2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)

Real-time shader-based shadow and occlusion rendering in AR

Agapi Chrysanthakopoulou Konstantinos Moustakas[*†]

Department of Electrical and Computer Engineering, University of Patras

Abstract

We present novel methods designed to elevate the realism of augmented reality (AR) applications focusing specifically on optical see-through devices. Our work integrates shadow rendering methods for multiple light sources and dynamic occlusion culling techniques. By creating custom surface shaders we can manage multiple light sources in real-time, augmenting depth perception and spatial coherence. Furthermore, the dynamic occlusion culling system handles occluded objects, ensuring a more convincing and seamless user experience. Several cases and methods are presented with their results for various lighting and spatial conditions, promising a more enhanced and immersive user experience in various AR domains.

Index Terms: Augmented Reality—Mixed Reality—Shadow Rendering—Dynamic Occlusion

1 Introduction

In augmented reality (AR), shadows and occlusion play pivotal roles in enhancing the realism of the digital elements within the physical world. By accurately casting shadows from virtual objects onto the physical environment, the illusion of depth and spatial relationships is significantly enhanced. Realistic shadow rendering helps ground virtual elements in the environment, making them feel more integrated and natural. Shadows provide important visual cues for users to perceive the positioning and interaction of virtual objects within their surroundings, contributing immensely to the overall realism and depth perception of the AR experience [1].

Occlusion, the technique of rendering virtual objects to appear behind real-world obstacles, creates a seamless integration between the virtual and real. It enables digital elements to behave realistically, hiding behind physical objects as they would in the natural environment. This aspect of AR enriches user immersion by fostering a sense of depth and tangibility, allowing users to interact more naturally with the augmented content [2]. Without occlusion, virtual objects may appear disjointed or disconnected from the surroundings, diminishing the overall sense of believability in the AR experience.

In this work, we present the combination and extension of known techniques for shadow rendering and dynamic occlusion culling for optical see-through (OST) augmented reality (AR) devices, for a more intuitive and immersive AR experience. We take advantage of the light estimation tools provided by the devices for the calculation of the current environment's main light direction and ambient intensity. Our approach adds the possibility of adding multiple light sources and material-dependent occlusion in a space. With the creation of custom shaders that utilize the "negative" shadow effect [3] and support multiple shadows and shadow-hue change along with the proper shaders for occlusion, the results give a more natural combination of the virtual and the real world.

[*]e-mail: moustakas@upatras.gr
[†]URL: www.vvr.ece.upatras.gr

2 Method Overview

2.1 Shadow Drawing

Notably, when using the HoloLens 2 (HL2), visualizing object shadows poses challenges due to the alpha channel behavior in the lights combiners technology. The device renders black as the alpha channel, causing darker colors to become less visible to the user. This limitation can lead to a paradox where completely black shadows or objects become entirely invisible to the observer. In this work, colored shadows have been also considered since they are reported to enhance the shadow experience in the presence of the aforementioned paradox [4].

So for the visualization of the shadow, we take advantage of an optic effect called "negative shadow". The area around the shadow is slightly lit so that the dark area that is the shadow of the object, which would normally appear transparent is distinguished from the rest (Fig. 1, Right). This material can be applied to surfaces created with the spatial mapping mechanism or any other way. The rest of each surface must be a true black (unlit), but using a directional light for shadows will result in everything appearing lighter, so the use of spotlights was necessary for multiple light sources.

Figure 1: **Left:** A transparent shader that receives multiple shadows and has shadow hue change. **Right:** A solid color unlit shader that receives multiple shadows

2.2 Occlusion Culling

One of the primary techniques used in HL2 for occlusion is spatial mapping with a black unlit shader. This technique allows for the creation of a 3D map of the user's environment with the proper material. Spatial surfaces play a vital role in this process by hiding holograms, which makes them appear more realistic and immersive within the physical space. By hiding holograms behind real-world surfaces, visual clutter is minimized, providing a cleaner interface. The creation of the spatial map also provides the ability of the real world to cast shadows on the virtual objects.

However, it's important to note that this tool has certain distance limitations that need to be taken into consideration. For large, static areas, such as warehouses, a scanned model of the space can be used to match with the real world, also allowing the marking of see-through areas (e.g., windows (Fig. 2, Left)). For closer distances (approximately $< 0.8m$), other techniques to consider include an additional depth camera for moving objects or simply utilizing the hand tracking with a proper material, which works well in most AR cases (Fig. 2, Right).

3 Implementation

As for now, the existing unlit and transparent shaders do not support shadows from multiple light sources or shadow gradients that are

979-8-3503-7449-0/24/$31.00 ©2024 IEEE
DOI 10.1109/VRW62533.2024.00281

Figure 2: **Left:** A virtual object behind the glass. **Right:** A hand model with a black color unlit material tracking the real hand.

vital for achieving the desired effects that will allow the user to perceive shadows and occlusion in AR.

