Geometric Calibration for a Multi-Camera-Projector System

Ricardo R. Garcia   Avideh Zakhor

UC Berkeley

Berkeley, CA

{rrgarcia,avz}@eecs.berkeley.edu

Abstract

In this paper, we describe a calibration method for multi-camera-projector systems in which sensors face each other as well as share a common viewpoint. We use a translucent planar sheet framed in PVC piping as a calibration target which is placed at multiple positions and orientations within a scene. In each position, the target is captured by the cameras while it is being illuminated by a set of projected patterns from various projectors. The translucent sheet allows the projected patterns to be visible from both sides, allowing correspondences between devices that face each other. The set of correspondences generated between the devices using this target are input into a bundle adjustment framework to estimate calibration parameters. We demonstrate the effectiveness of this approach on a multi-view structured light system made of three projectors and nine cameras.

1. Introduction

Camera and projector calibration is a crucial step in the deployment of many computer vision systems, especially in those with multiple cameras and projectors. Common applications of multi-camera-projector (MCP) systems include CAVE-like or large displays [14, 16] and multi-view geometry capture with structured light [9, 6]. There are publicly available tools such as Bouguet's Camera Calibration Toolbox [2] for calibrating cameras. Multi-camera calibration has also been addressed in the literature with publicly available tools [17, 15, 8]. Specifically, Svoboda et al. propose a multi-camera calibration method which tracks the position of an LED across synchronized images from each camera view [15]. Projector calibration has also been studied extensively. Since projectors cannot observe the scene, projector calibration has traditionally been preceded by camera calibration[20, 12, 1]. Alternatively, it is possible to calibrate a camera and projector simultaneously [3, 18]. There are also a number of approaches to calibrate multi-projector systems such as CAVE-like augmented reality systems [4],

and flexible reconfigurable projection screens [19]. In the majority of these approaches, the intrinsics and extrinsic parameters are not directly computed. Rather, a mapping between the projector pixels and projection surface is estimated in order to reduce interference between projectors.

A number of methods have been proposed for full intrinsic and extrinsic calibration of an MCP system [14, 10, 11]. In [14], cameras and projectors are arranged to create a large display surface. The geometry of the surface is reconstructed by using structured light patterns to determine correspondences between a pair of stereo cameras. These points are triangulated to 3D positions that can then be used to calibrate the projector. This process is repeated for each projector separately. In [10], all the cameras are first calibrated; then each projector is calibrated by projecting fiducials onto a board, capturing images of the projection, and robustly detecting the 2D and 3D positions of the fiducials to calibrate the projector. They develop an automated process to make calibration easy to complete. With this approach, it is not possible to generate projector-projector correspondences, except in configurations with highly overlapped views between the projectors; it is also not possible to generate correspondences between all devices in the system. Kobayashi et al. perform MCP calibration as well but require the cameras to be positioned such that they are directly illuminated by the projectors [11]. This greatly restricts the possible camera-projector system configurations. The downside to each of these MCP calibration methods is that the projector intrinsic and extrinsic parameters are estimated separately from the camera parameters, thus propagating the error in the camera calibration to the projector. Ideally, all parameters should be jointly recovered.

In this work, we address the problem of robust calibration of MCP systems. Our primary motivation is a multi-view structured light system, as shown in Fig. 1, which captures the dynamic geometry of a scene from multiple perspectives. In particular, we are interested in camera-projector configurations where the sensors in the system surround a central observation volume. As seen in Fig. 1, the cameras and projectors are oriented so that their cen-
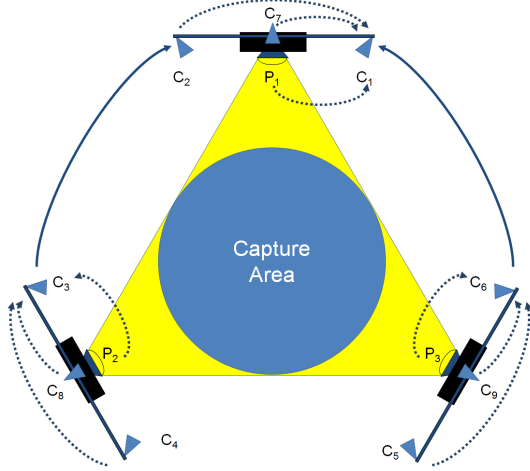
Figure 1. Layout of a multi-view structured light system. Arrows represent the pairs of devices whose extrinsics are computed during initial calibration estimation.
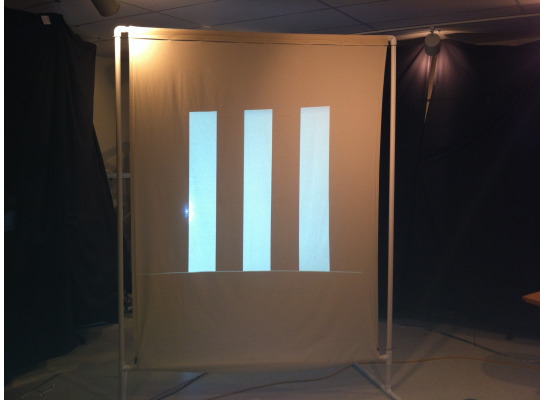

