Virtual Reality manuscript No.
(will be inserted by the editor)

Shadow Removal of Projected Imagery by Occluder Shape Measurement in a Multiple Overlapping Projection System

Daisuke Iwai · Momoyo Nagase · Kosuke Sato

Received: date / Accepted: date

Abstract This paper presents a shadow removal technique for a multiple overlapping projection system. In particular, this paper deals with situations where cameras cannot be placed between the occluder and projection surface. We apply a synthetic aperture capturing technique to estimate the appearance of the projection surface, and a visual hull reconstruction technique to measure the shape of the occluder. Once the shape is acquired, shadow regions on the surface can be estimated. The proposed shadow removal technique allows users to balance between the following two criteria: the likelihood of new shadow emergence and the spatial resolution of the projected results. Through a real projection experiment, we evaluate the proposed shadow removal technique.

Keywords Shadow removal · Multiple overlapping projection · Synthetic aperture capturing · Visual hull

1 Introduction

Projection-based mixed/augmented reality (MR/AR) visually augments physical objects by controlling the appearance of their surfaces with projected imagery from distributed multiple projectors (Raskar et al, 2001). This approach has several advantages such as wide field-of-view imagery, natural auto-stereoscopic vision and eye accommodation, and high spatial and geometric fidelity. Consequently, many researchers have employed it in their interactive systems where users can directly manipulate the appearance of physical objects by touching surfaces with tracked hand-held tools (Bandyopadhyay et al, 2001; Jones et al, 2010; Low et al, 2001). In such systems, the user's body (e.g., hands) or the tools can easily and inadvertently block the projected light, and consequently, cast shadows on the surfaces. These shadows must be removed to guarantee natural and smooth interaction, as they potentially occlude important projected information and detract from the users' visually immersive experience.

One of the most promising shadow removal solutions is to employ multiple overlapping projectors. When the shape and position of the occluder is known, its shadow from a projector onto a surface can be estimated and thus removed by projecting a compensation image from another unoccluded projector. However, in general, such geometric information about the occluder is not known a priori and thus needs to be measured online. The Kinect sensor[1] achieves this by measuring the shape of a scene on the basis of real-time dot pattern projection. However, we cannot estimate shadows from multiple projectors unless the whole occluder shape is measured. To solve this issue, we need to simultaneously use multiple Kinect sensors placed at different locations, but in this case, the projected patterns interfere with each other and the measurement becomes inaccurate. In addition, when the desired appearance for a shadow area is divided and then assigned to unoccluded projectors for shadow removal, the assigned pixel values affect the projected image quality (e.g., spatial resolution), and influence the likelihood of the emergence of a new shadow caused by the occluder's movement. However, previous studies have not focused on this issue.

D. Iwai · M. Nagase · K. Sato
Graduate School of Engineering Science, Osaka University
1-3 Machikaneyama, Toyonaka, 560-8531 Osaka, Japan
Tel.: +81-6-68506371
Fax: +81-6-68506341
E-mail: daisuke.iwai@sys.es.osaka-u.ac.jp

