Eurographics Conference on Visualization (EuroVis) (2013)
M. Hlawitschka and T. Weinkauf (Editors)

Short Papers

Automatic geometric calibration of projector-based light field displays

M. Agus and E. Gobbetti and A. Jaspe and G. Pintore and R. Pintus

CRS4 Visual Computing, Pula (Cagliari), Italy

Abstract

We present a novel calibration method for continuous multiview (light field) projection-based displays using a single uncalibrated camera and four fiducial markers. Calibration starts from a simple parametric description of the display layout. First, individual projectors are calibrated through parametric optimization of an idealized pinhole model. Then, the overall display and projector parameterization is globally optimized. Finally, independently for each projector, remaining errors are corrected through a rational 2D warping function. The final parameters are available to rendering engines to quickly compute forward and backward projections. The technique is demonstrated in the calibration of a large-scale horizontal-parallax-only 35MPixels light field display.

1. Introduction

The last decade has witnessed a proliferation of efforts in improving the quality of visualization systems by designing and manufacturing displays able to let people interact with stereoscopic scenes without the need for glasses. Among the various technologies available [GSP*05, WLHR12], multiview projection-based displays appear to be appealing for collaborative visualization systems, because of their ability to present a continuous image to many viewers within a large workspace [MGB*11, ABG*09, MAG*12], due to the high number of view-dependent pixels that contribute to a single image [IGM10, BGMP08]. They are typically composed of an array of projectors, lateral mirrors, and a selectively transmissive screen to produce a light field [ABF*06, BFA*05]. They can emulate the emission from physical objects at fixed spatial locations, providing multiple freely moving viewers natural parallax perceptual cues such that they feel the illusion of interacting with floating objects. Calibration of such kind of displays is a challenging problem since, in theory, the entire viewing workspace should be registered, and thus all the projectors need to be correctly modelled. The problem is further complicated when this needs to be achieved without involving special equipment. In this paper, we present a practical and effective method for geometric calibration of these displays which employs a single uncalibrated camera and does consider only four fiducial markers on the screen and a simple parametric description of the display geometry. The technique has been demonstrated in the calibration of a large scale 35 MPixels horizontal-parallax-only display composed by 72 projectors, and it provided good registration between displayed object space and physical space, thus resulting in high-quality real 3D imagery in the viewing workspace.

2. Related work

When considering the problem of calibrating autostereoscopic displays, the different physical properties and designs impose different alignment methods. Ge [GSP*05] et al. use a stereo camera to calibrate the auto-stereoscopic Varrier system. Annen et al. [AMZ*06] employ a digital camera observing a chessboard pattern to automatically calibrate their rear-projection parallax barrier display. For the calibration of their 360 degree light field display, Jones et al. [JMY*07] consider a simple procedure based on a linear approach requiring six correspondences between fiducial markers placed in the rotating mirror and 2D pixel positions. To the best of our knowledge, no automatic geometric calibration systems targeted to continuous projector-based light field displays have been presented so far. Our method falls in the category of techniques considering a single camera observing the structured patterns projected in the screen [RBY*99]. In this class, various methods have been presented to register multiple projectors displaying on planar surfaces [RP04], or on various kind of curved surfaces [SM11b, SM10, SM11a]. Our method is optimization-based and targeted for light field displays containing lateral mirrors and with planar or curved screens.

3. Light field display overview

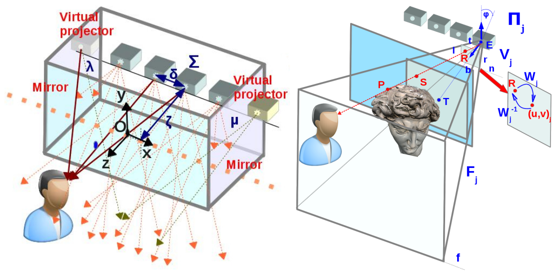

Figure 1: ***Display concept.*** *Left: Each projector emits light beams toward a holographic screen, and lateral mirrors increase the angular resolution. Right: given a ray coming from a projector and passing through the screen, the correspondent pixel is obtained by modeling the projector as a pinhole emitter.*

In a typical projector-based light field design, light emitters (projectors) are densely arranged in a grid behind the screen. Each projector emits light beams toward a subset of the points of the holographic screen, so that each screen point is hit by multiple light beams coming from different projectors (see figure 1 left). The screen is holographically recorded and performs selective directional transmission of light beams. The angular light distribution profile is characterized by a wide plateau and steep Gaussian slopes, resulting in a homogeneous light distribution and continuous 3D view with no visible cross-talk within the field of depth determined by the angular resolution. Two lateral mirrors are commonly located at the sides of the display and reflect back onto the screen the beams that would otherwise be lost, thus increasing the display angular resolution (see figure 1 left). In typical designs (see figure 1 left), generally few parameters are needed to define the display layout Σ: the separation δ between projectors, the distance ζ to screen, and the positions λ and μ of lateral mirrors.

