

Auto-Calibration of Multi-Projector Display Walls

Andrew Raij and Marc Pollefeys

{raij, marc}@cs.unc.edu

Department of Computer Science, University of North Carolina at Chapel Hill

Abstract

By treating projectors as pin-hole cameras, we show it is possible to calibrate the projectors of a casually-aligned, multi-projector display wall using the principles of planar auto-calibration. We also use a pose estimation technique for planar scenes to reconstruct the relative pose of a calibration camera, the projectors and the plane they project on. Together with assumptions about the pose of the camera, we use the reconstruction to automatically compute the projector-display homographies needed to render properly scaled and oriented imagery on the display wall. The main contribution of this paper is thus to provide a fully automated approach to calibrate a multi-projector display wall without the need for fiducials or interaction.

1. Introduction

There are two major geometric issues that must be addressed in the design of a multi-projector display wall. The first is how to align projectors so that imagery is contiguous across projector boundaries. Manual alignment is a tedious, time-consuming process but research in this area has led to several automatic alignment methods [3, 7, 8, 10, 1, 6] that use computer vision and graphics to correct for misalignment in software. Once projectors are aligned, what remains is to define the dimensions and orientation of the displayable area and relate this to the configuration of the projectors on the display surface.

Several display systems use fiducials to define the displayable area on a planar surface. Rehg et al [8] use four manually-placed fiducials to set a quadrilateral within a key-stoned projection as the displayable area. The corners of the projected imagery are mapped by homography to the four fiducials. Since the displayable area is an arbitrary quadrilateral, the projected imagery can be severely distorted. A real-time system tracks the fiducials, allowing the user to redefine the displayable area so that the image has whatever properties the user desires. The PixelFlex [10] and PixelFlex2 [3] systems also use four fiducials, although nei-

ther allows the user to redefine the displayable area on the fly. Both systems use fiducials that are carefully placed to form a rectangle aligned with the natural human sense of vertical and horizontal. A coordinate system is imposed on the rectangle whose relative dimensions match that of the real world, allowing imagery of the proper aspect ratio to be displayed. The fiducials are related by homography to the orientation of the projectors on the plane.

In this paper, we propose an automatic method for defining the display area on a plane, removing the need for physical fiducials and measurement of the area defined by them. By treating projectors as pin-hole cameras, we show that planar auto-calibration, proposed by Triggs in [9], can be used to determine the intrinsics of an array of projectors projecting on a single plane. We then reconstruct the camera, projectors and display plane using a relative pose estimation technique for planar scenes. This allows us to define the mapping from projector to display that is needed for proper rendering. Note that all three of these stages are especially challenging because the only scene observed by the camera and projectors is a plane.

Okatani and Deguchi [2] also estimate the relative pose of multiple projectors with respect to a planar display, but they require calibrated projectors. Raskar and Beardsley [5] treat a camera and projector mounted together as a stereo pair, and then estimate projector intrinsics and relative pose by observing the plane from two different poses. Tilt sensors in the camera-projector unit give the alignment of the projector image plane with respect to the world, allowing the projection of properly oriented imagery of a particular aspect ratio on vertical planes. In [7], Raskar et al extend this work to multiple projector-camera units for building ad-hoc multi-projector displays. Our work is different because it does not require tilt sensors, nor explicitly mounted projector-camera stereo pairs. We instead use one camera that can view all projections on the plane.

The testbed for our work consists of $n = 8\ 1024 \times 768$ LCD projectors projecting on a single plane. A mirror mounted on a pan-tilt unit is positioned in front of each projector, allowing the shape of the display to change. A calibrated camera (Sony SX900 black & white 1394 camera

with 1280×960 resolution), used for automatic projector alignment, is positioned such that it can view all projections on the plane. By projecting structured light, feature correspondences are made between the projectors and camera, and the homographies H_{cp} , $p = 1 \dots n$ from camera to projector are computed from the correspondences. However, for proper rendering (as in [4]) in a selected display area, one needs to compute H_{dp} , $p = 1 \dots n$, homographies that define the relationship between each projector and the display area in the plane. In this paper, we automatically define a display area, allowing computation of H_{dp} without the need for physical fiducials and manual user setup.

