Exploitation of human shadow perception for fast shadow rendering

Mirko Sattler Ralf Sarlette Thomas Mücken Reinhard Klein

University of Bonn
Institute for Computer Science II, Bonn, Germany
Email: {sattler,sarlette,muecken,rk}@cs.uni-bonn.de

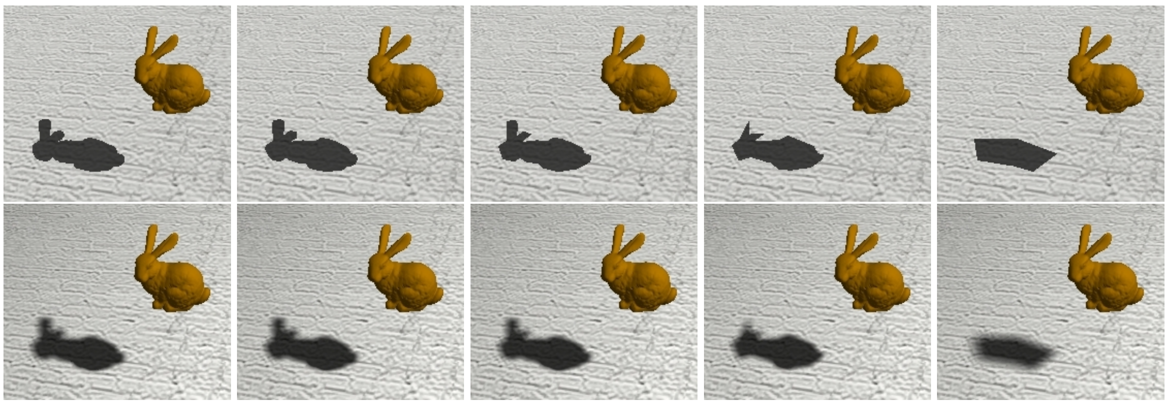

Figure 1: Visual perception of shadows. Decreasing level-of-detail for the shadow caster object from left to right. Hard shadows cast by a point light source (top row) and soft shadows cast by an area light source (bottom row) are shown.

Abstract

In this paper we describe an experiment to obtain information about the perceptual potential of the human visual system regarding shadow perception. Shadows play an important part for communicating spatial structures of objects to the observer. They are also essential for the overall realism of the rendered image. Unfortunately, most algorithms in computer graphics which are capable of producing realistic shadows are computationally expensive. The main idea behind the experiment is to use a simplified version of the shadow caster to generate hard and soft shadows, which would rapidly increase performance and to evaluate up to which degree a simplification is possible, without producing noticeable errors. Therefore, an experiment is performed, in which the test persons should mark the point of the just noticeable difference. First results show, that a mesh simplified to only 1% of its original complexity is capable to cast soft shadows that satisfy 90% of the test persons.

CR Categories: I.3.7 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism—;

Keywords: human perception, soft shadows, LOD, rendering

1 Introduction

Shadows are an important visual clue about the spatial structure of an object. Using virtual reality applications for example, such as life-sized cloth visualization, or medical surgery planning, they are important because they increase the presence of the virtual objects and the overall realism. For the entertainment industry, where realistic images and high frame-rates are desired, this is also an important aspect. However, the shadows need to be computed at interactive frame rates, otherwise usability and presence will break down. While local illumination models are the strength of modern graphics hardware, more advanced techniques which are capable of generating e.g. soft shadows are notoriously hard to perform efficiently. For the overall realism of a computer generated image, the used illumination is also important. While the classical graphics adapter pipeline only supports artificial point light sources, realistic images would require area light sources or image based illumination with environment maps. But, if the indispensable ingredients to generate a realistic image are shadows cast by area light sources, to what extend have these shadows to be correct, to be visually accepted by a human observer? If it is true, that the human observer is not very good in detecting slight errors in shadow visualization, it would be possible to use a simplified model of the object to cast the shadows, instead of the original one and achieve the same *acceptance rates* of the resulting images. Therefore, the main idea of the proposed experiment is to evaluate which level-of-detail (LOD) is sufficient for the shadow casting object to produce acceptable shadows. The rest of the paper is organized as follows. Section 2 gives a brief overview to related work, Section 3 describes the experimental setup and the program used. In Section 4 we present the results and conclude in Section 5.