[1] Microsoft, Kinect: `http://www.xbox.com/kinect/`

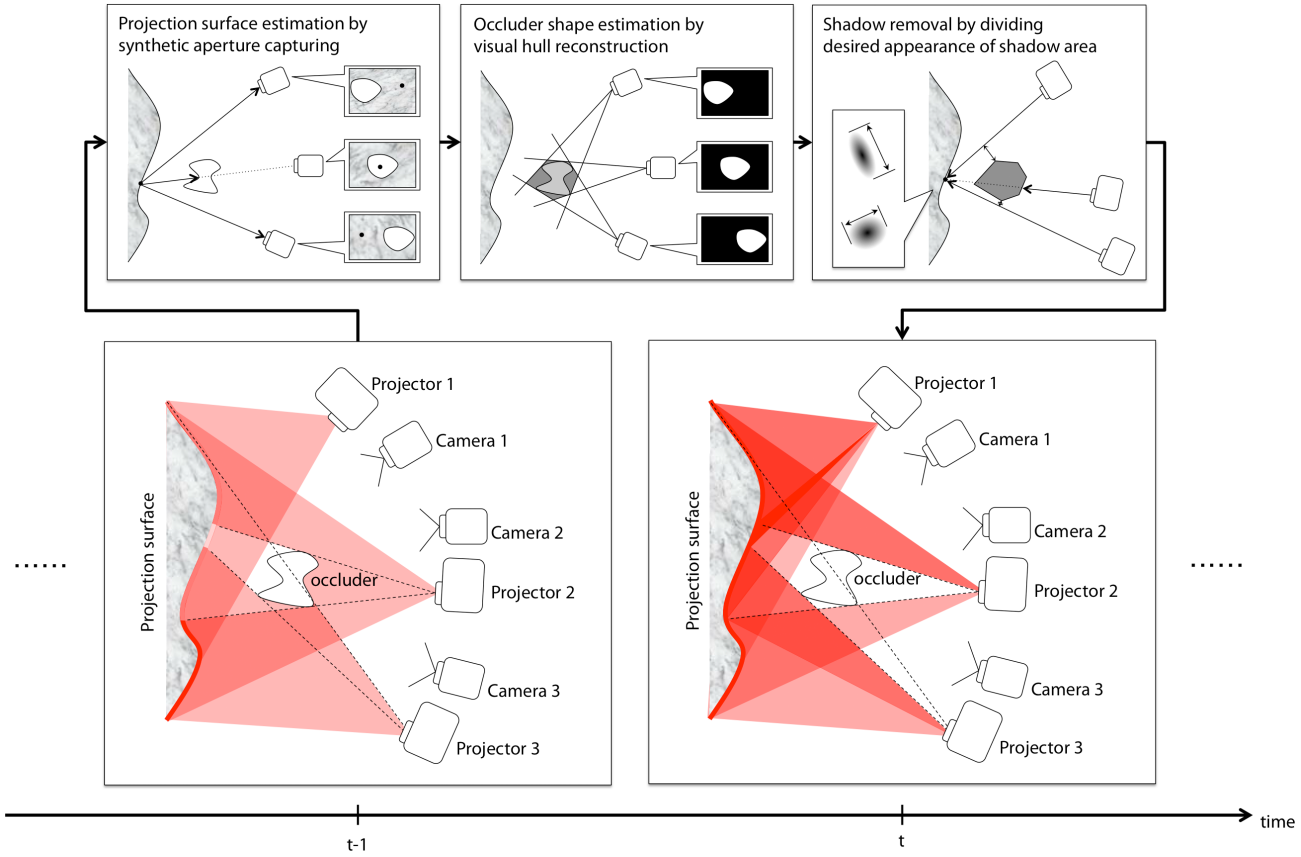

Fig. 1 Overview of the proposed technique.

This paper presents two basic technologies to achieve shadow removal in a multiple overlapping projection system by solving the issues described above. The first is an occluder shape reconstruction technique, and the second is a dividing technique to achieve the desired appearance that should be displayed in the shadow region.

To ensure that the shadows are completely removed, the whole shape of the occluder should be measured. This measurement can be approximate as long as the measured shape contains the occluder. Therefore, we propose to apply a visual hull technique (Laurentini, 1994) using multiple distributed cameras. For extracting the occluder's silhouette, we estimate its background (i.e., projection surface) on the basis of a synthetic aperture capturing technique (Vaish et al, 2006). Once the shape is estimated, we can remove the shadow from the projection surface by dividing the desired appearance (or pixel value) for the shadow area and then assigning the divided pixel values to unoccluded projectors. However, in the next frame, a new shadow may emerge when the occluder moves into the view frustum of an unoccluded projector. Therefore, it is obviously reasonable to assign a large pixel value to an unoccluded projector whose view frustum is at some distance from the occluder, although it is possible that such a projector provides poor image quality. Note that, when we refer to the image quality, we consider the spatial resolution of the projected imagery. Therefore, we propose a new projector selection technique that considers both the likelihood of a new shadow emerging and the projected image quality. Figure 1 outlines the proposed technique.

2 Related Works

In the past decade, several approaches have been proposed for shadow removal in front projection display systems. All these methods employ multiple overlapping projectors, and single or multiple cameras to capture either a shadow on the projection surface, or an occluder.

Sukthankar et al. proposed a method to remove shadows on a projection surface while displaying a still image (Sukthankar et al, 2001). They applied a feedback process where each iteration compared the current appearance of the projection surface with a reference image captured in advance and then generated a projected image that minimized the residual. The extended version of the method realized the suppression of blind-

ing light incident on a user (Cham et al, 2003). Another approach focused on displaying shadowless video footage by predicting an unoccluded appearance on a projection surface by considering the geometric and radiometric properties of the surface (Jaynes et al, 2001, 2004). It compared the captured scene with the predicted appearance to find a shadow region that was then illuminated by another unoccluded projector. Sugaya et al. also proposed a method that required only a single-shot image of the projection surface (Sugaya et al, 2010). They assigned different intensity values to each projector to identify occluded projectors from the shadows. The techniques described above assume that the cameras observing the projection surface are placed between the surface and the occluders.

Summet et al. relaxed this constraint (Summet et al, 2007). They detected an occluder instead of a shadow by illuminating the scene with infrared (IR) lights and capturing the projection surface with an IR camera. Then, the captured backlit silhouette of an occluder was warped to align with the projection surface. The camera does not have to be placed behind the occluder in this method. However, it is assumed that IR light sources are placed between the surface and occluder for robust silhouette extraction.

Audet et al. proposed recovering 3D occluder information with two cameras (Audet and Cooperstock, 2007). Once the 3D information was obtained, the position of the shadows of the occluder on the projection surface could be deduced from the geometric relationships among the projectors, surface, and occluder. There is no limitation on the placement of cameras in their method. As they focused on the compensation of shadows caused by walking people, the current dedicated system is designed to deal with humans standing vertically on the ground.

Most of the previous studies assumed that equipment, such as cameras or IR light sources, is placed between the occluder and the projection screen, thus constraining the user from being close to the projection surface. The central aim of this research is to relax this constraint. To this end, we reconstruct the shape of the occluder by applying a visual hull technique in order to address the shape reconstruction of various kinds of occluders, such as the user's body or a hand-held tool. In addition, none of the previous techniques explicitly considered how to divide the desired appearance for a shadow area among unoccluded projectors. This paper also presents a solution to this issue.

3 Shadow Removal Principle

The proposed shadow removal technique employs multiple overlapping projectors and cameras. This section describes the assumptions of the proposed technique, followed by the principle of each technical component of the proposal.

3.1 Assumptions

We assume that multiple projectors and cameras are distributed in such a way that they face toward a projection surface in our system. This assumption is reasonable as recent projection-based MR/AR applications employ the same setup (Bimber et al, 2005).

Geometric correction and radiometric compensation of projected images are required to display the desired images from the projectors. Because various techniques to solve these issues already exist, as summarized in (Bimber et al, 2008), we apply these existing techniques, in particular that in (Y.I. and H.M., 1971) for the geometric correction and that in (Yoshida et al, 2003) for radiometric compensation. To this end, we measure the shape of the projection surface as well as the poses and positions of the projectors and cameras in advance. However, the shape, position, and pose of the occluder remain unknown. We also assume that the reflectance property of the projection surface is Lambertian. Note that these assumptions are common in most projection-based MR/AR systems.