2. Planar Auto-Calibration

The planar autocalibration constraints first shown by Triggs in [9] can be used to calibrate the projectors of a single-plane multi-projector display wall. Given n projectors projecting on a plane, and a camera observing the plane, the calibrated image of the plane's direction basis must remain orthonormal in all projectors and the camera. Alternatively, the image of the circular points must be on the image of the absolute conic (IAC) in all projectors and the camera. If C_p is the inverse of a projector p 's intrinsics matrix K_p and C_c is the inverse of the camera's intrinsics matrix K_c , then $\omega_p^{-1} = C_p^T C_p$ is the IAC in projector p and $\omega_c^{-1} = C_c^T C_c$ is the IAC in the camera. Furthermore, if $X_c = \frac{1}{\sqrt{2}}(x_c + iy_c)$, $X'_c = \frac{1}{\sqrt{2}}(x_c - iy_c)$ and $X_p = H_{cp} X_c$, $X'_p = H_{cp} X'_c$ are the circular points in the camera and projectors, respectively, then we can express the constraints mathematically as

$$\begin{aligned} X_c^T \omega_c^{-1} X_c &= 0 \\ (H_{cp} X_c)^T \omega_p^{-1} (H_{cp} X_c) &= 0 \text{ for } p = 1 \dots n \end{aligned} \quad (1)$$

This yields $2n + 2$ constraints on n projectors and the camera image of the circular points. This is not enough to estimate the full five-parameter model of the intrinsics of the camera and each projector, so we make some reasonable assumptions to decrease the number of unknowns. We assume the camera has been calibrated so the intrinsics K_c are known. We also assume that projector pixels are square. Most commodity projectors have a principal point that is offset vertically so that the projection is not occluded by the ceiling or table the projector is mounted to. We assume that this vertical offset is unknown but that projectors of the same brand and zoom setting have the same offset. The horizontal component of the principal point is considered to be at the center of the image. In sum, we estimate n projector focal lengths, a single principal point offset, and four values for the camera image of the circular points, for a total of $n + 5$ unknowns. Note that if we estimate the n principal

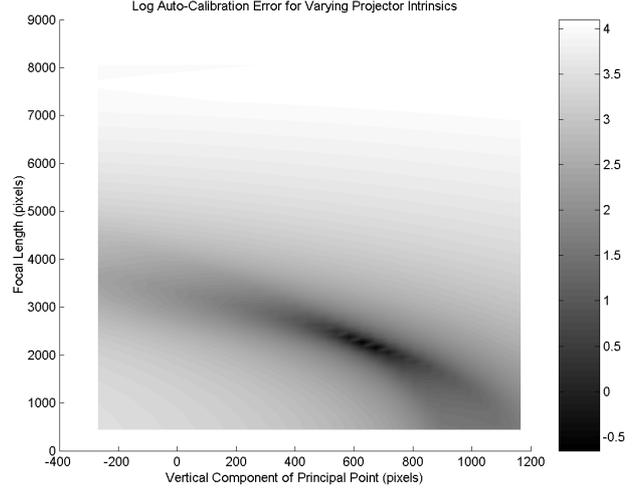


Figure 1. Our initialization algorithm finds clear minima in a search over the intrinsics.

point offsets separately, the number of unknowns becomes $2n + 4$, and the solution to the problem is underdetermined. We start the estimation with a search over the projector focal lengths and principal point offset and then refine the results using a non-linear minimization.

2.1. Initialization

The problem with iterative minimization is the need for reasonable initial values that will converge to a solution. We describe here an initialization algorithm that our experiments have shown works in practice. The algorithm uses a pose estimation technique for planar scenes proposed by Triggs in [9]. Given a calibrated homography between camera and projector $\bar{H}_{cp} = K_p^{-1} H_{cp} K_c$, the technique produces the relative pose of the camera, projector and two potential planes, only one of which is the plane of interest. For a single projector p , our initialization algorithm searches over a reasonable range of values for the projector's intrinsics K_p . From the current hypothesis of K_p and the known H_{cp} , we use Triggs' method to compute two potential planes that are compatible with this hypothesis and the known camera-projector homographies. We then determine the camera image of the circular points in each plane. Given these hypotheses of the circular points, we assume all projectors have the hypothetical intrinsics K_p and test the auto-calibration constraints for both sets of potential circular points. The K_p that best satisfies the auto-calibration constraints is the initial K_p for the projector p . Fig. 1 shows a plot of the error in the auto-calibration constraints for varying intrinsics. Note the clear minima in the error plot.

p	CS-PR	CC-PR	CC-PV	CS-PV
1	2267.81	2173.30	2149.39	2095.71
2	2200.50	2206.80	2171.94	2154.88
3	2148.76	2223.58	2122.13	2133.44
4	2229.53	2301.02	2165.01	2156.52
5	2219.13	2181.43	2172.84	2138.11
6	2199.33	2194.60	2148.33	2128.62
7	2253.56	2258.69	2191.39	2230.51
8	2250.98	2267.79	2177.96	2152.01
c_y	645.58	668.42	682.83	704.58

