Screen Space Ambient Occlusion for Virtual and Mixed Reality Factory Planning

Fabian Scheer
Daimler AG
Wilhelm-Runge-Str.11
89081 Ulm/Germany
fabian.scheer@daimler.com

Michael Keutel
Otto-von-Guericke-Universität Magdeburg
Hedwigstraße 5
04315 Leipzig/Germany
michael.keutel@googlemail.com

ABSTRACT

Ambient occlusion represents one aspect of global illumination, simulating the actual accessibility of surfaces for indirect illumination due to occlusion of nearby geometry. Recently the approximation in screen space made it feasible for realtime applications. We focus on the utilization of screen space ambient occlusion (SSAO) in virtual and mixed reality environments, especially considering scenes with a large spatial extent and massive amounts of data, typical for industrial factory planning scenarios. Therefore, a sophisticated screen space approach closer to the original definition of ambient occlusion is presented and compared with existing techniques considering the visual quality. Furthermore, we introduce a method to avoid the disappearance of the SSAO effect in depth for scenes with a large spatial extent. A user study compares the impact of our approach to standard SSAO and standard phong shading, regarding the several cues of human depth perception in a virtual and a mixed reality scenario. The evaluation underlines the benefit of our approach for pure virtual scenes and additionally illustrates a different perception of the cues regarding virtual scenes with a half transparent overlay-image of the real world, typical for mixed reality discrepancy check scenarios.

Keywords

Ambient Occlusion, Screen Space, global illumination, discrepancy check, mixed reality, factory planning.

1. INTRODUCTION

The approximation of global illumination (GI) effects in realtime has always been a major issue of computer graphics research. One technique simulating the occlusion term of indirect diffuse light is known as Ambient Occlusion (AO). Surface points being occluded by nearby geometry are shadowed, thus reflecting less light to the viewer, which can be observed in cavities or creases. Although AO is an indirect GI effect, the human visual system gains essential clues of depth, curvature and spatial arrangement from the resulting soft shadows [Lan99a]. Taking advantage of the evolved capabilities of modern graphics hardware [Mit07a] recently presented SSAO as an approximation of AO in image space. In contrast to former computationally expensive AO approaches, SSAO uses the Z-Buffer as a representation of the scene geometry and therefore involves significant advantages:

- computation independent from polygon count
- feasible for fully dynamic content
- calculation can entirely be done on the GPU

Due to an improved realism, SSAO nowadays is widely spread. Being independent from scene geometry and supporting fully dynamic content makes SSAO attractive for the visualization of mass data scenarios with a large spatial extent. Discrepancy checks in mixed reality environments, comparing digitally planned data with real build structures are an actual topic of research [Sch08a]. Typical scenes like complete facilities contain massive amounts of data [Bru06a], thus realtime GI effects are hard to achieve and sparsely considered. Screen space approaches can close this gap by improving the visualization of purely virtual environments and the process of discrepancy checks in mixed environments may also have a benefit. In this paper we present an alternative approach for SSAO, producing visually convincing results, especially in scenes with a large spatial extent. We give a comparison to existing approaches regarding the visual quality and present an evaluation of SSAO in mass data scenarios, showing how SSAO affects the several criteria of human depth perception in merely virtual and in mixed environments. Finally, our evaluation results for the use of SSAO in

discrepancy checks of industrial factory planning applications are described.

2. RELATED WORK

The idea of computing AO approximately in screen space by using a linear Z-Buffer originally evolved from Crytek [Mit07a]. They compute the occlusion value for every pixel *p* in the following way: The respective view space position *P* is computed to sample the vicinity with a set of 3D offset-vectors. Every obtained sample point is projected back to image space and an occluder is found if the depth stored in the linear Z-Buffer is closer to the viewer than the depth of the sample point itself [Kaj09a]. The final occlusion term represents the ratio of solid geometry to empty space in the vicinity of *P*. However, their method suffers from uniform shadowing of actually unoccluded faces and produces edge highlighting artefacts.

[Fil08a] present a similar approach, but apply an attenuation function to weight an occluder's influence dependent on the difference between the depth of the sample point and the one read from the Z-Buffer. The final occlusion value is formed by averaging the contribution of all samples. However, they do not assure that the detected occluders are actually located within the visible hemisphere of the point of interest. As a result they have to deal with self-occlusion occurring on unoccluded surfaces.

