Robust Blending and Occlusion Compensation in Dynamic Multi-Projection Mapping

Vanessa Lange, Christian Siegl, Matteo Colaianni, Marc Stamminger and Frank Bauer
Computer Graphics Group, University of Erlangen-Nuremberg

Abstract

Using multi-projection systems allows us to immerse users in an altered reality without the need to wear additional head-gear. The immersion of such systems relies on the quality of the calibration which in general will degenerate over time when used outside of a lab environment. This work introduces a novel balance term that allows us to hide high frequency brightness seams caused by self-shadowing of the projected geometry and the borders of the projection frustum. We further use this more robust blending between projectors to compensate for occluding spectators, who enter the projection volume, by filling the resulting shadows with light from other projectors.

Categories and Subject Descriptors (according to ACM CCS): H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities

1. Introduction

Merging the real world with virtual data is an emerging field in computer graphics. Recent developments have an impact on a wide variety of applications like training, art installations, product design or conferencing. Most available systems currently rely on virtual or augmented reality forcing participants to wear additional head-gear. In our research we focus on the advantages of multi-projection setups like Siegl et al. [SCT*15]. With such a system we augment and alter the physical appearance of objects in the real world. It is based on two important assumptions:

- The geometry of the real-world object is accurately captured.
- The calibration of the setup is almost perfect in every regard.

Generating precise models is easily accomplished. However it is not trivial to create and especially maintain a correct calibration for a fully dynamic multi-projection system. Both the self heating of the projectors and constant visitor traffic will degrade the calibration quality over time. Small unnoticed alterations in the projectors' position or orientation (less than the thickness of a sheet of paper) will have a visible impact on the projection. This introduces immersion breaking *brightness seams* (see Figure 1).
Another issue in a real world setup are the participants themselves. Our setup allows visitors to interact closely with the projections, for example by painting on the object [LSC*]. This however often turns them into occluders.

In this work, we introduce a new balance term for the original system, addressing both issues. Our new approach guarantees a smooth transition between projector contributions, hiding small, unavoidable calibration errors. It also enables us to use a low resolution depth camera to track occluders in the projection volume.

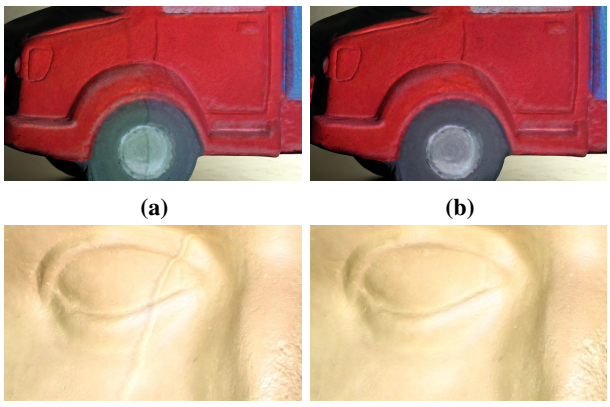

Figure 1: *In (a) and (c),* brightness seams *are visible. Our new blending hides these artifacts in (b) and (d).*

With this data, we alter the projection such that occluded light is compensated by other projectors in our setup (see Figure 7). Since our setup is dynamic, we cannot precompute critical areas and the required blending. Therefore, the projector contributions need to be determined in real-time.

2. Previous Work

Blending overlapping projectors is covered in a wide body of research. Brown et al. discuss several techniques for planar-like surfaces [BMY05]. The project *Shader Lamps* by Raskar et al. [RWLB01] showed that it is possible to blend multiple pro-

DOI: 10.2312/egsh.20171000

Figure 2: *The setup for our multi-projector system. Two projectors and a depth camera are directed at the Augustus bust and a truck.*

jectors on arbitrarily shaped objects. In contrast to our work, their multi-projection relies on a static setup. Lee et al. [LDMA*04] developed a technique for automatic projector calibration. However, light sensors need to be embedded in the target surface.The first system able to blend multiple projectors in real-time for dynamic setups and on arbitrary, markerless geometry was introduced by Siegl et al. [SCT*15].