Table 1. Estimated focal lengths of projectors p=1...8 and vertical principal point offset, each column representing a different configuration. (CC = Camera Center, CS = Camera Side; PR = Projectors Rectangular, PV = Projectors V-Shaped)

2.2. Non-Linear Refinement

The refinement is done with a Levenberg-Marquardt non-linear least squares minimizer. The following equation, derived from the planar auto-calibration constraints discussed previously, is minimized

$$\|X_c^T \omega_c^{-1} X_c\|^2 + \sum_{p=1}^n \|X_p^T \omega_p^{-1} X_p\|^2.$$

Ideally, we would also do a maximum likelihood estimation over all the parameters but we have not implemented this yet. Although we have not yet performed a precise evaluation of our estimation, we have repeated the experiment using several different geometric configurations of the scene and found that our intrinsics estimation consistently produces similar values. Table 1 summarizes these results.

3. Reconstruction

We can also use the pose estimation method from the previous section to determine the extrinsics of the cameras and projectors and reconstruct the display plane. Given a calibrated homography between the camera and a projector, the method will give the relative pose of the projector with respect to a canonical camera at position $[0, 0, 0]$ with orientation $I_{3 \times 3}$. Furthermore, the coordinate system is scaled such that the baseline of the camera and projector is normalized to a length of 1. The method will also produce two planes, only one of which is the real-world plane we are trying to reconstruct. The plane that yields the smallest error during the initialization algorithm is the plane we select. It is important to note that the distance of the plane from the

camera is in terms of the camera-projector baseline. Therefore, we first reconstruct the plane, camera and projector with each camera-projector pair. The n separate reconstructions obtained are in the same frame up to a scale factor. By normalizing each reconstruction by the distance from the plane to the camera, the reconstructions are merged into the same frame. Once the projectors and plane are determined, we can reconstruct the image of the projectors on the plane. For each projector, we cast a ray from the center of projection through the corners of the image plane. The intersections of the rays with the display plane are the projector images.

The left side of Figure 2 shows the reconstruction of one geometric configuration and the right side shows the camera image of the plane with the projectors on for this configuration. Note the close visual match between the reconstructed projections and the projections in the camera image. Although we have not yet performed a precise evaluation of the calibration accuracy, the reconstructed configuration corresponds visually to the system we have used.

4. Viewport Selection

The rectangular viewport should be aligned with the world horizontal and vertical. A knob could be provided for rotating the plane to the correct orientation (effectively one parameter), but we propose a more automatic approach. For cameras and projectors, the image y-axis is in general not vertical in the world because of tilting, but the image x-axis is in general horizontal. Thus we obtain the plane horizontal by orthogonal projection of the camera x-axis on the display plane. The plane vertical is then given by the cross product of the plane horizontal and the normal to the plane estimated during reconstruction. On the right side of Figure 2 is a camera image of the display with dashed lines extending towards the vanishing points. The vanishing points are computed by projecting the estimated direction vectors into the camera. Note that the lines parallel our impression of horizontal and vertical from other cues in the image.

Once the correct orientation is known, the largest fully inscribed rectangle, possibly of a specific aspect ratio, is selected as the display area. We use the method proposed by Raskar et al [7] to estimate the largest inscribed rectangle of the rendering application's aspect ratio. An example of a selected 16:9 viewport is shown in Figure 2.

5. Conclusions and Future Work

In this work, we have shown how to estimate the intrinsics of an array of projectors projecting on a planar display surface using only a set of camera-projector homographies. We have also shown that it is possible to estimate the extrinsics of the projectors and reconstruct the display plane.