2 Related Work

There exists a lot of work on shadow perception. Wanger et al.[1992; 1992] have done experiments about the context of object spatial position and size and shadow shape and sharpness for simple objects. Other experiments show that shadows are an import visual clue for object-object contact [Hu et al. 2000; Madison et al. 2001]. The importance of object motion and shadows for spatial perception was investigated by Kersten et. al [1996; 1997]. However, a vast amount of work has also been published on rendering shadows. Shadow maps [Williams 1978] or shadow volumes [Crow 1977] and all derivates are the classical approaches for point-like light sources. Both algorithms are well suited for todays graphics adaptors [Kilgard 2002]. An important detail of the shadow volume algorithm is the detection of silhouette edges. A silhouette edge is an edge, where one of the normals of the corresponding faces points towards the viewer (or the light source) and the other points away. Therefore it is necessary to test all edges, each time the light source moves. There exist some approximations [Markosian et al. 1997], which might lead to artifacts. Recently, more advanced techniques for soft shadow generation were developed, using various techniques, like wedges [Akenine-Möller and Assarsson 2002; Assarsson and Akenine-Möller 2003; Assarsson et al. 2003], smoothies [Chan and Durand 2003] or penumbra maps [Wyman and Hansen 2003]. A good overview of algorithms which produce soft shadows can be found in Hasenfratz et al. [2003]. Mesh simplification is one of the fundamental techniques used for polygonal meshes, there is an extensive amount of different methods. See e.g. [Luebke 2001] for a detailed review of simplification algorithms. As a measurement criterium for the simplification, the Hausdorff-Error can be used. The Hausdorff-Error describes the geometrical distance between the original and the simplified mesh and vice versa. Klein et al. [1996] first used it to control the simplification. Borodin et al. [2003] have produced high-quality results by combining generalized pair contractions an extension of the vertex pair contraction with the control of the distance between the original and simplified models during the whole simplification process. In 2005, the computationally expensive dense regular sampling was replaced with a significantly faster adaptive sampling method by Guthe et al. [2005].

3 Experimental setup

The program used for the experiment (see Figure 2) shows two different images of the same scene. The left side of the screen shows a high resolution version of the shadow casting object above a plane. This high resolution version is also used to calculate the shadows. The object shown on the right side is also the high resolution version, but in contrast to the left side, the object to calculate the shadows is a simplified version of the original. In the bottom area of the graphical user interface several buttons for program control can be seen.

To compute the shadows, we use a modified version of the soft shadow algorithm [Assarsson et al. 2003]. This version uses shadow volumes, to compute the umbra region. As test object the *Stanford Bunny* [Curless and Levoy 1996] consisting out of about 70000 triangles is used. The simplified versions of the shadow casting object were generated using techniques from [Borodin et al. 2003]. Thirty six level of detail were precalculated, ranging from 100 (LOD=1) up to 50000 triangles (LOD=36). In each level the number of triangles is increased. See Figure 3 for details. In the current version of the experiment, the shadow receiver is a plane. As the distance between the object and the plane we chose twice the object size, since this allows a wide range of viewing angles from which both object and shadows are fully visible. Increasing the distance further would make side views impossible, while reducing the distance would increase the minimum viewing angle above 30 degrees. Note, that the radial size of the area light source also is chosen accordingly. During the experiment, the test person is able to move the light source and the point of view around the object, while the viewing distance is fixed. Thus, it is possible to examine the generated shadows under several viewing angles. The current LOD can also be changed interactively. On the left side of the screen, always the highest LOD with shadows is shown for comparison.
The Hausdorff-Error between each LOD level and the original mesh is precalculated. This allows for the comparison between the original mesh and the simplified version. In Figure 3 the Hausdorff-Error is given in percentage of the bounding box diagonal.

