGRAPH 1999*, 385–392.

Watson, B., and Hodges, L. 1989. A Fast Algorithm for Rendering Quadratic Curves on Raster Displays. In *Proc. 27th Annual SE ACM Conference*, 160–165.

Wellner, P. 1993. Interacting with paper on the DigitalDesk. *Communications of the ACM 36*, 7, 86–97.

Wexler, Y., and Shashua, A. 1999. Q-warping: Direct Computation of Quadratic Reference Surfaces. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), June, 1999*, 333–338.

Yang, R., Gotz, D., Hensley, J., Towles, H., and Brown, M. 2001. PixelFlex: A Reconfigurable Multi-Projector Display System. In *IEEE Visualization 01*, 68–75.

Yotsukura, T., Morishima, S., Nielsen, F., Binsted, K., and Pinhanez, C. 2002. Hyper Mask - Talking Head Projected onto Real Object. *The Visual Computer 18*, 2, 111–120.