[Bav08a] interprete the Z-Buffer as an heightfield and search for the horizon, which is characterized by the slope in the heightfield surrounding a pixel. This horizon-based AO reveals visually pleasing shadows at interactive framerates. They further suggest to speed up computation time by rendering in half resolution and compensate the resulting low quality with a blur pass. We will give a comparison to this approach in section 4.

[Fox08a] consider the orientation of surface normals in the vicinity of a pixel to detect occlusion caused by creases. This approach is fast, but ignores low-frequency occlusion from distant objects.

[Sha07a] introduce an approach to calculate high- and low-frequency parts separately. They assume small spherical occluders to compute the high-frequency occlusion by nearby surface points entirely in screen space. Approximating scene geometry with spherical proxies is used to determine the low-frequency occlusion by distant geometry. Due to the need of approximating the scene geometry for distant occlusion, low frame rates for highly tessellated massive data scenes (e.g. > 20 Mio.) can be expected.

[Rei09a] address the limitation of viewpoint dependence. They perform a realtime voxelization of the scene geometry and compute occlusion by tracing sample rays through the voxelgrid, determining the amount of solid cells being hit. Thus, occlusion caused by hidden surfaces can be considered.

3. VECTOR BASED APPROACH

Referring to the definition of AO, occlusion at a point of interest *P* is caused by any surface point within its visible hemisphere, as this represents all directions of incoming light. We derive our approach from this background and determine occluders, by checking their position to be within the hemisphere of *P*.

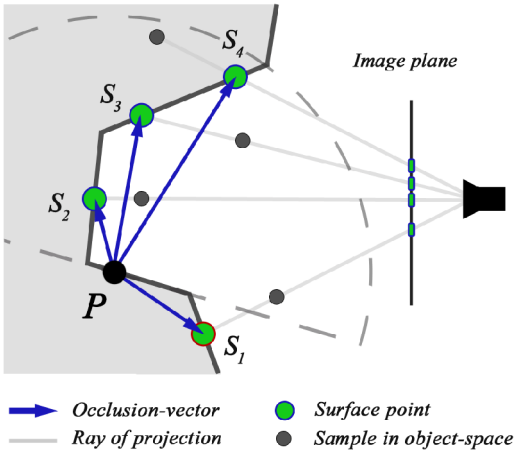

Figure 1. Concept of occlusion-vectors: Sample S_1 is located outside the visible hemisphere and does not contribute to occlusion of P.

Having a set of sampling positions on the image plane, we read the respective depth values from the linear z-buffer and further un-project the sample points into view-space. This gives us a set of surface points representing potential occluders. For any of them we finally create a so called occlusion-vector v_O, pointing from *P* to the respective occluder candidate. An occluder is found if the angle between the surface normal at *P* and v_O is less than 90°, checked by a simple dot-product (see figure 1).

The final occlusion value $o(v_O)$ is determined considering the length of v_O. This idea is straightforward, as long vectors representing distant objects occlude less than short vectors pointing to nearby objects. We apply a function similar to the one usually being used in computer graphics to model the attenuation of light:

$$o(v_O) = \frac{1}{1 + c \cdot \|v_O\|^2} \cdot \left\langle v_O^{norm}, n_P \right\rangle$$

The strength of the quadratic attenuation is controlled by an artistic parameter *c*. In addition a cosine weight expressed by the attached dot product is applied. This is based on [Lan02a] who applies Lamberts law on

the concept of ambient occlusion and concludes that the occlusion is maximal for surface points along the direction of the normal.

The utilization of occlusion-vectors is a simple but effective technique. We correctly ignore surface points not contributing to occlusion, represented by occlusion-vectors pointing outside the hemisphere. This enables us to detect fully unoccluded surface points and therefore avoids false shadowing effects on planar faces or the highlighting of edges, e.g. occurring in [Kaj09a]. For the same reason, our approach avoids the self-occlusion problem described by [Fil08a]. Finally, we acquire an enhanced visual quality of the AO with our approach. This is a result of a higher sample density effectively used for the occluder detection. This issue is clearly illustrated in figure 1, as we obtain two valid occluders for the sample points S_2 and S_3 in contrast to the approaches of [Kaj09a, Fil08a]. Furthermore, the length of the occlusion-vector is a more effective measure for the occlusion strength than the view-direction dependent depth deltas used by [Fil08a]. Hence, we are able to properly represent the distance from the point of interest to the respective occluder.