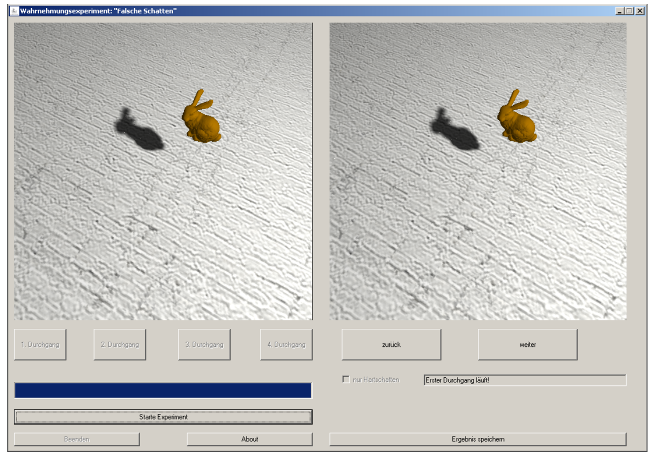

Figure 2: Program used for the perception experiment

3.1 Experimental procedure

The experiment is performed in the following way. Perception is tested in two directions. First, starting with the highest level-of-detail of the mesh (LOD=36), the LOD number decreases. That is, the number of triangles used for calculating the shadows, decreases. Following the tradition of the "Just Noticeable Difference" experiment by Weber [1834], the test person can mark the level, which seems to be the first to produce noticeable errors in the shadows. The second part of the experiment starts with the lowest LOD level and the test person this time marks the level, which seems to produce correct shadows for the first time. These two parts are repeated with a different size of the light source. This allows predictions about the influence of the size of the penumbra on the shadow perception. Then, the area light source is substituted by a point light source and the experiment is performed again, now with hard shadows. For the experiment, 20 test persons were interviewed.

4 Results & Discussion

Because the perception of shadows by an observer differs from human to human, there can not be an exact transition point between *realistic* and *artificial* perception, but more or less a transition region, in which most observers will fall. The results of the first two parts of the experiment (area light sources) are shown in Figure 4. In the first part of the experiment 17 test persons (=85%) are between

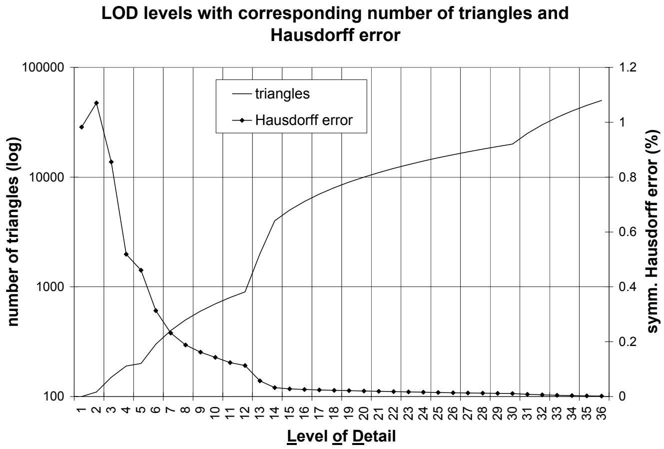

Figure 3: LOD numbers with corresponding triangle count and Hausdorff error.

200 (LOD=5) and 600 (LOD=9) triangles. 8 (=40%) of these persons were satisfied with a shadow cast from only 200 triangles and 25% from 300 triangles. In the second part of the experiment (increasing LODs) nearly the same results were observed. The mean value for the large area light sources is around LOD=5.6. The third and fourth part of the experiment (larger light source) show similar results as the first and second part, as shown in Figure 5. It seems, that the size of the penumbra region has only little influence on the perception. The mean value for the large area light sources is around LOD=5.4. To summarize the results for soft shadows and the object *Stanford Bunny*: less than 1% of the the original number of triangles for shadow calculation are sufficient to produce realistic shadows for a majority (=90 %) of the test persons. On the computer which was used for the experiment, the rendering time per frame for the original mesh is about 1.15 seconds. Using only LOD=9, a speed-up of factor 16 can be achieved (0.07 seconds rendering time). In the last two parts of the experiments *hard shadows* are rendered. The scattering of the hits is more obvious, as in the parts with soft shadows, as can been seen in Figure 6. For both directions, the majority of test persons (=87.5 %) is between LOD=5 and LOD=17. Besides that, some outliers can be seen at LOD=30 and LOD=36. The mean value for the point light source is around LOD=10.9. It seems, that errors in hard shadows are noticed more early, than in soft shadows, which is reasonable, since soft shadows blur fine details. The possible performance gain in rendering speed is about factor 7 from 0.129 seconds for the original mesh to 0.019 seconds using LOD=17 with 7000 triangles.