Figure 2. Calibration sheet used to increase overlapping view between cameras and projectors. The projected pattern on the sheet appears sharp and in focus from both sides of the sheet.

tral axes of observation all roughly point towards the center of a capture volume. Thus, some of the sensors are facing each other rather than pointing in the same general direction. The key to joint calibration of a system with multiple cameras and projectors is to generate a set of correspondences across all devices. To accomplish this, we use a simple calibration target, namely a sheet stretched over a PVC pipe frame, as shown in Fig. 2. Binary patterns encoding the vertical and horizontal pixel locations of the projector are projected onto this target to establish camera-projector pixel correspondences across all cameras. The patterns that illuminate the sheet are clearly visible from both sides of the sheet making it possible to establish correspondences between cameras and projectors even when they are positioned in opposite facing directions. The sheet does not significantly scatter the projected light, resulting in the projected pattern appearing in focus from both sides

of the calibration sheet. The key to our approach is to decode the projector to camera correspondences in order to generate camera to camera and projector to projector pixel correspondences. The final set of global correspondences consists of the projected coordinates of observed 3D points in each of the devices in the system. This set along with its corresponding set of 3D positions is input to a bundle adjustment framework along with a coarse initial estimate of each device's intrinsic and extrinsic measurements. The output is the refined intrinsic and extrinsic parameters for all the devices in an MCP system. Of existing work in the literature, our proposed approach comes closest to that of [7] except that the translucent sheet enables us to find correspondences among all devices, rather than only a subset of devices with overlapping views.

The outline of the remainder of this paper is as follows: Section 2 explains the pattern projection and decoding process. Section 3 presents strategy for generating correspondences across all sensors in the system. Section 4 includes results on calibrating an actual system, and Section 5 provides concluding remarks.

2. Pattern Projection and Capture

To establish global correspondences between all devices, first correspondences between each projector and all cameras are generated. This is done by encoding the horizontal and vertical position of each projector pixel with binary structured light patterns. The overall calibration process requires multiple correspondence sets to be generated by placing the projection target in multiple positions. In doing so, correspondences are found at 3D locations that are spread throughout the capture volume. At each target position, each projector projects a set of vertical and horizontal binary coded patterns which are captured by all the cameras. These images are used to establish global correspondences across all devices. During calibration, the target is positioned in the center of the capture area within the limited depth of field of the projectors. This ensures the projected patterns appear in focus and allows for accurate decoding.

The outline of this section is as follows: We begin by describing the binary patterns in Section 2.1. The process for decoding the binary patterns is explained in 2.2. Finally, the correspondence mappings between the cameras and projectors are presented in Section 2.3.

2.1. Pattern Projection

For a single camera and projector, pixel correspondences can be determined by projecting binary coded patterns [14]. Most commonly, the patterns are generated by taking the index value of each column of pixels in the projector and representing it as as a binary number. A binary image is projected for each bit of the binary representation. The result is a set of striped patterns of decreasing width as shown

in Fig. 3. Cameras observing a scene illuminated by the binary patterns can determine the projector column illuminating the scene point observed by each pixel by decoding the set of binary patterns. This creates a correspondence between each pixel in the camera and the columns of the projector. If a set of vertical and horizontal binary coded patterns are projected, then pixel to pixel correspondences between the camera and projector are established. This process is repeated for all cameras observing the binary projected patterns. [1]

We opt to use a translucent, planar sheet as the projection target to calibrate all sensors. Translucency is needed so that even though the target is being projected from one side, the cameras viewing it from the opposite side can still view and decode the projected patterns. The projected patterns appear equally focused on both sides of the sheet, although with slightly lower brightness when viewed on the side not directly illuminated. Furthermore, to ensure proper decoding, the sheet must approximate a planar surface. As such, we have opted to frame it with PVC pipes as shown in Fig. 2. Using a translucent sheet as a projection target greatly increases the positions from which cameras can observe the projected patterns. In spite of this, cameras whose central axis are nearly perpendicular to the normal of the calibration sheet are unable to decode the projected patterns. To resolve this, we capture image sets with the target in different positions so as to obtain correspondences between most devices. Specifically, for the MCP system in Fig. 1, we place the sheet at the center of the capture area and orient it in multiple directions near the center of the capture area for each pattern projection and camera capture. In doing so, we are able to determine correspondences between a projector and all of the surrounding cameras. We use two sets of patterns to determine full pixel to pixel correspondences: a vertical set of binary codes, $B_n^V, n \in \{0, ..., K_V - 1\}$, and a horizontal one, $B_n^H, n \in \{0, ..., K_H - 1\}$, where K_V and K_H are related to the resolution of each projector, i.e. $K_V = \lceil \log_2(\text{projector width}) \rceil$ and $K_H = \lceil \log_2(\text{projector height}) \rceil$. In addition to the frames that represent the binary value, the inverse of each frame is also projected. As explained later, capturing a binary frame and its inverse makes the decoding process more robust and obviates the need to choose a global threshold to decode each frame.