Erasing shadows from a projection thus far is limited to planar or planar-like surfaces. In [JWS*01] and [JWS04], Jaynes et al. use cameras to compare the actual projection to the expected result. While this approach gives impressive results, it is not suited for complex surfaces. A similar concept is employed by Sukthankar et al. [SCS01]. They compare reference images with a live view of the projection to find and compensate shadows.

A different approach to shadow removal is detecting the occluding object itself. For planar surfaces, Tan et al. [TP02] use infrared background illumination. Audet et al. [AC07] employ computer vision methods to robustly track occluding objects. However, they are restricted to vertical occluders.

While those methods work well in their specific settings, we propose a simple, more general solution based on a low resolution depth-sensor.

3. System Overview

As a basis for this work, we use the system introduced by Siegl et al. The original system optimizes the luminance of every projector ray such that multiple projectors are correctly blended within overlapping regions. This results in a uniform illumination of the target object. Each ray's luminance contribution is determined by minimizing the following energy functions:

- A physical term, representing real-world attenuation of light (incident angle and distance)
- A balance term, giving advantage to the projector ray that will likely give the best projection quality
- A *Laplacian* regularization term
- A bounding term, limiting the projector luminance to the possible 0 to 1 range

To reach real-time rates, the resulting non-linear optimization problem is solved efficiently on the GPU. For more details we refer the reader to [SCT*15].

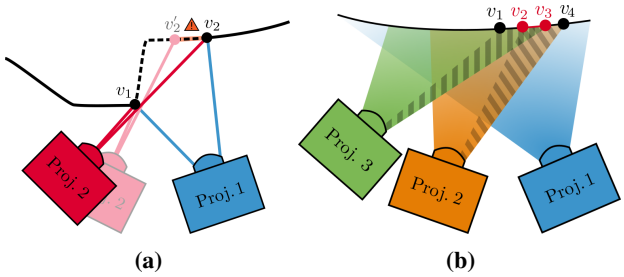

Figure 3: *In (a), the red projector cannot project onto the area between v_1 and v_2 (dashed line). The assumed, incorrect projector position (lighter red) will cause a* brightness seam *between v_2' and v_2. In (b), the green and orange projector fade out towards their borders (hatched areas), which overlap between v_2 and v_3.*

4. Blend Weights

The balance term implemented by Siegl et al. sets the incident angle of a ray in ratio to the incident angles of every other ray hitting the same surface point. This results in a linear combination of pixel brightnesses prone to high frequency discontinuities where the set of contributing projectors changes. See Figures 1a and 1c for the resulting artifacts caused by a slight miscalibration. The influence of the *Laplacian* regularizer, which is intended to counteract this behavior, is limited due to the small filter size. However, we will continue to use the regularization to improve numerical stability.

In Figure 3a, two projectors illuminate the surface. The dashed line between v_1 and v_2 describes a part of the surface that is occluded for the red projector in the real world. Due to a slight miscalibration our system assumes $[v_1, v_2']$ to be the occluded area instead. Since both projectors are expected to illuminate the area $[v_2', v_2]$, the blue projector's brightness output is reduced here. However, this area is in fact not illuminated by the red projector, resulting in a dark artifact (*brightness seam*). The human eye is very sensitive to such high frequency brightness discontinuities which breaks the immersion. Every surface point (*transition point*) with a change in the set of incident projectors (v_1 and v_2 in the example) is an artifact-prone transition.

We present an improved geometry aware balance term that incorporates a smooth, C^0 continuous blending around these *transition points*. The blending transforms the critical areas into low frequency seams which are imperceivable to the spectator.

The first step in our algorithm is to identify the *transition points*. Therefore, we create a depth map from the viewpoint of every projector (see Figure 4a) and mark depth discontinuities (see Figure 4b). The discontinuities are dilated (Figure 4c) and blurred (Figure 4d) to soften the critical areas. The resulting *blend weight map* in Figure 4d shows a smooth transition between 0 (black) and 1 (white). The closer to 0 a pixel's blend weight w, the less brightness contribution (luminance) we wish to assign.