Sampling technique

Correctly sampling the vicinity of a surface point forms the basis for a well approximated occlusion value. Regarding the space wherein the samples are created, view and image space sampling can be distinguished. View space sampling uses a set of 3D offset-vectors to sample the environment of the point of interest [Mit07a][Fil08a], while image space sampling utilizes 2D vectors to directly create sample positions on the image plane [Bav08a]. The latter makes the back-projection of sample points redundant and offers direct access to the depth- and normal-buffer.

However, we obtain an enhanced visual quality with view space sampling due to a better distribution of the sample positions. According to [Fil08a], we flip an offset-vector in the sampling process if it points outside the visible hemisphere of *P*, obtaining a higher density of sample positions within the hemisphere. Although the direction of the acquired occlusion-vectors is not predictable due to the un-projection with an unknown depth value, the higher initial sampling density significantly enhances the probability to find potential occluders. Figure 2 compares the AO obtained by both sampling methods. The higher noise using image space sampling is clearly observable.

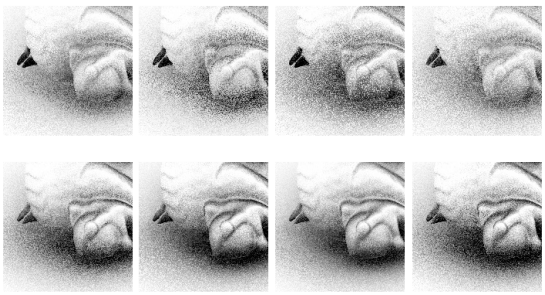

Figure 2. Comparison of image (top) and view space sampling (bottom). Four random sampling patterns applied to Stanford Dragon

Aliasing artefacts due to a constant sampling pattern, are avoided by reflecting the offset-vectors on a random plane as described by [Mit07a]. A Gaussian bilateral filter commonly used for SSAO is further applied to blur the noise, while preventing AO from bleeding through strong object edges [Pet04a].

Accounting for scene scale and extent

Virtual and mixed reality factory planning applications generally involve scenes of a large spatial extent, exhibiting geometry of different metric unit scales due to the individual creation process of the CAD authoring tools. These scenes reveal difficulties for a robust computation of occlusion. We therefore pay attention to the following issues:

1. Definition of the sampling radius
2. Normalization of the occlusion measure

In general the sampling pattern can be scaled by a radius multiplier s_{scene} to give an artistic control over the area being considered by AO. This sampling radius is defined to be constant in view space throughout all currently known approaches. Hence, the area of sampling decreases for surfaces being farther from the viewer, and the visible effect finally vanishes with distance. Therefore, we chose the sampling radius to linearly grow with increasing depth values, keeping its projection constant in screen space. Although this is physically incorrect, it reveals two advantages: First, the larger sampling area produces soft shadowing on distant objects as well and thus makes the AO effect observable within the entire scene (figure 5). An improved visual perception of depth has been verified in a user study presented in section 5. Second, setting the radius in relation to depth couples the radius to individual scene scales, making a further adjustment by the artist redundant. The relation between the initial radius r_B and the resulting sampling radius r_S is derived as follows:

$$r_S = r_B \cdot s_{scene} \cdot \frac{d}{z_P} \qquad d : \text{depth of image plane}$$

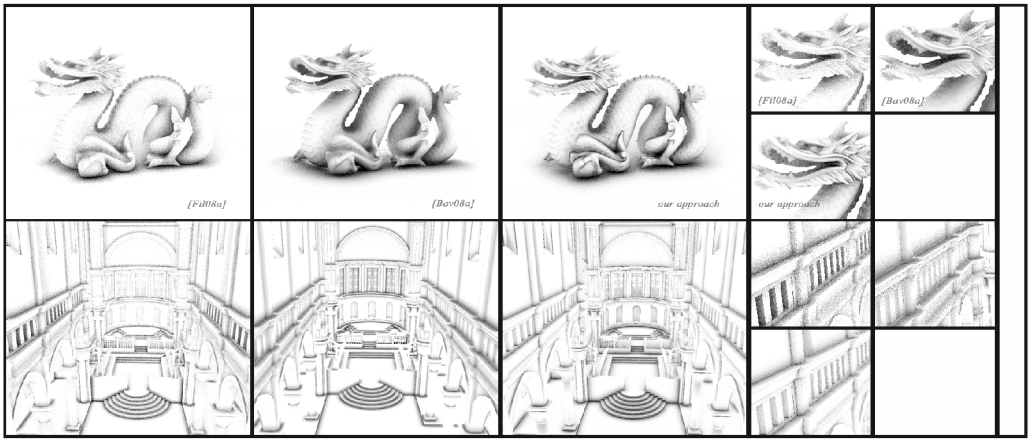

Figure 3. Comparison of our results with the approaches by [Fil08a] and [Bav08a]