3.2 Projection Surface Estimation by Synthetic Aperture Capturing

To ensure that shadows from the projectors on the projection surface are completely removed, the whole shape of the occluder should be measured, but the measurement can be approximate as long as the measured shape contains the occluder. To meet this requirement, we apply a visual hull technique to reconstruct the occluder shape. Visual hull reconstruction obtains the shape of an object as an intersection of multiple silhouettes captured from different positions.

We need to estimate the background appearance for the silhouette extraction of the visual hull reconstruction. Note that in this study the background is identical to the projection surface. A fixed background appearance captured in advance is not suitable in our context, because it is highly possible that while interacting with the system, the user's movement affects the illumination from the environment lighting, and consequently, the background subtraction fails. Therefore, we

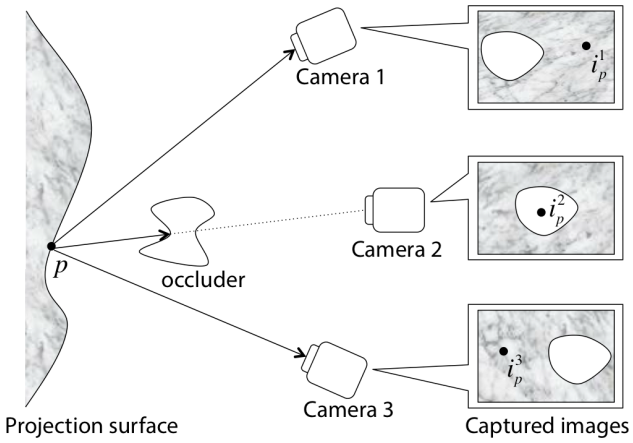

Fig. 2 Synthetic aperture capturing technique.

employ a synthetic aperture capturing technique (Vaish et al, 2006) to estimate the projection surface for every frame. The synthetic aperture capturing technique realizes a virtual camera with an incredibly large aperture by digitally aligning multiple images taken from distributed cameras. Owing to the large aperture, the camera can "see-through" objects or occluders that are smaller than the aperture and placed between the cameras and the virtual focal plane.

In particular, we independently estimate the appearance a_p of each point p $(= 1, \ldots, N_p)$ on the projection surface. For each point p, we select a color value from each camera c $(= 1, \ldots, N_c)$ at the corresponding pixel in its captured image $I_c(u, v)$. The point-to-pixel correspondence can be acquired on the basis of geometric information, such as the shape of the surface and position/pose of the camera relative to the surface, which is calibrated offline. We represent i_p^c as the color value of the corresponding pixel in the captured image of camera c (Fig. 2). We finally take the median of the color values of all cameras in the system to estimate the appearance a_p of point p.

$$a_p = median(i_p^1, \ldots, i_p^{N_c}). \tag{1}$$

We can estimate the appearance of the entire surface by simply repeating the process described above for the other points on the surface. Then, the background image $B_c(u, v)$ of each camera c is synthesized by aligning the estimated appearance of each point a_p on the basis of the point-to-pixel correspondence described above.

The offline geometric calibration is done as follows. First we place a physical calibration object, on which visual markers that define the world coordinate system are painted, in front of the projection surface. Second, the cameras capture the calibration object to acquire 2D (camera coordinates) and 3D (world coordinate) correspondences. From the correspondences, we calibrate the cameras' intrinsic and extrinsic parameters using the direct linear transformation (DLT) method (Y.I. and H.M., 1971). Third, the projectors project structured light patterns, such as the gray code pattern (Sato and Inokuchi, 1987), onto the calibration object to acquire 2D (projector coordinates) and 3D (world coordinate) correspondences. The projectors' intrinsic and extrinsic parameters are then calibrated from the correspondences using the DLT method. Finally, we measure the shape of the projection surface in the world coordinate system. We project structured light patterns from one of the calibrated projectors to the projection surface, and capture the reflections by one of the cameras, to acquire pixel correspondences between the projector and the camera. From the correspondences and the calibration results of the projector and camera, we measure the 3D shape of the surface in the world coordinate system (Sato and Inokuchi, 1987).

3.3 Visual Hull Reconstruction

Our visual hull reconstruction technique is based on the technique previously proposed by (Laurentini, 1994). The binary silhouette image $S_c(u, v)$ of camera c is computed by thresholding the subtraction of the background image B_c from the captured image I_c as follows:

$$S_c(u, v) = \begin{cases} 1, \text{ if } |I_c(u, v) - B_c(u, v)| > t_s, \\ 0, \text{ otherwise.} \end{cases} \tag{2}$$

where t_s represents the predefined threshold.

After acquiring the silhouette images from all the cameras in the system, we reconstruct the shape from them. We assume a tessellation of the reconstruction space into discrete voxels $V(x, y, z)$. The voxels are binarily labeled as either transparent or opaque, where the latter represents the element of the occluder's volume.

We project the voxels on camera image planes and carve V by labeling them as transparent when at least one projection corresponds to the pixel value of 0 in a silhouette image S_c. We denote the opaque and transparent voxels as V_o and V_t, respectively. All opaque voxels V_o belong to the visual hull that encloses the object (Fig. 3).