The balancing linear combination by Siegl et al. does not support an unilateral dimming of a single projector. Thus a new balance term E_{bal} is introduced. We maintain the idea that the light contribution p_i of pixel i should be equivalent to its expected projection quality. The quality is measured by $s_i = \langle \vec{n}_i, \vec{i}_i \rangle$, where $\vec{n}_i$ is

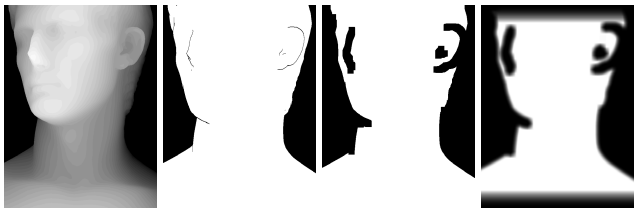

(a) (b) (c) (d)

Figure 4: *Based on the expected depths in (a), depth discontinuities and the object silhouette are detected in (b). These regions are dilated (c) and then blurred afterwards (d). An additional gradient blends at the borders.*

the normalized surface normal and $\vec{i_i}$ the normalized negative light direction. For every ray this quality measure is set in relation to the rays it is interacting with (if the set of reprojections R_i exists):

$$\frac{p_i}{\sum_{r\in R_i} p_r} \stackrel{!}{=} \frac{s_i}{\sum_{r\in R_i} s_r}$$

By multiplying the blend weight w_j to each quality s_j direct influence on the expected ray quality is taken. As a result, pixels prone to causing *brightness seams* are assigned less contribution and a smooth blending between projectors is achieved.

Consider a pixel i with high projection quality s_i close to the depth discontinuity at the nose (see Figure 4d). In the original system this results in a high brightness contribution. However, due to its corresponding small blend weight, the product $s_i w_i$ is small and likewise the assigned luminance p_i.

For all pixels, the balance error is described as follows:

$$E_{\text{bal}} = \sum_{i=1}^{N} \left(p_i \cdot \sum_{r\in R_i} s_r w_r - s_i w_i \cdot \sum_{r\in R_i} p_r \right) \stackrel{!}{=} 0$$

Figure 5 shows the new term's brightness balancing for the setup depicted in Figure 3b where two critical areas (hatched) overlap. In this example the green projector transitions in $[v_1, v_3]$ and the orange projector in $[v_2, v_4]$. Without loss of generality we consider a constant quality distribution s_i.

While the old balance term enforces a combined brightness of 1, abrupt changes occur at the transition points v_3 and v_4 (dashed lines). In contrast, our new balance term ensures C^0 continuity while still maintaining a combined brightness of 1 (solid lines).

In the following we demonstrate the continuity at the *transition point* v_3, which is illuminated by a blue (b), orange (o) and green (g) projector pixel. The point v_3 is an element of $[v_2, v_3]$ and $[v_3, v_4]$. When approaching v_3 from the left, $e_{\text{bal},\,g} = s_o w_o p_g + s_b w_b p_g - s_g w_g p_o - s_g w_g p_b$. Since w_g is 0 (due to the *blend weight* at the projection border), the remaining error term $e_{\text{bal},\,g} = (s_b w_b + s_o w_o) p_g$ is minimized by $p_g = 0$. From the right, the green projector has no influence and the luminance p_g is therefore implicitly 0. Hence the green projector is C^0 continuous at v_3.
For the orange projector, the balance term at v_3 in $[v_2, v_3]$ is $e_{\text{bal},\,o} = s_b w_b p_o + s_g w_g p_o - s_o w_o p_b - s_o w_o p_g$. As shown above, $p_g = 0$ and $w_g = 0$ and therefore $e_{\text{bal},\,o} = s_b w_b p_o - s_o w_o p_b$. From the right, without the green projector the balance term is also $e_{\text{bal},\,o} = s_b w_b p_o - s_o w_o p_b$, making the orange projector C^0 continuous as well. The same holds for the blue projector.