In order to account for different scene scales, we normalize the measure of occlusion given by the length of the occlusion-vectors due to their dependence on the underlying scene unit. We hence avoid the individual adjustment of the artistic parameter *c* in extreme ranges of up to $[0,10^7]$. Yet the normalization mathematically requires the maximum length of the occlusion-vector to be known. As these vectors are formed by the un-projection with an afore unknown depth value, the maximum is not given in advance. We deal with this issue by taking the respective sampling radius in view space as an approximate measure and thus doing a quasi-normalization:

$$\|v_O\|' = \frac{\|v_O\|}{r_S}$$

This means the normalized length might exceed the range of [0,1]. Nevertheless, we observe robust results with this approach, being able to deal with different scene scales without the need of manual adjustment for every scene.

4. COMPARISON

In figure 3 we show the comparison of the SSAO approaches acquired with the Stanford Dragon (871K polygons) and the Sibenik Cathedral (80K polygons). 36 Samples were used at a resolution of 800x800 pixels. The unblurred effect of SSAO is shown to allow a proper evaluation. As SSAO generally is a strongly parameter-dependent effect, we attempted to appropriately set the parameter for every approach allowing a visible comparison. The analysis of the images reveals a high amount of noise for [Fil08a] in both scenarios. Surface details and contours are not properly underlined. [Bav08a] in contrast produces better results, tending to show less noise overall. Nevertheless, some regions exhibit a sudden volatile noise that can be observed in the close up illustrations at the chin and head of the dragon and around the row of small archs in the cathedral. However, we obtain a satisfying overall quality of occlusion. [Bav08a] indentified the following criteria affecting computation time of SSAO: screen resolution, resolution of AO computation, sample count and size of the applied blur kernel. We benchmarked the Sibenik Cathedral scenario using a constant resolution of 1280x1024 pixels while increasing the sample count stepwise from 16 to 128 samples per pixel. We captured the time for the mere calculation of the SSAO-Pass listed in table 1. As [Bav08a] and [Rei09a] favor a half resolution computation to significantly speed up, we consider this aspect as well.

Samples	**SSAO time [ms]**	
	Full Res.	Half Res.
16	12.4	1.6
36	24.6	3.5
64	47.0	6.1
128	96.9	12.4
Platform: Intel Core i7 920 2,67GHZ 3GB RAM, Nvidia Geforce GTX 285, OpenGL 2.1, GLSL		

Table 1. SSAO computation time and visual quality

Geo.	**SSAO**	**Blur**	**Application**
3.6ms	12.4ms	3.6ms (13x13)	19,6ms / **51.02FPS**
3.6ms	24.6ms	2.8ms (9x9)	31,0ms / **32.2FPS**

Table 2. Overall performance for 16 (top) and 36 Samples (bottom) per pixel

We gain an acceptable performance feasible for realtime applications and present the overall performance for 16 and 32 samples in table 2. The clear loss of quality applying half resolution

rendering is usually compensated by using big blur kernels. We do not favor this method, as satisfying quality is only obtained using high screen resolutions and the quality loss is too significant.

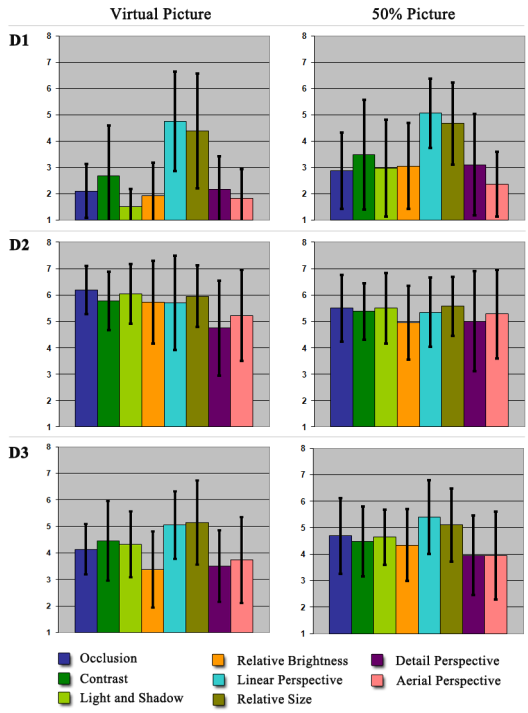

Figure 4. Mean values and standard deviation for test 1 (first col.) and test 2 (second col.)