3.4 Shadow Removal

Once the 3D information of an occluder is recovered, we can compute which projector is visible from each point on the projection surface. Our technique can project compensation images onto the shadow region from unoccluded projectors to remove it. However, in some cases,

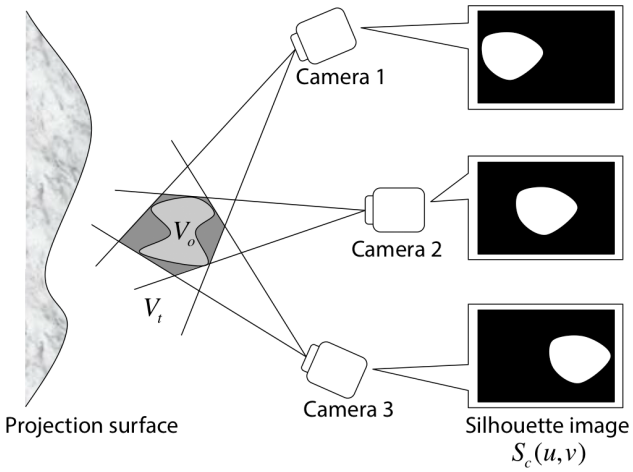

Fig. 3 Visual hull reconstruction technique.

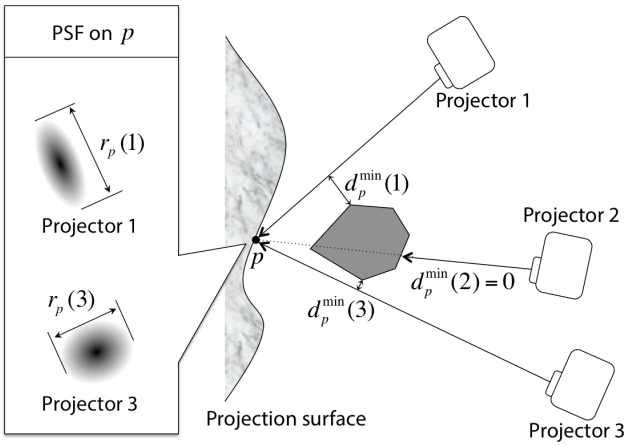

Fig. 4 Parameters for shadow removal.

this is not desirable if the target appearance is divided and assigned to the unoccluded projectors by considering only visibility. This is because a new shadow may emerge in the next frame when the occluder moves into the view frustum of the unoccluded projector. In such a case, it is obviously advisable to assign a large pixel value to the unoccluded projector whose view frustum is distant from the occluder in order to avoid the emergence of another shadow in the next frame. However, such a projector may provide poor image quality (hereafter referred to as low spatial resolution) due to a steep grazing angle of incident light rays or a large distance from the projector to the surface.

Therefore, we propose a new dividing technique for the desired appearance of a shadow area, which considers the following three criteria: (1) the visibility of projectors, (2) the likelihood of new shadow emergence, and (3) the spatial resolution of projected imagery. In particular, the target appearance $a_p^t(l)$ at point p that a projector l $(= 1, \ldots, N_l)$ should display is decided by dividing the final target appearance a_p^t according to the above-mentioned criteria. In other words, our technique decides the weights $w_p(l)$ $(0 \leq w_p(l) \leq 1)$ to divide the target appearance:

$$a_p^t(l) = a_p^t w_p(l), \tag{3}$$

where

$$\sum_{l=1}^{N_l} w_p(l) = 1. \tag{4}$$

We design the weights so that the second and third criteria can be balanced by each user. Therefore, the weights are decomposed as per the following equation:

$$w_p(l) = \alpha w_p^L(l) + (1-\alpha) w_p^S(l), \tag{5}$$

where

$$\sum_{l=1}^{N_l} w_p^L(l) = 1, \ \sum_{l=1}^{N_l} w_p^S(l) = 1, \ 0 \leq \alpha \leq 1. \tag{6}$$

Here, α, w_p^L and w_p^S are the user-defined balance coefficient, the weight for the second criterion (i.e., the likelihood of new shadow emergence), and that for the third criterion (i.e., spatial resolution), respectively.

Visibility of projector. For each point p on the surface, we check the visibility of each projector l $(= 1, \ldots, N_l)$. This is done by comparing the actual distance from the point to the projector and the depth value from the depth map computed by rendering the scene, including the projection surface and occluder, from the projector's viewpoint. If the distance and depth values are different, the projector is regarded as invisible from the point. We define a binary function $vis_p(l)$ that takes the value 1 if the projector l is visible from the point p.

$$vis_p(l) = \begin{cases} 1, \text{ if } l \text{ is visible from } p, \\ 0, \text{ otherwise.} \end{cases} \tag{7}$$

Likelihood of new shadow emergence. For each point p, we calculate a distance $d_p(l, V_o)$ from each opaque voxel $V_o(x_o, y_o, z_o)$ to the light ray emitted from each visible projector to the point. Then we search for the minimum distance $d_p^{min}(l)$ for each visible projector:

$$d_p^{min}(l) = \min_{V_o} d_p(l, V_o) vis_p(l). \tag{8}$$

When we consider the likelihood of new shadow emergence, we assign weighting values to the visible projectors so that the weights are proportional to the minimum distances (Fig. 4). Thus,