Projecting graphics. In order to project 3D geometry on the display, the light beams leaving the various screen pixels need to be software controlled to be propagated in specific directions, as if they were emitted from physical objects at fixed spatial locations. Thus, reconstructing the light field of a rendered scene amounts to defining functions which map projector pixels to light beams and vice versa. Typically, 3D scenes are projected by casting rays from projector emitters to object points [AGG*08]. For the characteristics of light field displays, the most effective way to model a projector is to consider it similar to a linear pin-hole camera [RP04] composed by a perspective transformation F_j and a rigid body map V_j, and to correct the remaining distortion errors as 2D warping functions (see figure 1 right). Since rational distortion models have been universally proven to effectively correct various kinds of distortions [TGVGMM12], we considered a quadratic rational function, similar to the model of Claus and Fitzgibbon [CF05]:

$$(u,v) = \left(\frac{\Omega_1^T \chi(u_R, v_R)}{\Omega_3^T \chi(u_R, v_R)}, \frac{\Omega_2^T \chi(u_R, v_R)}{\Omega_3^T \chi(u_R, v_R)} \right) \qquad (1)$$

where $R = (u_R, v_R) = F_j V_j S$ is the pixel position after linear transformation, $\chi(u_R, v_R) = [u_R^2, u_R v_R, v_R^2, u_R, v_R, 1]^T$ is the lifting of the pixel position to a six dimensional space, and Ω_i^T are the rows of a $3x6$ matrix containing the rational distortion coefficients. Since the rational warping function is not analytically invertible, a way to obtain the backward correction function is to model it with independent rational distortion coefficients. These simple forward and backward projection schemes can be easily implemented in GPU shaders, and since they rely on a limited number of parameters, they provide a fast and effective method for rendering on projector-based light field displays.

4. Calibration procedure

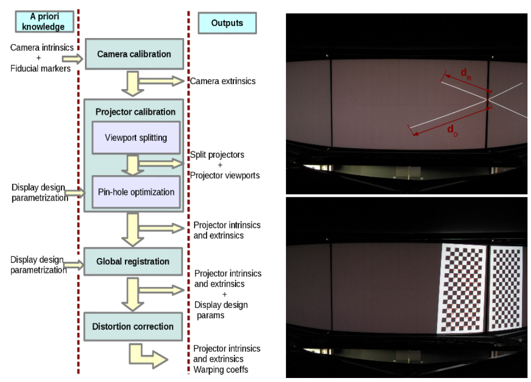

Figure 2: *Left: pipeline of the automatic calibration method. Right: patterns employed for detecting the interference of lateral mirrors (top), and for acquiring control points for calibration (bottom).*

The automatic calibration method finds the correct values of the parameters for the projectors, by employing a single digital camera with the same orientation of projectors and a white diffuser covering the screen in order to enable projecting patterns over it. The technique assumes as priors the intrinsics of the camera, the position coordinates of the screen corners, and a parametric representation of the light field display design (screen and projectors layout). The outputs of the calibration technique are the model parameters (frustums, poses and direct and inverse warping coefficients) of all projectors composing the light field display, which can be used by rendering engines to quickly compute forward and backward projections. The steps of the procedure are indicated in figure 2 left.

Camera calibration. Our method assumes that a single camera C with known intrinsic matrix K_C is employed. An intrinsic matrix with distortion parameters is needed to ensure the best measurement precision possible, and could be obtained by camera calibration methods [Zha00]. The first step of the calibration pipeline consists of finding an estimation of the extrinsic matrix of the camera, in order to map camera pixels to world positions. Given the world coordinates of the fiducial screen corners, the calibration procedure consists of rendering a white uniform image, detecting the screen corners defining the region of interest of the screen, and computing the camera pose V_C.