5. EVALUATION

Even though our approach of keeping the sampling radius constant in image space produces results of enhanced visual quality for scenes with a large spatial extent, we intended to evaluate the impact on the several cues of human depth perception in detail. We further wanted to examine whether there is a change in the perceived criteria of depth using our approach in a mixed reality discrepancy check scenario in contrast to standard SSAO.

Discrepancy checks in industrial applications are performed in-situ at a factory floor with a tracked mobile device [Sch08a]. Therefore, the device is moved at the desired position, and the discrepancy is examined by interactively adjusting the transparency of a plane, textured with the real world video image laying in front of the virtual scene. The visualization is shown on a touch screen mounted on the device. Even if the tracked device can be moved, a single check is mostly done at a fixed position similar to an examination of a freezed image. Due to this setup, binocular or kinetic depth cues [Dra96a] like motion parallax and stereopsis can be neglected in our experiments, keeping focus on the perception of pictorial depth cues. The human perception of depth in virtual and augmented reality environments is a well researched topic in the community [Dra96a, Cut95a, Swa07a, Mur95a]. Nevertheless, the effect of an indirect illumination approximation like SSAO for the several cues of human depth perception has yet not been investigated. According to [Mur95a], we consider depth cues of a constructionist point of view, because the discrepancy check cannot fulfill the key point of the ecological approach regarding the impact of motion. [Mur95a] and [Dra96a] describe various depth cues and additional criteria that have a strong effect on the perception of depth cues. In our experiment we focused on the pictorial depth cues, listed as follows: *Occlusion*, *contrast*, *light and shadows*, *relative brightness*, *linear perspective*, *relative size*, *detail perspective* and *aerial perspective*.

During a period of one week we tested 22 participants (4 female, 18 male) recruited at the Daimler Research Center in Ulm, consisting of students, Phd candidates and scientific assistants aged between 23 and 47 years (M=31.05, SD=6.0). Their experience with virtual or augmented reality applications was asked and yielded a mean value of 4.14 (SD=2.57) on an ordinal scale from 1 to 7. The candidates sat in front of a Toshiba plasma tv-screen with an image diagonal of 42 inches and a distance to the viewer of approximately 55 inches.

Each experiment started by showing every participant three different renderings of a typical factory scene with a large spatial extent. We labeled the renderings *D1* for standard phong shading, *D2* for our approach to SSAO, and *D3* for standard SSAO with a decreasing size of the sample radius in depth. Additionally, the same scene was presented with a 50% alpha blended image taken in the real factory at a viewpoint matching the one of the virtual scene. The several depth cues were explained to the participants, being further instructed to take care of how the three rendering methods support their perception of the different cues, in both the mere virtual and the mixed 50% transparent representation. Next to the tv-screen we put an additional flat screen showing the definitions and sample pictures (not related to any of the test scenes) of the various depth cues for explanation and as a reminder during the tests.

Study Design

Overall, we ran two tests with five different factory planning scenarios. The scenes were randomized to eliminate any scene related effects. In the first test *D1*, *D2*, and *D3* were shown for a virtual scene (figure 5). The three scene representations were arranged side by side on the tv-screen. The resolution of *D1*, *D2* and *D3* was set to match the video resolution (1024x768) of the image processing