2.2. Decoding

Once both sets of binary patterns are projected and captured by all the cameras, the images are decoded for a given target position. To decode each bit in the projected binary
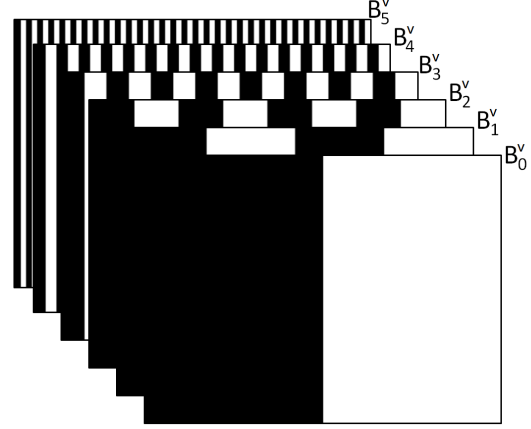

Figure 3. Vertical binary coded projection patterns.

code, the image for each bit and its inverse representation are both used. We first compute the absolute difference between the two frames; the larger this difference the more reliable a binary labeling of 0 or 1. A small absolute difference occurs when camera pixels are not directly illuminated by the projected patterns or when the transition between black and white stripes in the projected binary images falls on a camera pixel. Only bits where the absolute difference is greater than an assigned minimum threshold are considered to be properly decoded. For such pixels, the bit is decoded as a 1 if the intensity from a frame is greater than its inverse, and as 0 when the inverse frame has a greater intensity. [2]

After decoding all bits, we only keep the correspondences for pixels where all K_V and K_H bits of the vertical and horizontal patterns respectively are properly decoded. For cases where the projector resolution is higher than that of the camera, camera pixels can still be kept as correspondences even if the least significant bits are not decoded. [3]

2.3. Mapping

Assume a system consists of the set of cameras C, projectors P, and target plane positions T. For each camera $c \in C$ at a given plane position $t \in T$, correspondence matrices $Q_{x:c,p}^t$ and $Q_{y:c,p}^t$ are defined between the pixels of the camera and the x and y coordinates of each projector $p \in P$. The dimensions of $Q_{x:c,p}^t$ and $Q_{y:c,p}^t$ are both equal to the resolution of the camera c. Each element in the correspondence matrix holds the coordinates, either x or y, of its corresponding pixel in projector p. These correspondences

---

[1] An alternative to direct binary coding for the binary patterns is Gray codes. We have empirically found direct encoding of the column indices to be sufficient since we only keep the decoded positions for pixels that are confidently decoded.

[2] In our camera-projector configuration, the light sources of the projectors are not visible in the captured camera images. For the configurations where this does not hold, the lamp needs to be masked out and camera gain needs to be adjusted accordingly.

[3] The projector-camera resolution mismatch for generating correspondences makes the decoding of the least significant bits difficult [16]. It is possible to use additional sparse high frequency patterns to obtain more accurate correspondences.

are generated from the decoded binary patterns as described in Section 2.2. Jointly, $Q_{x:c,p}^t$ and $Q_{y:c,p}^t$ are referred to as $Q_{c,p}^t$. In the remainder of this paper, we will drop the x and y notation for each correspondence matrix and merely refer to the union set. The correspondence matrices from the projector to the cameras are defined as $\bar{Q}_{p,c}^t$ [4] with the dimensions equal to the resolution of projector p. Each entry in these matrices holds the coordinates to the corresponding pixel in the camera c. We use $Q_{c,p}^t$ to populate correspondence matrix $\bar{Q}_{p,c}^t$. Specifically for each camera pixel in $Q_{c,p}^t$, its corresponding projector pixel in $\bar{Q}_{p,c}^t$ is populated with the location of the original camera pixel. Fig. 4 shows an example of $Q_{c,p}^t$ graphed as an image. The color in this picture represents the value of the corresponding coordinate in the projector. The figure illustrates how densely correspondences can be determined between devices. Also, the calibration target is large enough to allow correspondences to be established throughout most of the camera's field of view. The target is placed at multiple positions to ensure not only correspondences across all devices, but also correspondences for all projector pixels. Mask matrices $M_{c,p}^t$, $\bar{M}_{p,c}^t$ are generated to indicate which pixels in $Q_{c,p}^t$ and $\bar{Q}_{p,c}^t$ have a correspondence. The sets of correspondence matrices and masks are generated for all $|T| \times |C| \times |P| \times 2$ combinations of $t \in T$, $c \in C$, and $p \in P$.

3. Full Correspondence Generation

So far, the sets of camera-projector correspondences for each target orientation have been determined. The goal is to use the individual pairwise correspondence sets to find global correspondences across all, or nearly all, devices. For each projector, the projector pixels with correspondences in multiple cameras are used to identify a set of potential global correspondences. For each of the projector pixels in this set, the corresponding pixels in each of the cameras are identified. These camera pixels together with the camera to projector correspondence matrices are used to find potential correspondences to the other projectors in the system. A detailed explanation of the above process follows.