$$w_p^L(l) = \frac{d_p^{min}(l)}{\sum_{l'=1}^{N_l} d_p^{min}(l')}. \tag{9}$$

Spatial resolution. The spatial resolution of the projected result depends on the point spread function (PSF) of the projected pixel, because the result can be represented as the convolution of the original image and the PSF. Researchers have proposed projector selection techniques for multi-projector systems in the context of extending the depth of field of the systems (Bimber and Emmerling, 2006; Nagase et al, 2011). In these works, for each point on a projection surface, one of the projectors is selected to display a pixel on that point. For example, one of the previous works (Bimber and Emmerling, 2006) selected the projector with the smallest PSF. An alternative technique was proposed by (Nagase et al, 2011); it approximated the projector pixel PSF as an isotropic Gaussian function, and consequently, the shape of the PSF as an ellipse. Then, it selected a projector whose major axis of the PSF ellipse had the shortest length. Our method is based on the latter technique because it has been shown in (Nagase et al, 2011) that it provides better image quality. In particular, for each point p, we calculate the major axis length $r_p(l)$ of the PSF of a pixel projected from each projector l to the point (Fig. 4). Then, we compute the maximum length r_p^{max} of the major axes at the point p among the projectors. Thus,

$$r_p^{max} = \max_l r_p(l). \tag{10}$$

We assign weighting values to the visible projectors so that the weight is high when the length is small. Thus,

$$w_p^S(l) = \frac{\frac{1}{r_p(l)} vis_p(l) - \frac{1}{r_p^{max}}}{\sum_{l'=1}^{N_l} (\frac{1}{r_p(l')} vis_p(l) - \frac{1}{r_p^{max}})}. \tag{11}$$

If there is no occlusion, we assign the evenly divided appearance of the original image to the projectors. Thus,

$$w_p(l) = \frac{1}{N_l}. \tag{12}$$

4 Experiment

We conducted a proof-of-concept experiment to evaluate the proposed shadow removal technique on a real projection system prototype. The prototype projector-camera system was implemented as shown in Fig. 5. The system consisted of four projectors (Acer K10, 100 ANSI Lumen, 800×600 pixel) and five cameras (Point Grey Research Chameleon, 1290×960 pixel), which were connected to and controlled by a single PC. We assigned a square region on a planar surface as the projection surface of the experiment. The projectors and cameras were placed so that they were facing the projection surface. We calibrated the projectors and cameras in advance as described in Section 3.1 and 3.2. We prepared a sphere as an occluder. Because our system estimated the background in every frame, we could stably extract the silhouette of the occluder without having to assign a large value to the threshold t_s. We particularly used $t_s = 20$ in the range of $0 \leq t_s \leq 255$. In a general shape reconstruction context, the size of the voxel tessellation in the visual hull reconstruction must be small enough to measure an object's shape as precisely as possible. However, in the case of our system, the voxel does not need to be so small because the measurement can be approximate as long as the measured shape contains the occluder. Therefore, we set the size of the voxel tessellation as 5 mm in the following experiment.

The experiment involved four successive frames; we moved the occluder in every frame. Figure 6 shows the scene of each frame captured by a camera in the system. We projected an image of two parrots (Fig. 7) onto the rectangle region, with shadows caused by the occluder in every frame. From the first frame to the

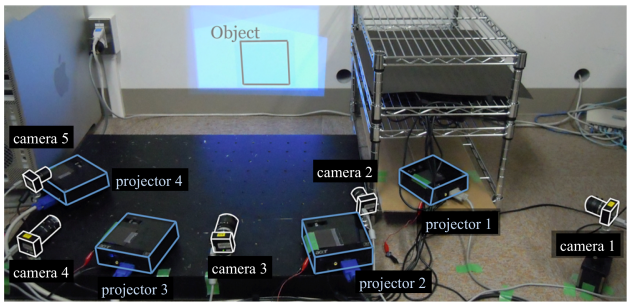

Fig. 5 Overview of the system.

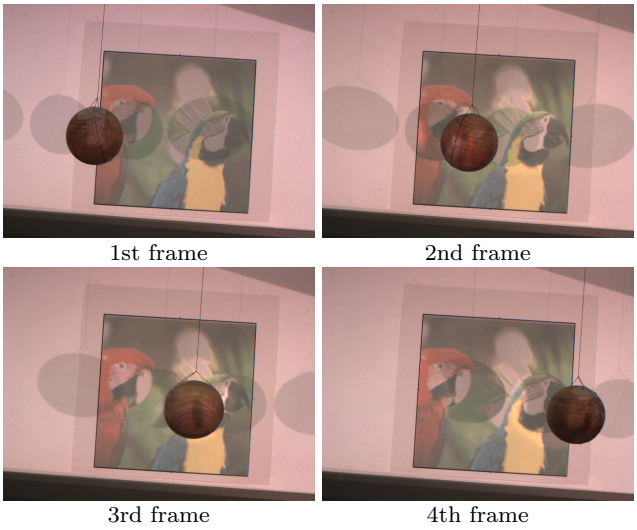

Fig. 6 Captured scene of each frame.

Fig. 7 Target image.

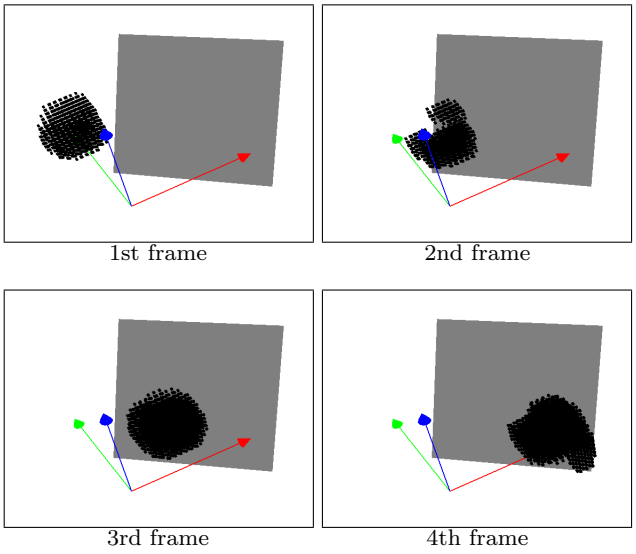

Fig. 8 Results of visual hull reconstruction (gray quadrilateral: the projection surface, black dots: reconstructed voxels).