camera used for discrepancy checks to obtain comparable results. Since SSAO is a mere indirect effect of GI, and every individual has a different preference of the intensity of such an effect, we gave the participants the opportunity to switch between four slides showing *D2* and *D3* in a varying intensity of the effect, parameterized from weak to strong. *D1* remained unaffected. On a questionnaire the participants marked which of the three representations mostly fulfilled the depth cues as a whole. Next they evaluated their perception of the depth cues for each of the three representations on an ordinal point scale from 1 to 7, or marked an "X" if a conclusion was not found. The participants were asked to assign the points 1 to 7, representing how properly they could perceive the respective depth cue in the whole scene. In the second test *D1*, *D2* and *D3* where shown with a 50% transparent overlay of an image of the real factory, matching the viewpoint of the virtual scene (figure 5). Again the participants were asked to judge which representation mostly fulfills the various depth cues as a whole and assign points ranging from 1 to 7 or mark an "X". Due to former observations, we wanted to prove the hypothesis that some of the depth cues may be perceived differently in the 50% representation and that the general preference for one of the three representations might slightly shift in this scenario. The user study confirmed this assumption, as we will show in the next section. To prove any significance of our results we ran an analysis of variance (ANOVA) for repeated measures of pair wise comparisons with bonferroni correction, as well as a paired t-test to compare the results for the several cues between the virtual and the 50% scenario regarding the respective representation.

Study Results

For the virtual scene all participants marked *D2* as the representation mostly supporting the perception of the depth cues as a whole. For the 50% overlay scene 1 of the participants (4.54%) chose *D1*, 14 (63.63%) chose *D2* and 7 (31.81%) chose *D3*, indicating that SSAO is generally preferred in the discrepancy check scenario. However the advantage of perceiving depth with *D2* compared to *D3* is not that distinct as it was for the mere virtual scenario. The ANOVA for test 1 and test 2 revealed many main effects summarized in table 3 and table 4. The assessment of the several depth cues given by the participants is depicted in figure 4, including the mean value and the standard deviation. Comparing both representations of SSAO to pure phong shading, the ANOVA yields numerous main effects from highly significant ($p<.001$) over very significant ($p<.01$) to significant ($p<.05$), except the cue of linear perspective, where no significant change of the perception between the three representations could be found. Thus, the perception of the linear perspective does not significantly profit from the effect of SSAO. Furthermore, the cue of relative size is unaffected regarding *D1* and *D3*. Comparing *D2* and *D3*, the ANOVA revealed significant or highly significant results, except the cue of relative size yielding only a trend ($p<.10$). The reason for that lies in a stronger perceived effect of the blur and an increased shadowing in depth of our approach. The main effects for *D2* and the 100% selection of *D2* by the participants show that the essential pictorial cues for human depth perception like occlusion [Cut95a] are better perceived with our approach for pure virtual scenes with a large spatial extent. The ANOVA of test 2 revealed numerous main effects as well, but a slight shift in the perception of the several cues can be noticed. For the cue of contrast, we could only assess a trend between *D1* and *D3*, but no significance at all for the cues of detail perspective and relative size. Comparing both SSAO representations with *D1*, no effect was observable regarding the linear perspective, what proved that this cue in reference to test 1 does not take any advantage of the SSAO effect. The perception of the depth cues for the two SSAO representations revealed further interesting differences. Yet no significant change emerged for the cue of occlusion, relative brightness and relative size. The cue of contrast only reveals a trend. However the results for the cue of light & shadow and for the aerial perspective shift from highly significant to very significant, compared to the results of test 1. This could be an explanation why only 65% of the participants chose *D2*, whereas 32% preferred *D3*. Between *D2* and *D3* the most effective depth cue of occlusion [Cut95a] can not be perceived significantly different. The ANOVA of other important cues and additional criteria necessary for the perception of depth, like contrast, relative brightness and relative size also revealed no main effects. Nevertheless, 96% of the participants chose SSAO in general to be the preferred representation for the discrepancy check scenario. *D2* and *D3* do not differ in the perception of important cues, but a tendency to our approach *D2* is still noticeable. Additionally, we ran a paired t-test to evaluate significant changes in the perception of the several cues for *D1*, *D2* and *D3* between the two scenarios. The results are listed in table 5 and show various significant differences, indicating that the perception of depth seriously changes in the 50% overlay-image. According to the significance levels, the perceptual changes in the virtual and 50% scenario appear the most for *D1*. The perception of the cues for *D1* mainly improves in the mixed scene. For *D2* the important cues of occlusion and relative brightness change significantly, with lesser mean values for the