To start, we use the set of pairwise correspondences and masks from a single target position t, $Q_{c,p}^t, M_{c,p}^t, \bar{Q}_{p,c}^t, \bar{M}_{p,c}^t, \forall c \in C, p \in P$ to generate a set of global correspondences $G^t = \{G_x^t, G_y^t\}$. Global correspondences are those seen across many devices in the system. The matrix G_x^t has $|P| + |C|$ rows; its number of columns is equal to the number of global correspondences found for target position t. The global correspondence sets for each target position are concatenated to create a final set of global correspondences $G = \bigcup_{t \in T} G^t$. The matrix G
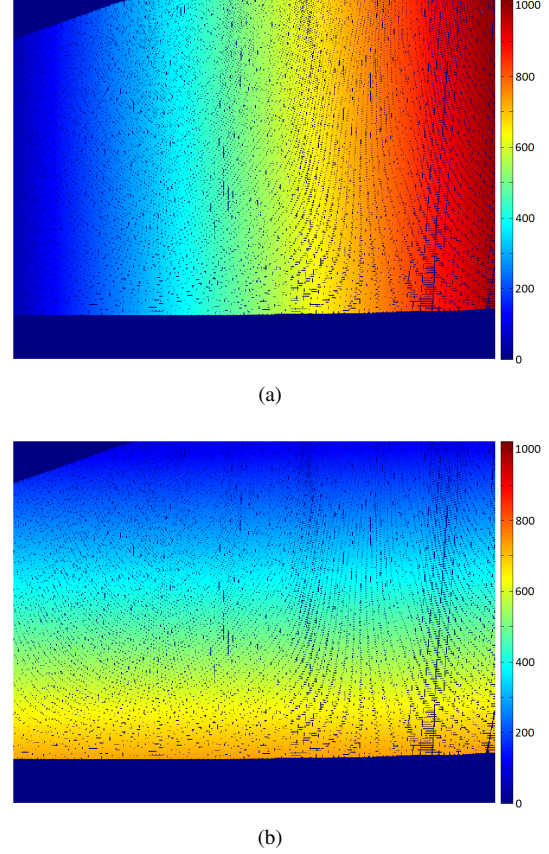

(a)

(b)

Figure 4. Mapping of correspondences from camera c to (a) x-coordinates $Q_{x:c,p}^t$ and (b) y-coordinates $Q_{y:c,p}^t$ of projector p.

has $|P| + |C|$ rows and the number of columns is equal to the sum total of correspondence points found in all target positions.

Fig. 5 shows the block diagram of the steps to generate G^t for the first projector p'. The processing for the remaining projectors in the system is more or less identical to the first one except for a minor detail to be described shortly.

3.1. Correspondences using Visibility Matrices

We start by examining the projector to camera correspondences for the projector p', denoted by $\bar{Q}_{p',c}^t$, $c \in C$. The set of correspondence masks, $\bar{M}_{p',c}^t, c \in C$, are element-wise summed together, $\sum_{c \in C} \bar{M}_{p',c}^t$ to obtain the projector visibility matrix, denoted by $\bar{V}_{p'}^t$, for projector p'. The dimensions of $\bar{V}_{p'}^t$ are equal to the resolution of projector p'. The values in the projector visibility matrix represents the number of cameras that are able to view each pixel of the examined projector p'. Only projector pixels that are visible to enough cameras, i.e. larger than $C_{\min}$, are considered as candidate global correspondences. These pixels are used to generate the visible projector pixel mask for projector p' and target t, denoted by $\bar{V}_{M:p'}^t$. Intuitively,
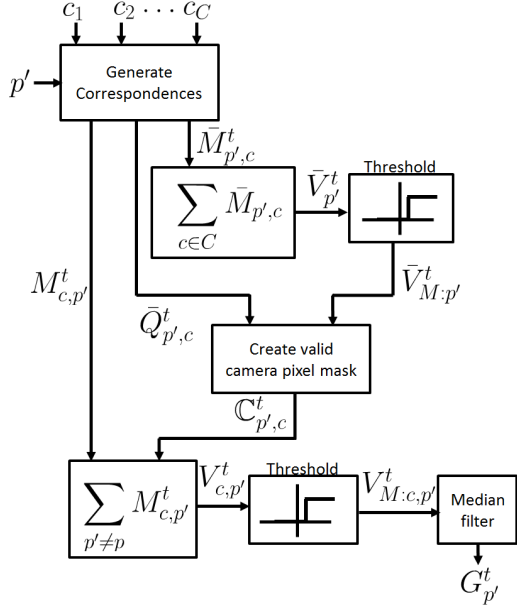
Figure 5. Block diagram of the processing steps for the first projector p'.

the projector pixels in $\bar{V}_{M:p'}^t$ have correspondences in a sufficiently large number of cameras.