Projector calibration. In light field displays, the behavior of lateral mirrors complicates the calibration of projectors. In fact, two different types of projectors can be recognized: the central projectors, directly projecting images to the screen, and lateral projectors which project a part of the viewport directly to the screen, while the the other part is projected through a mirror. In order to correctly model these projectors, viewports and frustums need to be split. For this reason, we introduce the concept of **virtual projectors**, whose positions are obtained by reflecting the original projector positions with respect to the mirror (see figure 1 left). In order to take into account the effect of the mirror, the first step of the projector calibration consists of recognizing whether the projector needs to be split or not, and to compute the size of the direct viewport, and eventually of the mirrored one. The process of mirror detection and viewport splitting consists of the following steps: the projector renders a diagonal line pattern, and Hough transform is employed to detect how many line segments are inside the region of interest defined by screen corners: one segment if the projector does project without mirror interference, two segments otherwise (like in figure 2 top right). When two segments are detected, direct and mirrored viewport widths are computed considering their proportionality with respect to the segment lengths (d_D and d_R in figure 2 top right), and frustums are horizontally split using the same proportionality ratios. Once projector viewports are computed, each projector needs to be calibrated according to the projector model considered. The following procedure is carried out: the projector renders a chessboard pattern covering the viewport and containing a significant number of control points $f_j = (u_j, v_j)$ (figure 2 bottom right), and the projector pin-hole model is calibrated by finding the optimal parameters of the mapping function between the screen 3D world coordinates of the control points $w_j = (x_j, y_j, z_j)$ and the correspondent 2D projector pixel coordinates f_j. In order to find the optimized pin-hole model parameters a least-square minimization is carried out, which employs the Levenberg-Marquardt algorithm (LMA) [Lou04].

Global registration: bundle adjustment. After the independent calibration of all projectors, given the redundancy of projector model parameters, and the fact that control points are detected on the screen surface, geometry incongruities between projectors happen. Specifically, projector positions are not guaranteed to be registered with respect to the display layout, and with respect to the lateral mirrors. Since light field projection schemes are dependent on emitter positions (see [AGG*08]), if projectors are not correctly registered, displayed light fields could be non-uniform, and artifacts in the viewing workspace are likely to occur. Thus, a registration step is needed to fix the projector poses with respect to the display layout parameterization. To do so, a global optimization is carried out which computes: the display geometry parameters (layout parameters, lateral mirror equations), and the projector pin-hole parameters (poses and frustums). Bundle adjustment is performed on a complete system containing all chessboard control points employed for the independent calibration of all projectors. The initial values of parameters are set by considering approximate design approximate measures, and forced to respect geometry constraints according to the display layout design. Since the initial values of projector parameters are the results of the previous calibration step, a limited number of iterations is sufficient to ensure convergence to the optimal solution.

Distortion correction. After global registration of all projectors, the remaining error is corrected by considering a 2D warping function modeling the distortion caused by the characteristics of projector lenses as well as other non-linear distortions caused by defects in projectors or mirrors. In this final step, each projector is calibrated separately, by considering the rational warping function in equation (1). Least square optimization is carried out by employing LMA, where projectors frustums and poses are constrained to be close to the values obtained after the global registration, and warping coefficients are optimized. The inverse warping function is modeled with a similar rational function, whose coefficients are computed by swapping input and output points and by employing least-square optimization LMA.

5. Results

The calibration technique has been implemented on a Linux system in C++ employing OpenCV library for the pattern detection and the *levmar* C++ library [Lou04] for the Levenberg-Marquardt non-linear least square optimization. The calibration system has been tested on a large scale horizontal-parallax-only light field display model HV640RC manufactured by Holografika [ABF*06] (see figure 4 right). The display can visualize 35MPixels by composing images generated by 72 SVGA LED commodity projectors SAMSUNG SP-P300ME with 800x600 resolution illuminating a 160x90cm holographic screen. The system is driven by a cluster of 18 commodity PCs, each one driving 4 projectors. The display provides continuous horizontal parallax within an approximately 50° horizontal field-of-view, with 0.8° angular accuracy. The pixel size on the screen surface is 1.5mm. The geometry design of the HV640RC light field display has the following characteristics (see figure 4 left): a cylindrical screen; a cylindrical layout with projec-

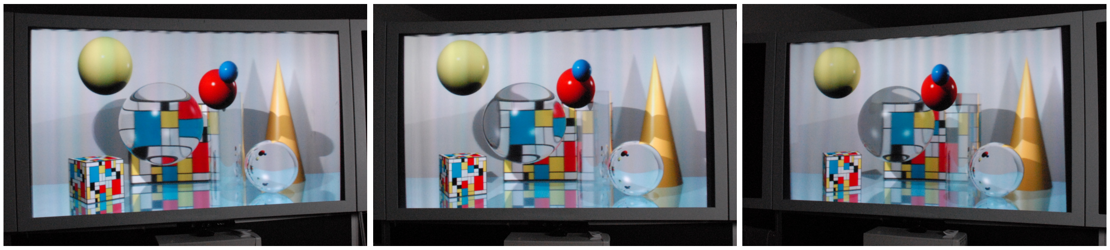

Figure 3: ***Static scenes rendered in a 35MPixel large scale HPO light field display.*** *Pictures acquired live with a hand held camera. The calibrated system is able to provide correct motion parallax effects and a good registration between displayed object space and physical space.*