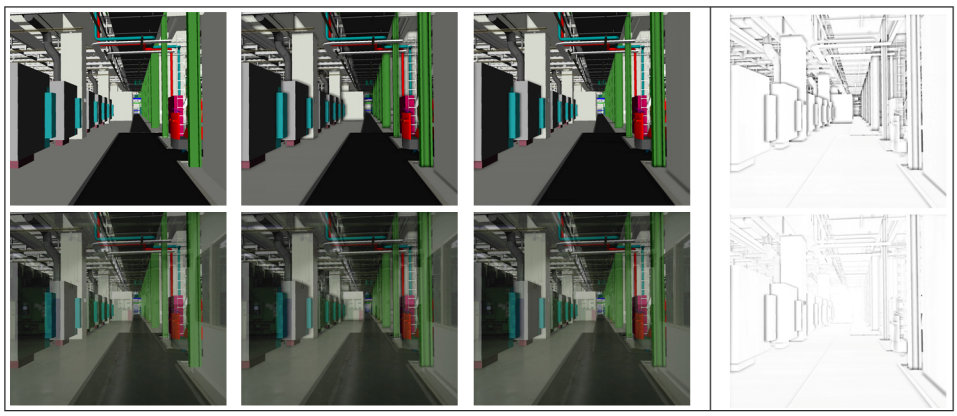

Figure 5. Sample Scene of the user study showing the pure virtual scene (top) and the 50% mixed scene (bottom). The occlusion term for D2 (top right) and D3 (bottom right) is also presented.

mixed scene, and even the clue of light and shadow shows a trends of being perceived worst. *D3* shows a trend of improving the perception for occlusion and also reveals a significant better perception of relative brightness. The strong changes in the perception of such important cues for *D2* and *D3* with *D3* showing a better benefit in the mixed scene explain the choice of 32% of the participants for *D3*. This observation is furthermore confirmed by the results of the ANOVA for test 2, where no main effects between *D2* and *D3* emerged for the important cues of occlusion and relative size.

6. Conclusion

We presented a more accurate realtime approach of SSAO according to the original definition of ambient occlusion. Keeping the sampling radius constant in screen space seriously improves the effect of SSAO for scenes with a large spatial extent. Even if the representation is physical incorrect, a user study confirmed the benefit for the human perception of depth among various cues for pure virtual scenes. Furthermore, a general preference of SSAO for the representation of virtual scenes was observed and compared for several depth cues. For mixed reality scenes with a 50% transparent overlay SSAO also proved to be the appropriate representation with a tendency to our approach. Nevertheless, mixed scenes reveal more balanced differences in perception, especially for *D2* and *D3*. Therefore, our future work will focus on extending SSAO with alternate forms of visualizations and their benefit for the perception in mixed scenarios.

7. ACKNOWLEDGMENTS

This work was partially funded by the AVILUS project of the german Bildungsministerium für Bildung und Forschung (BMB+F). The Dragon is courtesy of the Stanford University and the Sibenik Cathedral is courtesy of Marko Dabrovic

8. REFERENCES

[Bav08a] Bavoil, L., Sainz, M., Dimitrov, R. Image-Space Horizon-Based Ambient Occlusion, ACM SIGGRAPH 2008 talks, 2008.

[Bru06a] Brüderlin, B., Heyer, M., Pfützner, S.: Visibility Guided Rendering to Accelerate 3D Graphics Hardware Performance, in Course Notes EGSR 2006, ACM Eurographics, 2006.

[Cut95a] Cutting,. J.E., Vishton, P.M.: Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Conceptual Use of Different Information about Depth. In: Epstein, W., Rogers, S. editors. Perception of Space and Motion, p. 69-117, 1995.

[Dra96a] Drascic, D., Milgram, P.: Perceptual Issues in Augmented Reality, in: Proc. Stereoscopic Displays and Virtual Reality Systems SPIE 1996, p 123-134, 1996.

[Eng09a] Engel, W. Shader X7: Advanced Rendering Techniques, Charles River Media, 2009.

[Fil08a] Filion, D., McNaughton, R. Effects & Techniques, ACM SIGGRAPH 2008 classes, pp.133-164, 2008.

[Fox08a] Fox, M., Compton, S. Ambient Occlusive Crease Shading, Game Developer Magazine, 2008.

[Kaj09a] Kajalin, V. Screen Space Ambient Occlusion, in: Engel, W. (Editor): ShaderX7: Advanced Rendering Techniques, pp.413-424, 2009.

[Lan99a] Langer, M.S., Bülthoff, H.H. Perception of shape from shading on a cloudy day, Max-Planck-Institut für biologische Kybernetik, 1999.

[Lan02a] Landis, H. Production-Ready Global Illumination, ACM SIGGRPAH 2002 courses, 2002.