last, we moved the occluder from the left to the right of the scene, occluding the projection lights from the first frame. Note that there was no occluder before the first frame. The movement of the occluder could be repeated so that we could compare the shadow removal results under different conditions. The experiment was conducted under two conditions where the desired appearance of the shadow area was divided and assigned to unoccluded projectors on the basis of either (1) the likelihood of new shadow emergence (i.e., $\alpha = 1.0$ in Eq. 5) or (2) the spatial resolution of the projected result (i.e., $\alpha = 0.0$). We denote the former condition as the **LNSE condition** and the latter as the **SR condition**. Under both conditions, the movement of the occluder was the same as that shown in Fig. 6.

Figure 8 shows the results of visual hull reconstruction. Figure 9 shows the experimental results for shadow removal under the two conditions. The figures were obtained by the synthetic aperture capturing technique described in Section 3.2. The projected results under

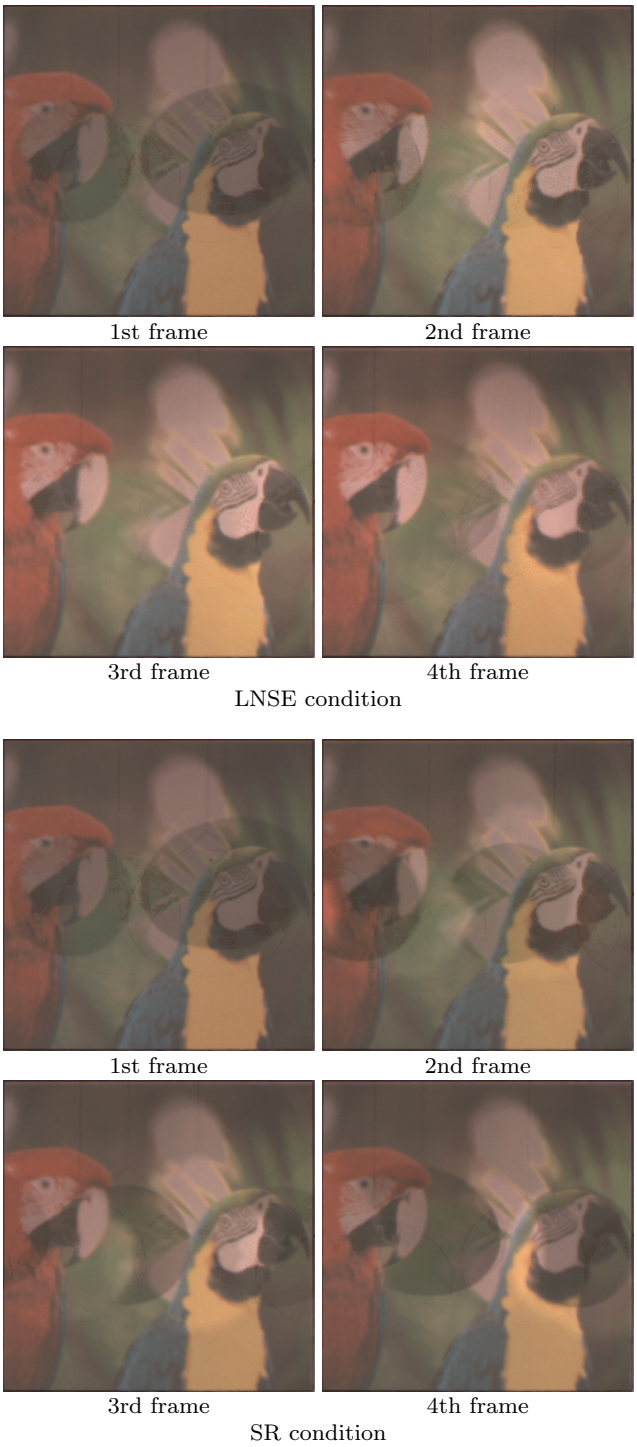

Fig. 9 Shadow removal results obtained by synthetic aperture capturing.

Frame	SSIM (LNSE condition)	SSIM (SR condition)
1st	0.652	0.652
2nd	0.690	0.661
3rd	0.696	0.667
4th	0.695	0.664

Table 1 SSIM evaluation results.

LNSE condition SR condition

Fig. 10 Magnified view around the face of the right parrot in the fourth frame. The SR condition provides a better image quality than the LNSE condition, that displays unclear smeared pixels.

the SR condition showed that the new occluder shadows could not be removed in all the frames. On the other hand, the results under the LNSE condition showed that the shadows could be removed in the second and fourth frames more effectively than for the SR condition; furthermore, they were completely removed in the third frame. Figure 10 shows two magnified views around the face of the right parrot in the fourth frame. The SR condition decreases the weights of projectors that project a pixel over a wider area to achieve the best image quality in terms of spatial resolution. Therefore, at a place in the projection surface on which no shadow is cast, the SR condition can provide a better image quality than the LNSE condition, which does not explicitly consider the spatial resolution. Consequently, as shown in Fig. 10, the LNSE condition degrades the image quality of the displayed result that displays unclear smeared pixels. We objectively evaluated the image quality using the structural similarity index (SSIM), which is a method for assessing the perceptual quality of a distorted image when compared to the original (Fig. 7 in this experiment) (Wang et al, 2004). Table 1 shows the result. In summary, we confirmed that the newly emerged shadows could be removed more effectively and the overall image quality was better under the LNSE condition, while the spatial resolution of the projected result was better under the SR condition.