The visible projector pixel mask $\bar{V}_{M:p'}^t$ along with the projector to camera correspondences $\bar{Q}_{p',c}^t$ are used to identify the pixels in each camera, $\mathbb{C}_{p',c}^t$, that correspond to the pixels in $\bar{V}_{M:p'}^t$. Intuitively, the camera pixels in $\mathbb{C}_{p',c}^t$ identify the correspondences to those projector p' pixels that are "visible" in a sufficiently large number of cameras.Next, we determine the number of projectors that can see each camera pixel in $\mathbb{C}_{p',c}^t$ denoted by matrix $V_{c,p'}^t$. The matrix $V_{c,p'}^t$ is defined by setting the $(i,j)^{th}$ value to $\sum_{p'' \neq p'} M_{c,p''}^t(i,j)$ for positions where $\mathbb{C}_{p',c}^t(i,j)$ is non-zero. Next, we determine the subset of camera pixels in $V_{c,p'}^t$ that are visible in a sufficiently large number of projectors, i.e. in more than $P_{\min}$ projectors. We denote this set by matrix $V_{M:c,p'}^t$ and compute it for every camera. Intuitively the pixels in $V_{M:c,p'}^t$ represent pixels in camera c that are sufficiently visible and can therefore be used to calibrate the projector extrinsically with respect to the other devices.

In most cases, if a correspondence between projectors exists, multiple cameras are able to observe it. Sometimes small errors in the camera to projector pairwise correspondences can occur due to the mismatch in resolution of cameras and projectors. The camera redundancy in these projector to projector correspondences can be used to remove decoding errors. To do so, we stack the set of redundant observations from multiple cameras to find correspondences between projector p' and p''. For each set of redundant observations corresponding to a single projector to projector

correspondence, the median x and y coordinates are used as the actual correspondence location and stored in matrix $G_{p'}^t$ along with correspondence coordinates for the cameras.

3.2. Global Correspondence Matrix

We now proceed to find correspondences $G_{p''}^t$ for the second projector p''. A matrix $V_{M:p''}^t$ for a second projector p'' is generated in the same way as $V_{M:p'}^t$ for the first projector p' and its coordinates are compared against those in $G_{p'}^t$ from the first projector p'. Duplicate points are removed and the processing is continued using the same method as for p'. We remove the redundant points so as to ensure that each global correspondence equally contributes to the final calibration solution. Note that generating G_p^t for each individual projector is only necessary when $P_{\min}$ is less than the number of projectors. Otherwise, if it is equal to the number of projectors, it is sufficient to only perform the processing on a single projector. In practice, the value chosen for $P_{\min}$ is dependent on the layout of projectors in the MCP system. If all projectors illuminate a common area, then $P_{\min}$ should be set to the number of projectors in the system. If the projectors do not all share a common observation region, then $P_{\min}$ should be set to the smallest number of projectors illuminating any one portion of the capture region.

Next, the global correspondence points generated from each projector are concatenated into a single matrix to generate the total set of correspondences for the plane position t, $G^t = \{G_1^t, G_2^t, \ldots, G_N^t\}$. The process for plane t is repeated for all the other plane positions, and a final set of correspondences is generated by concatenating the correspondence matrices from each target, $G = \{G^1, G^2, \ldots, G^T\}$. Matrix G represents the correspondence matrix between every sensor in the system. Before this data is applied to a bundle adjustment (BA) framework, we need to provide two other pieces of information as initial conditions to the BA: 1) a coarse estimate of the calibration parameters and 2) an estimate of the 3D positions of the correspondence points.

3.3. Bundle Adjustment

The coarse intrinsic calibration estimate of each camera is generated via Bouguet's camera calibration toolbox [2]. Additionally, we use the stereo calibration feature of the toolbox to estimate the extrinsic relationships between neighboring pairs of cameras. We refer to the process of solving for this extrinsic relationship as "pairwise" calibration. Enough extrinsic pairwise estimates are generated to convert each camera to a defined world coordinate system, which is conveniently chosen to be the same as the coordinate system of one of the cameras. To generate the coarse intrinsic and extrinsic calibration for the projector, we use the Projector-Camera Calibration Toolbox [5] which is built on top of Bouguet's toolbox [2].

The 3D points provided to the bundle adjustment process are generated by triangulating the existing corresponding
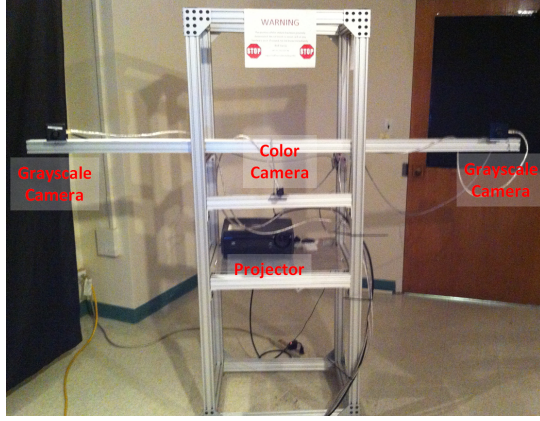
Figure 6. Single station from our multi-view structured light system.