[Mit07a] Mittring, M. Finding next gen: CryEngine2, ACM SIGGRAPH 2007 courses, pp. 97-121, 2007.

[Mur94a] Murray, J.: Some Perspectives on Visual Depth Perception, ACM SIGGRAPH Computer Graphics, Vol. 28, Issue 2, 1994.

[Pet04a] Petschnigg, G., Szeliski, R., Agrawala, M. Digital Photography with flash and no-flash image pairs, ACM SIGGRAPH 2004, pp.664-672, 2004.

[Rei09a] Reinbothe, C., Boubekeur, T., Alexa, M. Hybrid Ambient Occlusion, EUROGRAPHICS 2009 Areas Papers, 2009.

[Sch08a] Schoenfelder, R. and Schmalstieg, D. Augmented Reality for Industrial Building Acceptance, Virtual Reality Conference VR 2008, IEEE Computer Society, 2008.

[Sha07a] Shanmugam, P., Arikan, O. Hardware accelerated ambient occlusion techniques on GPUs, I3D'07 (ACM) , pp.73-80, 2007.

[Swa07a] Swan II, J.E., Adam, J., Kolstad, E., Livingston, M.A. and Smallman, H.S. in: Egocentric Depth Judgements in Optical See-Through Augmented Reality, IEEE Transactions on Visualization and Computer Graphics, Vol. 13, No. 3, 2007.

	Test 1		**Test 2**	
Depth Cue	ANOVA	Bonferroni	ANOVA	Bonferroni
Occlusion	F(2;46) = 87.48	D1<D2, p < .001 D1<D3, p < .001 D2>D3, p < .001	F(2;46) = 21.22	D1<D2, p < .001 D1<D3, p < .001 D2>D3, NO
Contrast	F(2;46) = 20.77	D1<D2, p < .001 D1<D3, p < .001 D2>D3, p < .05	F(2;44) = 9.22	D1<D2, p < .01 D1<D3, p < .1 D2>D3, p < .1
Light and Shadow	F(2;46) = 101.5	D1<D2, p < .001 D1<D3, p < .001 D2>D3, p < .001	F(2;46) = 19.82	D1<D2, p < .001 D1<D3, p < .01 D2>D3, p < .01
Relative Brightness	F(2;46) = 53.16	D1<D2, p < .001 D1<D3, p < .01 D2>D3, p < .001	F(2;44) = 12.33	D1<D2, p < .01 D1<D3, p < .05 D2>D3, NO
Linear Perspective	F(2;42) = 2.41	D1<D2, NO D1<D3, NO D2>D3, NO	F(2;42) = 0.81	D1<D2, NO D1<D3, NO D2<D3, NO
Relative Size	F(2;44) = 6.57	D1<D2, p < .05 D1<D3, NO D2>D3, p < .1	F(2;44) = 7.08	D1<D2, p < .01 D1<D3, NO D2>D3, NO
Detail Perspective	F(2;42) = 21.0	D1<D2, p < .001 D1<D3, p < .01 D2>D3, p < .05	F(2;44) = 10.95	D1<D2, p < .01 D1<D3, NO D2>D3, p < .05
Aerial Perspective	F(2;44) = 39.25	D1<D2, p < .001 D1<D3, p < .001 D2>D3, p < .001	F(2;42) = 28.34	D1<D2, p < .001 D1<D3, p < .01 D2>D3, p < .01

Table 3. Statistical results of test 1 and test 2

Depth Cue	D1	D2	D3
Occlusion	T=-3.37, S1<S2, p<.01	T=3.07, S1>S2, p<.01	T=-1.78, S1<S2, p<.10
Contrast	T=-2.07, S1<S2, p<.10	T=1.06, S1>S2, NO	T=-0.12, S1<S2, NO
Light and Shadow	T=-4.18, S1<S2, p<.001	T=1.74, S1>S2, p<.1	T=-0.90, S1<S2, NO
Relative Brightness	T=-3.67, S1<S2, p<.01	T=2.83, S1>S2, p<.05	T=-2.83, S1<S2, p<.05
Detail Perspective	T=-2.39, S1<S2, p<.05	T=-0.53, S1<S2, NO	T=-0.87, S1<S2, NO
Aerial Perspective	T=-1.99, S1<S2, p<.10	T=-0.43, S1<S2, NO	T=-0.73, S1<S2, NO

Table 4. Statistical results for the t-test (S1/S2 being the mean value of test 1/test 2)