Figure 11(a) shows the weights $w_p(l)$ (see Eqs. 3 and 4) for each projector in each frame. The projected images based on the weights are also shown in Fig. 11(b). We confirmed that the weights were dynamically changed according to the movement of the occluder. In addition, we confirmed that the weights were evenly distributed among the projectors in the first frame because there was no occluder before the first frame. Therefore, in the first frame, the same number of shadows emerged in the projected results under both conditions, as shown in Fig. 9. The running time was 8.1 seconds for each frame in the experiment.

5 Discussion

We compared our method with the most relevantly related work (Audet and Cooperstock, 2007). The authors of the work developed their 3D tracking model by assuming that the occluder would stand vertically on the floor, which is usually the case with people. They also applied prediction of the occluder ' s movement based on a Kalman filter. Therefore, in the context of virtual rear projection screen for presentations, their method may work better than our method. On the other hand, their model had a special geometrical process that only worked for the assumed occluders, i.e., people. Because our method does not make any assumption about the occluder, it should work better in other contexts than this related work.

Nevertheless, there is a problem with measuring the 3D shape of a moving occluder at the current frame, which is then used to update the weights of the projectors for the next frame. There will always be a latency between the actual new position of the occluder and the estimated positions. In case the occluder moves too fast, the shadow cannot be perfectly removed, particularly on the shadow edge, even under the LNSE condition. To reduce the errors on the shadow edge, shadow prediction can be used. We will apply this method to improve our system in the future.

The running time of the experiment showed that the proposed approach did not run in real-time (30 frames per second). However, the current implementation was not well optimized (i.e., all processes ran on a CPU). The most time consuming part was the 2D image warping process in synthetic aperture capturing. We can make this process significantly faster using GPU parallel architecture. Another improvement can be done by applying GPU optimization, such as in the work of (Ladikos et al, 2008), to our visual hull reconstruction. In the future, we will improve the implementation to realize real-time processing.

6 Conclusion

In this study, we proposed a shadow removal technique using a multiple overlapping projection system. In particular, we focused on the case where cameras cannot be placed between an occluder and a projection surface. We applied a synthetic aperture capturing technique to estimate the appearance of the projection surface, and a visual hull reconstruction technique to measure the