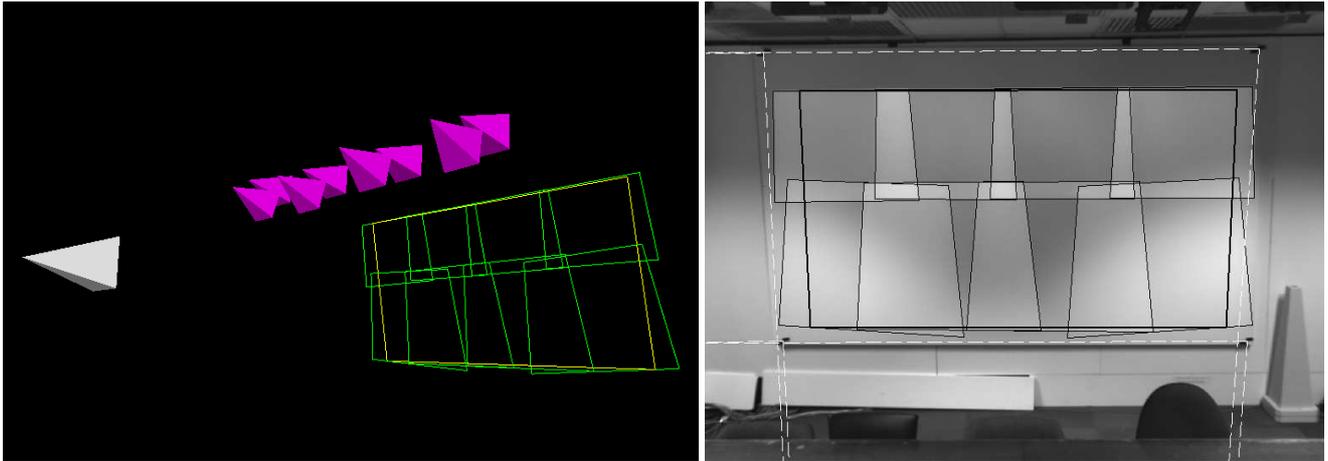


Figure 2. The reconstructed Camera Center - Projectors Rectangular configuration (left) and corresponding camera image. Computed outlines of the projectors and a 16:9 viewport of maximum inscribed area is shown. Dashed lines in the image extend towards computed horizontal and vertical vanishing points of the plane.

Using the reconstruction and some assumptions on camera placement, we also discuss how to select a viewport in the plane that results in the proper orientation and aspect ratio for rendered imagery. Therefore, our approach allows a fully automatic auto-calibration of a multi-projector/camera system without the need for fiducials in the scene. It is also important to note that for this purpose a precise metric calibration is not necessary, since it is sufficient if the viewport on the display roughly corresponds to a horizontally aligned rectangle in the real world.

Our future plans are to improve the calibration accuracy by developing a maximum likelihood estimation based on bundle-adjustment. We also intend to calibrate for camera intrinsics and model the effect of radial distortion in both the camera and projectors.

6. Acknowledgements

This work was partially supported by NSF Career award IIS 0237533, Department of Energy ASC VIEWS Program B519834 and DARPA DARWARS program ONR N00014-03-1-0589. We want to thank the past and present members of the PixelFlex team, whose hard work made it possible to use PixelFlex as a testbed for this research.

References

[1] Y. Chen, D. W. Clark, A. Finkelstein, T. Housel, and K. Li. Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera. In *Proc. IEEE Visualization*, 2000.

[2] T. Okatani and K. Deguchi. Autocalibration of projector-screen-camera system: Theory and algorithm for screen-to-camera homography estimation. In *Proc. IEEE Ninth International Conference on Computer Vision*, pages 774–781, 2003.

[3] A. Raij, G. Gill, A. Majumder, H. Towles, and H. Fuchs. Pixelflex2: A comprehensive, automatic, casually-aligned multi-projector display. In *Proc. IEEE International Workshop on Projector-Camera Systems*, 2003.

[4] R. Raskar. Immersive planar display using roughly aligned projector. In *Proc. IEEE Virtual Reality*, 2000.

[5] R. Raskar and P. Beardsley. A self correcting projector. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 626–631, 2001.

[6] R. Raskar, M. S. Brown, R. Yang, W.-C. Chen, H. Towles, B. Seales, and H. Fuchs. Multi-projector displays using camera based registration. In *Proc. IEEE Visualization*, 1999.

[7] R. Raskar, J. van Baar, P. Beardsley, T. Willwacher, S. Rao, and C. Forlines. ilamps: Geometrically aware and self-configuring projectors. In *Proc. ACM SIGGRAPH*, 2003.

[8] J. M. Rehg, M. Flagg, T.-J. Cham, R. Sukthankar, and G. Sukthankar. Projected light displays using visual feedback. In *Proc. Seventh International Conference on Control, Automation, Robotics and Vision*, 2002.

[9] B. Triggs. Autocalibration from planar scenes. In *Proc. Fifth European Conference on Computer Vision*, pages 89–105, 1998.

[10] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. S. Brown. Pixelflex: A reconfigurable multi-projector display system. In *Proc. IEEE Visualization*, 2001.