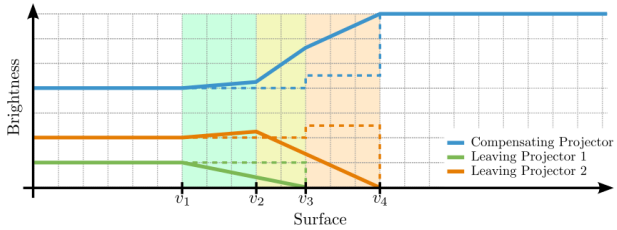

Figure 5: *Brightness distribution of three projectors with our (solid) and the original (dashed) balance term around overlapping critical areas (see Figure 3b).*

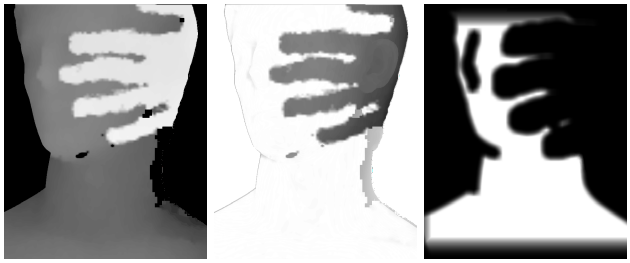

(a) (b) (c)

Figure 6: *By comparing the converted depth image from the depth camera (a) with the rendered depth (see Figure 4a), the occluding hand is extracted (b). The affected pixels are then included in the blend weights map creation process (c).*

Dynamic Shadows Our second objective in this work is removing dynamic shadows. These are usually caused by the users themselves (e.g. occluding hand) and can appear anywhere on the target object. Inherently, eliminating dynamic shadows is only possible if there is another projector which can illuminate the shadowed region with enough power. Using the depth camera of our setup, we first detect the occluder (see Figure 6a) in real-time. For every projector, the rendered depth map is compared to the recorded depth map that is transformed for each projector's point of view. Any measured depth value smaller than the expected value belongs to an occluder (see Figure 6b) and is easily segmented. We amend our blend weight map with this information to integrate shadow elimination as follows: The occluded pixels are assigned a blend weight of 0 (before filtering the blend weight map, see Figure 6c). The affected projector therefore fades out towards the occluder and the shadows are filled with light from the other (unaffected) projectors.

5. Results and Discussion

Using our novel blending approach we remedy brightness artifacts as can be seen in Figure 1b and 1d. The new balance term also enables us to eliminate dynamic shadows cast by the user, e.g. an occluding hand (see Figure 7). Our additions are easily integrated into an existing system without adding a noticeable impact on the real-time performance.

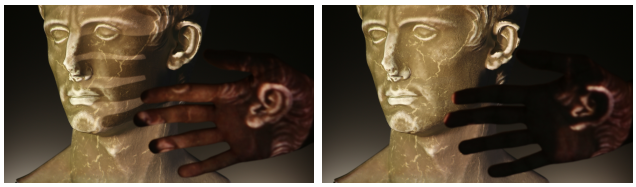

(a) (b)

Figure 7: *The hand shadow visible in (a) is eliminated in (b). Some light remains on the occluder if no other projector can compensate.*

(a) (b) (c)

Figure 8: *When projecting sharp edges like a checkerboard pattern (a), the actual misalignment of the projectors becomes apparent (b). If the loss of a projector cannot be compensated the shadow remains visible (c).*

Hardware and Performance Our multi-projector system (see Figure 2) consists of two `NEC NP-P451WG` projectors (1280x800) and an `ASUS Xtion PRO Live` depth camera for object tracking. These components are connected to a standard desktop workstation with an `Intel Core i7 4771 (3.5GHz)`, 32GB of RAM and an `NVidia GeForce GTX 980` graphics card.

A detailed overview over the performance of various parts of the base system is given in [SCSB16]. All operations added to the pipeline in this work operate on projector-pixel level and therefore are independent of the target mesh size. For the Augustus bust, generating the map and blending adds 1.44 ms, shadow elimination 0.98 ms to the pipeline. The truck model takes 1.15 ms for blending and also 0.98 ms for shadow elimination. Faster blending for the truck is due to its smaller size in pixel space. The performance of the core solver part is not affected by the new balance term. With overall runtimes of well below 30 ms, the system remains real-time. Note, that any visible latency arises is due to the used consumer-grade hardware. For a discussion see [SCT*15].