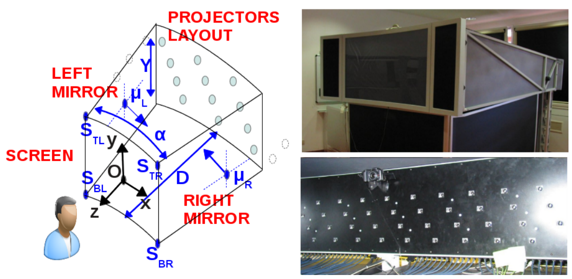

Figure 4: ***Test case: calibration of a cylindrical HPO display.*** *Left: schematic layout. Right: pictures taken from the light field display HV640RC highlighting the screen layout, the projectors layout and the position of camera.*

tors spaced uniformly along the horizontal direction on four rows; two lateral mirrors connecting the border of the projectors cylindrical layout to the border of the screen layout.

Calibration results. The complete automatic calibration of the large scale light field display required around 25 minutes: around 15 minutes and 30 seconds for the independent projector calibration, 7 minutes and 45 seconds for the global registration, and one minute and 45 seconds for the distortion correction. The projector calibration was carried out on the light field display server PC equipped with Athlon64 3300+ CPU. Global registration and distortion correction were performed offline on a Laptop DELL Studio XPS. Global registration required 63 iterations. To evaluate the effectiveness of the calibration procedure, we performed a quantitative measurement of the pixel error of the chessboard control points employed for the calibration, and computed statistics with respect to the root mean square error for each projector of the HV740RC light field display. Figure 5 shows the errors for projectors, comparing the calibration accuracy after the three main steps of the overall calibration. It is evident the improvement of root mean square error for each projector, after the global registration step and the final distortion correction. Another indicator of the calibration accuracy is the number of samples with error above a given pixel threshold.

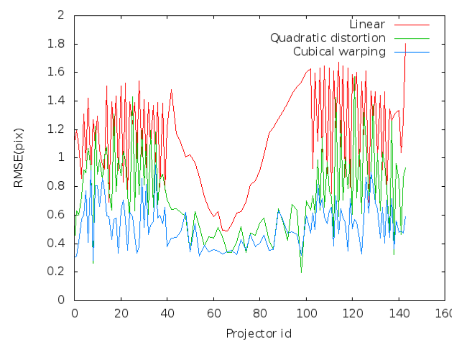

Figure 5: ***Comparison of accuracy after the different steps of calibration of HV640RC display.*** *Errors (RMSE) on chessboard control points.*

For a threshold of 2 pixels, over a total number of 33012 control points employed for calibration (more than 400 for each projector), the number of samples with big error was 8722 after projector calibration, 6528 after global registration, and only 1 sample after distortion correction. From this indicator and figure 5, it is evident that the rational warping model dramatically reduce distortion errors.

Qualitative evaluation. In order to evaluate the overall quality of projected scenes, we also rendered a number of static virtual scenes on the light field display and we inspected them from different positions in the viewing workspace. We recorded the rendered scenes using a hand-held video camera freely moving in the display workspace (see the accompanying video). It is obviously impossible to fully convey the impression provided by the light field display on paper or video. Pictures acquired live with a hand held camera are shown in figure 3. In any case, it appears evident that the calibrated system is able to provide correct motion parallax effects and a good registration between displayed object space and physical space, which demonstrate the correct registration of projectors as well as the multi-user capability of the display.

Acknowledgments. This work is partially supported by the EU FP7 Program under the DIVA project (REA Agreement 290277).