Figure 7. Comparison of the reprojection error using the proposed calibration method vs. pairwise calibration.

points from a pair of cameras. If all correspondences are not visible from a single camera pair, the remaining points are triangulated with another camera pair and then transformed into the world coordinate system. The cameras used to triangulate the 3D points should have a similar field of view. It is convenient to define the world coordinate system with respect to one of the cameras in the pair. The estimates for 3D position, sensor intrinsics, sensor extrinsics, and the correspondence information are input into a bundle adjustment framework [13] to refine sensor intrinsics and extrinsics. The output of the BA provides us the intrinsics of each device, including radial distortion parameters, as well as the extrinsic parameters.

4. Results

The proposed calibration method is demonstrated on a multi-view structured light system consisting of three stations surrounding a central capture area. Each station is equipped with two grayscale cameras, i.e Point Grey Dragonfly Express 640×480, a color camera, i.e. Point Grey Flea 2 640×480, and a video projector, i.e. Optoma TX780 1024×768, as illustrated in Fig. 6. Each camera and projector is pairwise calibrated with respect to a single camera in each station to be used as initial conditions in BA. In Fig. 1, the pairwise calibrations within each station are shown as dashed arrows, and those between stations are shown as solid arrows. In total, nine cameras and three projectors are calibrated. The calibration plane is positioned in eight distinct orientations at the center of the capture area in order to generate the global correspondence points, as described in Section 3, to be used in bundle adjustment. C_{min} and P_{min} are set to 8 and 3, respectively. All three projectors are able to clearly see the projection target positioned in the central capture area, so P_{min} is set to three. We opted to set C_{min} to 8 rather than 9 to accommodate the cases where a camera may not be able to decode the projected pixels due to the
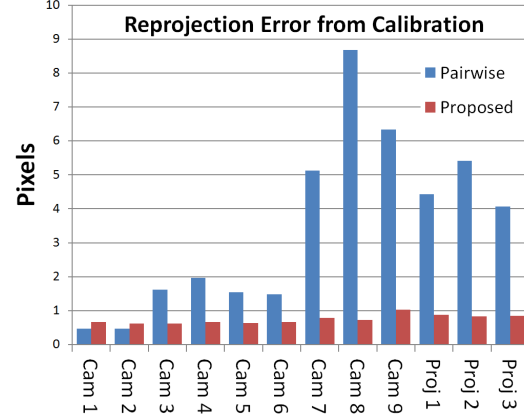
large angle between the central axis of the camera and the normal of the calibration target. In total, 17,262 points are used for calibration. Fig. 7 compares the reprojection error of our proposed calibration method against pairwise calibration. As seen, the reprojection error of sensors is considerably smaller for our method than pairwise calibration. Fig. 8(a) illustrates the projection of 3D correspondence points onto the coordinate frame of camera 7. The misalignment between the projected points and the actual point locations are especially visible on the left hand side of the image. Fig. 8(b) shows an improved alignment with the new calibration method.

MCP calibration greatly affects the quality of 3D reconstruction in the multi-view structured light system of Fig. 1. We compare the accuracy of pairwise calibration with our proposed method on a spherical test object of radius 31 cm by applying phase shifted sinusoidal patterns for 3D reconstruction. Partial reconstruction from each of the six grayscale cameras C_1 through C_6 in Fig. 1 are merged and a least square fit to the spherical point cloud is obtained. Average distance and standard deviation of the points in the point cloud to the sphere are then used as two accuracy metrics to compare the two methods. The average distance of the proposed method and pairwise calibration are 2.7 mm and 11.0 mm respectively. The standard deviation of the proposed method and pairwise calibration are 2.0 mm and 5.9 mm respectively. The results indicate a closer fit to the sphere and hence superior accuracy of the proposed calibration method. Both calibration methods generate an estimate of the sphere's radius that is within the margin of error, i.e. 1.6 mm, of the manually measured radius. Fig. 9 shows slices of the reconstructed sphere along x,y, and z axes for both calibration methods. The double surfacing is clearly visible in Figs. 9(b), 9(d), 9(f) and 9(h).

Fig. 10 shows three views of merged point clouds for a moving person generated by the sinusoidal phase shifted,
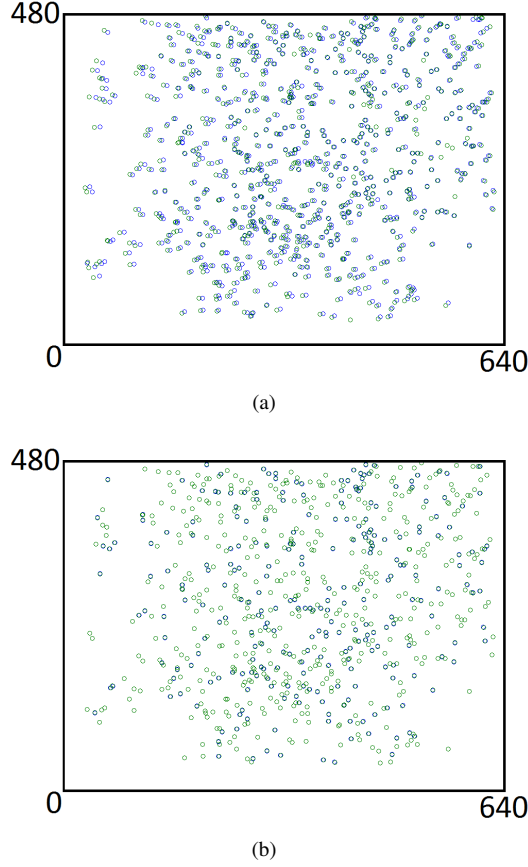
(a)

(b)

Figure 8. Reprojection error using (a) pairwise calibration; (b) proposed method. Image locations are shown in green and the projected image locations in blue.

three-view structured light system shown in Fig. 1, calibrated with our proposed method. The resulting point cloud from each camera is shown in a different color. As seen, the resulting point clouds from all six cameras are well aligned.