The proposed method was implemented across two distinct platforms: an OST AR head-mounted display (HMD)(HL2), and a smart device utilizing the MRTK toolkit and AR Foundation toolkit respectively within the Unity3D engine. As the user moves through space, the Hl2 camera captures images to create a Cubemap which is then assigned to a Reflection Probe to calculate the ambient and directional lighting from the reflection calculations. For the light estimation, we also tested a sun-direction calculation based on the user's location (Fig. 3, Left). For outdoor applications, this can be combined with information about weather conditions for the light intensity.

To achieve precise and realistic shadow mapping of multiple light sources, we tested two distinct scenarios:

1. A directional virtual light in Unity emulates the real-world primary light source, including direction and intensity, as identified from the devices, while virtual point lights represent any supplementary light sources.
2. All lights, including the primary source, were modeled using point lights, positioned and oriented in alignment with the real-world light sources for enhanced accuracy and realism (Fig. 3, Right).

In each case the spotlight's position and direction are calculated from the following:

$$\begin{aligned} SpotlightPos &= ObjectPos - LightDir * distance + offset, \\ SpotlightRot &= LightRot, \end{aligned} \tag{1}$$

where offset and distance depend on the size of the virtual object and the range of the spotlight respectively. Currently the additional lights are placed manually and are not calculated from any algorithm.

Figure 3: **Left:** Light estimation based on user's location, captured with smart device. **Right:** Virtual and real object comparison with two light sources, captured with HL2.

For HL2 two shaders gave the best results depending on the case and both shaders took under consideration the desired "negative" shadow effect. The first is an unlit shader that receives shadows and creates a gradient lighter outline around the shadow (Fig. 4, Left). The second is a solid color shader that can receive shadows from multiple light sources by adding multiple spotlights. The hue of the shadow can be altered in both cases (Fig. 4, Right).

For the effect to work properly, the solid color must be a true black (unlit) in each case. The reason why multiple spotlights are preferred over directional lights is that each directional light brightens the entire virtual scene, making it difficult to distinguish any shadows.

In order to use the same technique with a smart device a different custom shader must be created. By default Unity's transparent shader cannot receive shadows, so the proper alterations give the potential to receive multiple shadows and change the shadow color if one needs to use the complementary color generation algorithm. The application on a smart device was complementary work for comparison and exploring different possibilities.

Finally, the virtual objects have the proper materials to be able to receive shadows as well as cast them. In the present work we only used the spatial mapping produced by the HL2 and the results are already very promising as we can see the changes of the shadow along the geometry of the object (Fig. 4, Right).

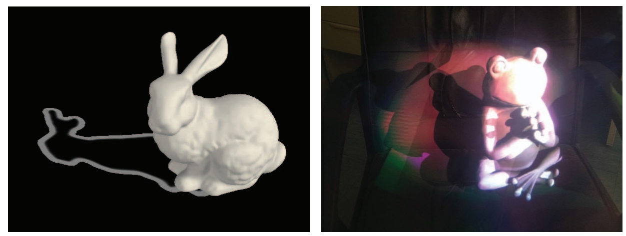

Figure 4: **Left:** A shader with shadow outline for "negative" shadow. **Right:** A virtual object casting shadows from two light sources on a real chair and receiving shadow from the handle.

4 Conclusion

In summary, our work combines and extends the possibilities of known techniques for shadow rendering and occlusion culling in AR applications, significantly enhancing visual fidelity and immersion. Our methods highlight the importance of having the virtual world seemingly merging with the real one in various lighting and spatial occasions. Unlit and transparent shaders by default cannot receive shadows so we created custom ones that include new capabilities, such as integrating multiple light sources and managing occlusion based either on spatial mapping or detailed models of a space that take into consideration the transparency of real objects. The potential of these advancements underscores the evolving landscape of augmented reality, paving the way for more convincing and integrated virtual experiences.

Acknowledgments

This work was supported by the European Union's Horizon Europe Research and Innovation Program under Grant 101092875—DIDYMOS-XR: Digital DynaMic and respOnsible twinS for XR.

References

[1] Haley Adams, Jeanine Stefanucci, Sarah Creem-Regehr, Grant Pointon, William Thompson, and Bobby Bodenheimer. Shedding light on cast shadows: An investigation of perceived ground contact in ar and vr. *IEEE transactions on visualization and computer graphics*, 28(12):4624–4639, 2021.

[2] Haley Adams, Jeanine Stefanucci, Sarah Creem-Regehr, and Bobby Bodenheimer. Depth perception in augmented reality: The effects of display, shadow, and position. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 792–801. IEEE, 2022.

[3] Francesco Osti, Gian Maria Santi, and Gianni Caligiana. Real time shadow mapping for augmented reality photorealistic rendering. *Applied Sciences*, 9(11):2225, 2019.

[4] Chun Wei Oio and John Dingliana. Perceptually enhanced shadows for ost ar. In *Proceedings of the 14th International Workshop on Immersive Mixed and Virtual Environment Systems*, pages 30–36, 2022.