References

[ABF*06] AGOCS T., BALOGH T., FORGACS T., BETTIO F., GOBBETTI E., ZANETTI G.: A large scale interactive holographic display. In *Proc. IEEE VR Workshop on Emerging Display Technologies* (2006). CD ROM Proceedings. 1, 3

[ABG*09] AGUS M., BETTIO F., GIACHETTI A., GOBBETTI E., IGLESIAS GUITIÁN J., MARTON F., NILSSON J., PINTORE G.: An interactive 3d medical visualization system based on a light field display. *The Visual Computer 25*, 9 (2009), 883–893. 1

[AGG*08] AGUS M., GOBBETTI E., GUITIÁN J. A. I., MARTON F., PINTORE G.: GPU accelerated direct volume rendering on an interactive light field display. *Computer Graphics Forum 27*, 2 (2008), 231–240. 2, 3

[AMZ*06] ANNEN T., MATUSIK W., ZWICKER M., PFISTER H., SEIDEL H.-P.: Distributed rendering for multiview parallax displays. In *Proceedings of Stereoscopic Displays and Virtual Reality Systems XIII* (2006), pp. 231–240. 1

[BFA*05] BALOGH T., FORGACS T., AGOCS T., BALET O., BOUVIER E., BETTIO F., GOBBETTI E., ZANETTI G.: A scalable hardware and software system for the holographic display of interactive graphics applications. In *EUROGRAPHICS 2005 Short Papers Proceedings* (Conference Held in Dublin, Ireland, August 2005, 2005). 1

[BGMP08] BETTIO F., GOBBETTI E., MARTON F., PINTORE G.: Scalable rendering of massive triangle meshes on light field displays. *Computers & Graphics 32*, 1 (February 2008), 55–64. 1

[CF05] CLAUS D., FITZGIBBON A. W.: A rational function lens distortion model for general cameras. In *IEEE CVPR* (2005), vol. 1, pp. 213–219. 2

[GSP*05] GE J., SANDIN D., PETERKA T., MARGOLIS T., DEFANTI T.: Camera based automatic calibration for the varrier-system. *CVPR* (2005), 110. 1

[IGM10] IGLESIAS GUITIÁN J. A., GOBBETTI E., MARTON F.: View-dependent exploration of massive volumetric models on large scale light field displays. *The Visual Computer 26*, 6–8 (2010), 1037–1047. 1

[JMY*07] JONES A., MCDOWALL I., YAMADA H., BOLAS M. T., DEBEVEC P. E.: Rendering for an interactive 360 degree light field display. *ACM Trans. Graph 26*, 3 (2007), 40. 1

[Lou04] LOURAKIS M.: Levmar: Levenberg-marquardt nonlinear least squares algorithms in C/C++. www.ics.forth.gr/ lourakis/levmar/, Jul. 2004. 3

[MAG*12] MARTON F., AGUS M., GOBBETTI E., PINTORE G., BALSA RODRIGUEZ M.: Natural exploration of 3d massive models on large-scale light field displays using the fox proximal navigation technique. *Computers & Graphics 36*, 8 (December 2012), 893–903. 1

[MGB*11] MARTON F., GOBBETTI E., BETTIO F., IGLESIAS GUITIÁN J., PINTUS R.: A real-time coarse-to-fine multiview capture system for all-in-focus rendering on a light-field display. In *Proc. 3DTV Conference: The True Vision - Capture, Transmission and Display of 3D Video (3DTV-CON)* (2011), pp. 1–4. 1

[RBY*99] RASKAR R., BROWN M. S., YANG R., CHEN W.-C., WELCH G., TOWLES H., SEALES W. B., FUCHS H.: Multi-projector displays using camera-based registration. In *IEEE Visualization* (1999), pp. 161–168. 1

[RP04] RAIJ A., POLLEFEYS M.: Auto-calibration of multi-projector display walls. In *Proceedings of the Pattern Recognition* (2004), ICPR '04, pp. 14–17. 1, 2

[SM10] SAJADI B., MAJUMDER A.: Automatic registration of multiple projectors on swept surfaces. In *Proceedings of the 17th ACM VRST* (2010), pp. 159–166. 1

[SM11a] SAJADI B., MAJUMDER A.: Autocalibrating tiled projectors on piecewise smooth vertically extruded surfaces. *IEEE TVCG 17*, 9 (Sept. 2011), 1209–1222. 1

[SM11b] SAJADI B., MAJUMDER A.: Automatic registration of multi-projector domes using a single uncalibrated camera. *Comput. Graph. Forum 30*, 3 (2011), 1161–1170. 1

[TGVGMM12] TANG Z., GROMPONE VON GIOI R., MONASSE P., MOREL J.-M.: Self-consistency and universality of camera lens distortion models. Electronic version in hal, 2012. 2

[WLHR12] WETZSTEIN G., LANMAN D., HIRSCH M., RASKAR R.: Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting. *ACM Trans. Graph. (Proc. SIGGRAPH) 31*, 4 (2012), 1–11. 1

[Zha00] ZHANG Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22* (2000), 1330–1334. 3