5. Conclusions

We have proposed a method for calibrating an MCP system by generating correspondences across all sensors in the system. Aside from generating correspondences, our calibration target opens up the possibility of generating correspondences between projectors. The method is effective even in environments where the viewpoints of all sensors surround a central area and do not necessarily share a field of view.

References

[1] S. Audet and M. Okutomi. A user-friendly method to geometrically calibrate projector-camera systems. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 47–54, June 2009. 1
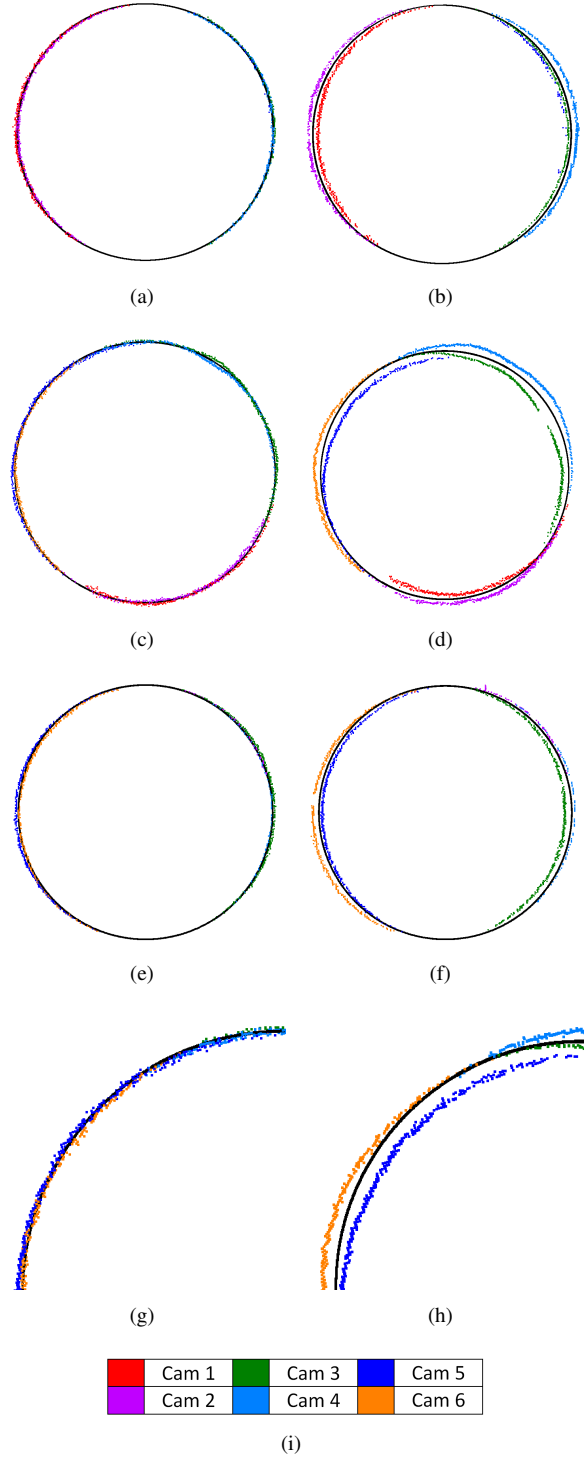
Figure 9. Projection of sphere for the proposed method along (a) x axis, (c) y axis, (e) z axis, and for pairwise calibration for (b) x axis, (d) y axis, (f) z axis; (g) zoomed in portion of (a); (h) zoomed in portion of (b). The points for each camera are assigned a unique color (i).
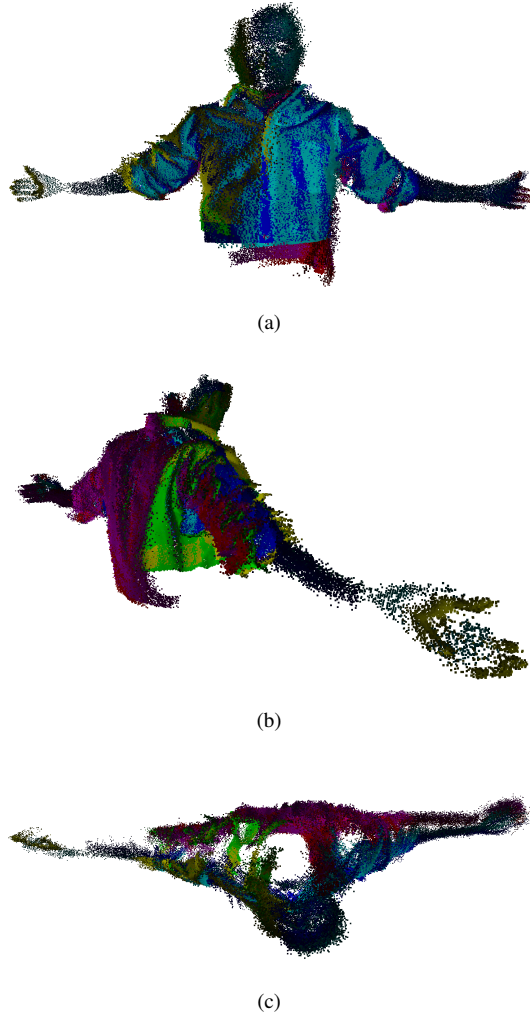
(a)