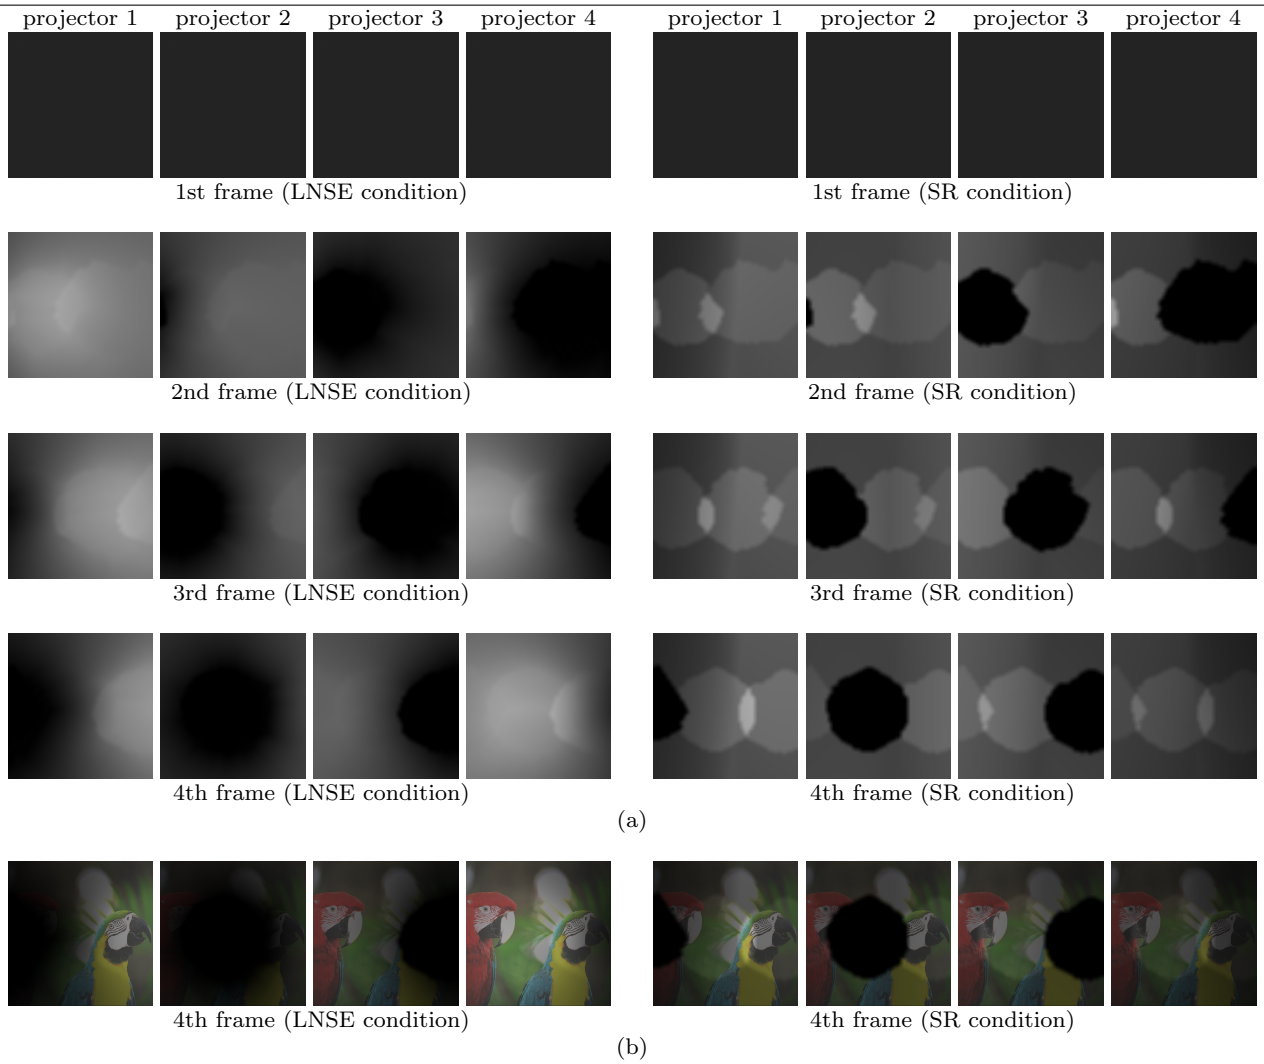

Fig. 11 (a) Weights $w_p(l)$ for each projector in each frame. The four weight maps in each frame correspond to the four projectors of the system. Higher intensity indicates a higher weight. (b) Projection images in the fourth frame.

occluder shape. Once the shape was acquired, shadow regions on the surface could be estimated. Our proposed shadow removal technique allows users to balance between the following two criteria: the likelihood of new shadow emergence and the spatial resolution of projected results. Through a real projection experiment, we confirmed that the technique that considers the likelihood of new shadow emergence (LNSE) could remove the shadows well, while that considering the other criterion (SR) could display clear images on the surface.

References

Audet S, Cooperstock JR (2007) Shadow removal in front projection environments using object tracking. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, pp 1–8

Bandyopadhyay D, Raskar R, Fuchs H (2001) Dynamic shader lamps: Painting on movable objects. In: Proc. IEEE/ACM International Symposium on Augmented Reality, pp 207–216

Bimber O, Emmerling A (2006) Multi-focal projection: A multi-projector technique for increasing focal depth. IEEE Transactions on Visualization and Computer Graphics 12(4):658–667

Bimber O, Wetzstein G, Emmerling A, Nitschke C (2005) Enabling view-dependent stereoscopic projection in real environments. In: Proc. IEEE/ACM International Symposium on Mixed and Augmented Reality, pp 14–23

Bimber O, Iwai D, Wetzstein G, Grundhöfer A (2008) The visual computing of projector-camera systems. Computer Graphics Forum 27(8):2219–2254

Cham TJ, Rehg JM, Sukthankar R, Sukthankar G (2003) Shadow elimination and occluder light sup-

pression for multi-projector displays. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, vol II, pp 513–520

Jaynes C, Webb S, Steele RM, Brown M, Seales WB (2001) Dynamic shadow removal from front projection displays. In: Proc. IEEE Visualization, pp 175–182

Jaynes C, Webb S, Steele RM (2004) Camera-based detection and removal of shadows from interactive multiprojector displays. IEEE Transactions on Visualization and Computer Graphics 10(3):290–301

Jones BR, Sodhi R, Campbell RH, Garnett G, Bailey BP (2010) Build your world and play in it: Interacting with surface particles on complex objects. In: Proc. IEEE International Symposium on Mixed and Augmented Reality, pp 165–174

Ladikos A, Benhimane S, Navab N (2008) Efficient Visual Hull Computation for Real-Time 3D Reconstruction using CUDA. In: Proc. IEEE Workshop on Visual Computer Vision on GPU

Laurentini A (1994) The visual hull concept for silhouette-based image understanding. IEEE Transactions on Pattern Analysis and Machine Intelligence 16(2):150–162

Low KL, Welch G, Lastra A, Fuchs H (2001) Life-sized projector-based dioramas. In: Proc. ACM symposium on Virtual Reality Software and Technology, pp 93–101

Nagase M, Iwai D, Sato K (2011) Dynamic defocus and occlusion compensation of projected imagery by model-based optimal projector selection in multi-projection environment. Virtual Reality 15(2):119–132

Raskar R, Welch G, Low KL, Bandyopadhyay D (2001) Shader lamps: Animating real objects with image-based illumination. In: Proc. Eurographics Workshop on Rendering, pp 89–102

Sato K, Inokuchi S (1987) Range-Imaging System Utilizing Nematic Liquid Crystal Mask. In: Proc. IEEE International Conference on Computer Vision, pp 657–661

Sugaya Y, Miyagawa I, Koike H (2010) Contrasting shadow for occluder light suppression from one-shot image. In: Proc. International Workshop on Projector-Camera Systems, pp 13–18

Sukthankar R, Cham TJ, Sukthankar G (2001) Dynamic shadow elimination for multi-projector displays. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, vol II, pp 151–157

Summet J, Flagg M, Cham TJ, Rehg JM, Sukthankar R (2007) Shadow elimination and blinding light suppression for interactive projected displays. IEEE Transactions on Visualization and Computer Graphics 13(3):508–517

Vaish V, Levoy M, Szeliski R, Zitnick CL, Kang SB (2006) Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition, vol II, pp 2331–2338

Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Perceptual image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing 13(4):600–612

YI AA, HM K (1971) Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. In: Proc. the Symposium on Close-Range Photogrammetry (American Society of Photogrammetry), pp 1–18

Yoshida T, Horii C, Sato K (2003) A virtual color reconstruction system for real heritage with light projection. In: Proc. International Conference on Virtual Systems and Multimedia, pp 161–168