Limitations As we have demonstrated, our system compensates immersion breaking *brightness seams* introduced by calibration errors. However, when combined with a high frequency projection, e.g. a checkerboard pattern (see Figure 8a and 8b), *ghosting* artifacts become visible on the target object. This is a result of blending the misaligned textures between projectors. To reduce the effect, the balancing could be adjusted to prefer one projector alone wherever possible. For reasonably small calibration errors (as we would expect from daily usage) such errors mostly result in a barely noticeable softness in the critical areas.

The shadow elimination is very robust (please refer to the accompanying video). However, there are still artifacts in some areas as the shadow is only compensated in regions where the remaining projectors contribute enough light (see Figure 8c where only the shadow of the fingertips is eliminated).

6. Conclusion

In this work we have presented a remedy for two common problems in multi-projection mapping. First, we introduced a new blending approach that efficiently solves the problem of *brightness seams*, caused by even small errors in the calibration. A second common problem in real-world setups are occluders casting shadows on the target object. Using our new smooth balance term, we were able to solve this efficiently. With these two important additions to an existing multi-projection mapping system we substantially improve the perceived quality and prevent artifacts from breaking the user's immersion. The real-world applicability is enhanced due to a more robust system that better handles the deterioration of the calibration over time and an increased tolerance for dynamic occluders.

References

[AC07] AUDET S., COOPERSTOCK J. R.: Shadow removal in front projection environments using object tracking. In *2007 IEEE Conference on Computer Vision and Pattern Recognition* (June 2007), pp. 1–8. 2

[BMY05] BROWN M., MAJUMDER A., YANG R.: Camera-based calibration techniques for seamless multiprojector displays. *Visualization and Computer Graphics, IEEE Transactions on 11*, 2 (2005). 1

[JWS*01] JAYNES C., WEBB S., STEELE R. M., BROWN M., SEALES W. B.: Dynamic shadow removal from front projection displays. In *Proceedings of the Conference on Visualization '01* (Washington, DC, USA, 2001), VIS '01, IEEE Computer Society, pp. 175–182. 2

[JWS04] JAYNES C., WEBB S., STEELE R. M.: Camera-based detection and removal of shadows from interactive multiprojector displays. *IEEE Transactions on Visualization and Computer Graphics 10*, 3 (2004). 2

[LDMA*04] LEE J. C., DIETZ P. H., MAYNES-AMINZADE D., RASKAR R., HUDSON S. E.: Automatic projector calibration with embedded light sensors. In *Proceedings of the 17th annual ACM symposium on User interface software and technology* (2004), ACM. 2

[LSC*] LANGE V., SIEGL C., COLAIANNI M., KURTH P., STAMMINGER M., BAUER F.: Interactive painting and lighting in dynamic multi-projection mapping. In *Augmented Reality, Virtual Reality, and Computer Graphics. AVR 2016*, Springer, Cham, pp. 113–125. 1

[RWLB01] RASKAR R., WELCH G., LOW K.-L., BANDYOPADHYAY D.: *Shader Lamps: Animating Real Objects With Image-Based Illumination*. Springer Vienna, Vienna, 2001, pp. 89–102. 1

[SCS01] SUKTHANKAR R., CHAM T.-J., SUKTHANKAR G.: Dynamic shadow elimination for multi-projector displays. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on* (2001), vol. 2. 2

[SCSB16] SIEGL C., COLAIANNI M., STAMMINGER M., BAUER F.: Stray-light compensation in dynamic projection mapping. In *SIGGRAPH ASIA 2016 Technical Briefs* (New York, NY, USA, 2016), SA '16, ACM, pp. 21:1–21:4. 4

[SCT*15] SIEGL C., COLAIANNI M., THIES L., THIES J., ZOLLHÖFER M., IZADI S., STAMMINGER M., BAUER F.: Real-time pixel luminance optimization for dynamic multi-projection mapping. *ACM Trans. Graph. 34*, 6 (Oct. 2015), 237:1–237:11. 1, 2, 4

[TP02] TAN D. S., PAUSCH R.: Pre-emptive shadows: Eliminating the blinding light from projectors. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2002), CHI EA '02, ACM, pp. 682–683. 2