(b)

(c)

Figure 10. Point cloud of moving person generated from a sinusoidal phase shifting three-view structured light system of Fig. 1 calibrated with our proposed method: (a) front, (b) back, and (c) top view; the partial point clouds for each of the six cameras in three stations are shown in 6 colors: red, green, blue, cyan, magenta, yellow.

[2] J.-Y. Bouguet. Camera calibration toolbox for MATLAB, 2004. 1, 5

[3] X. Chen, J. Xi, Y. Jin, and J. Sun. Accurate calibration for a camera-projector measurement system based on structured light projection. *Optics and Lasers in Engineering*, 47(34):310–319, 2009. 1

[4] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: Audio visual experience automatic virtual environment. *Commun. ACM*, 35(6):64–72, 1992. 1

[5] G. Falcao, N. Hurtos, J. Massich, and D. Fofi. Projector-camera calibration toolbox, 2009. 5

[6] R. Furukawa, R. Sagawa, A. Delaunoy, and H. Kawasaki. Multiview projectors/cameras system for 3D reconstruction of dynamic scenes. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1602–1609, Nov. 2011. 1

[7] R. Furukawa, R. Sagawa, H. Kawasaki, K. Sakashita, Y. Yagi, and N. Asada. One-shot entire shape acquisition method using multiple projectors and cameras. In *Image and Video Technology (PSIVT), 2010 Fourth Pacific-Rim Symposium on*, pages 107 –114, nov. 2010. 2

[8] Y. Furukawa and J. Ponce. Accurate camera calibration from multi-view stereo and bundle adjustment. *International Journal of Computer Vision*, 84:257–268, 2009. 1

[9] A. Griesser, T. Koninckx, and L. Van Gool. Adaptive real-time 3D acquisition and contour tracking within a multiple structured light system. In *Computer Graphics and Applications. Proceedings. 12th Pacific Conference on*, pages 361–370, Oct. 2004. 1

[10] A. Griesser and L. Van Gool. Automatic interactive calibration of multi-projector-camera systems. In *Computer Vision and Pattern Recognition Workshop. CVPRW '06. Conference on*, page 8, June 2006. 1

[11] S. Kobayashi, F. Sakaue, and J. Sato. Multiple view geometry of projector-camera systems from virtual mutual projection. 1

[12] Z. Li, Y. Shi, C. Wang, and Y. Wang. Accurate calibration method for a structured light system. *Optical Engineering*, 47(5):053604, 2008. 1

[13] M. A. Lourakis and A. Argyros. SBA: A Software Package for Generic Sparse Bundle Adjustment. *ACM Trans. Math. Software*, 36(1):1–30, 2009. 6

[14] R. Raskar, M. Brown, R. Yang, W.-C. Chen, G. Welch, H. Towles, B. Scales, and H. Fuchs. Multi-projector displays using camera-based registration. In *Visualization '99. Proceedings*, pages 161–522, Oct. 1999. 1, 2

[15] T. Svoboda, D. Martinec, and T. Pajdla. A convenient multicamera self-calibration for virtual environments. *Presence: Teleoper. Virtual Environ.*, 14(4):407–422, 2005. 1

[16] J.-P. Tardif, S. Roy, and M. Trudeau. Multi-projectors for arbitrary surfaces without explicit calibration nor reconstruction. In *3-D Digital Imaging and Modeling. 3DIM 2003. Proceedings. Fourth International Conference on*, pages 217–224, Oct. 2003. 1, 3

[17] T. Ueshiba and F. Tomita. Plane-based calibration algorithm for multi-camera systems via factorization of homography matrices. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, volume 2, pages 966–973, Oct. 2003. 1

[18] S. Yamazaki, M. Mochimaru, and T. Kanade. Simultaneous self-calibration of a projector and a camera using structured light. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2011 IEEE Computer Society Conference on*, pages 60–67, June 2011. 1

[19] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. Brown. Pixelflex: a reconfigurable multi-projector display system. In *Visualization, 2001. Proceedings*, pages 167–554, Oct. 2001. 1

[20] S. Zhang and P. S. Huang. Novel method for structured light system calibration. *Optical Engineering*, 45(8):083601, 2006. 1