5 Conclusions & Future Work

We have presented an experiment to exploit the human shadow perception for acceleration of the rendering of shadows in a computer generated image. Our first results suggest, that it is possible to use a highly simplified instead of the original model to generate soft shadows. As a direction of future work, we like to evaluate, if our results can be generalized to other objects. Therefore, we will increase the number of objects used. Additionally, we want to use a 3D shape, over which we have parametric control. This could be e.g. an object consisting of few cylinders. Furthermore, we only used a planar surface as shadow receiver. Therefore, no irregularities can distract the viewer from shadowing errors. However, this can be the case, if the plane is substituted by an irregularly shaped surface. Another direction is the usage of real-world scenes with a much higher complexity and larger shadow regions.

References

AKENINE-MÖLLER, T., AND ASSARSSON, U. 2002. Approximate soft shadows on arbitrary surfaces using penumbra wedges. In *Proceedings of the 13th Eurographics workshop on Rendering*, Eurographics Association, 297–306.

ASSARSSON, U., AND AKENINE-MÖLLER, T. 2003. A geometry-based soft shadow volume algorithm using graphics hardware. *ACM Trans. Graph. 22*, 3, 511–520.

ASSARSSON, U., DOUGHERTY, M., MOUNIER, M., AND AKENINE-MÖLLER, T. 2003. An optimized soft shadow volume algorithm with real-time performance. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, Eurographics Association, 33–40.

BORODIN, P., GUMHOLD, S., GUTHE, M., AND KLEIN, R. 2003. High-quality simplification with generalized pair contractions. In *Proceedings of GraphiCon'2003*, 147–154.

CHAN, E., AND DURAND, F. 2003. Rendering fake soft shadows with smoothies. In *Proceedings of the Eurographics Symposium on Rendering*, Eurographics Association, 208–218.

CROW, F. C. 1977. Shadow algorithms for computer graphics. In *SIGGRAPH '77: Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, ACM Press, 242–248.

CURLESS, B., AND LEVOY, M. 1996. A volumetric method for building complex models from range images. *Computer Graphics 30*, Annual Conference Series, 303–312.

D. KERSTEN, P. M., AND KNILL, D. C. 1997. Moving cast shadows induce apparent motion in depth. In *Perception*, vol. 26, 171–192.

GUTHE, M., BORODIN, P., AND KLEIN, R. 2005. Fast and accurate hausdorff distance calculation between meshes". 41–48.

HASENFRATZ, J., LAPIERRE, M., HOLZSCHUCH, N., AND SILLION, F. 2003. A survey of real-time soft shadows algorithms. In *Eurographics*, Eurographics, Eurographics. State-of-the-Art Report.

HU, H. H., GOOCH, A. A., THOMPSONA, W. B., SMITS, B. E., SHIRLEY, P., AND RIESER, J. J. 2000. Visual cues for imminent object contact in realistic virtual environments. In *IEEE Visualization 2000*.

KERSTEN, D., KNILL, D., MAMASSIAN, P., AND BÜLTHOFF, I. 1996. Illusory motion from shadows. *Nature 379*, 31.

KILGARD, M. J. 2002. Shadow mapping with todays opengl hardware. In *SIGGRAPH 2002 Course*.

KLEIN, R., LIEBICH, G., AND STRASSER, W. 1996. Mesh reduction with error control. In *IEEE Visualization '96*, R. Yagel and G. M. Nielson., Eds., 311–318.

LUEBKE, D. 2001. A developers survey of polygonal simplification algorithms. vol. 23, 24–35.

MADISON, C., THOMPSON, W. B., KERSTEN, D. J., SHIRLEY, P., AND SMITS, B. S. 2001. Use of interreflection and shadow for surface contact. In *Perception and Psychophysics*, vol. 63.

MARKOSIAN, L., KOWALSKI, M. A., GOLDSTEIN, D., TRYCHIN, S. J., HUGHES, J. F., AND BOURDEV, L. D. 1997. Real-time nonphotorealistic rendering. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 415–420.

WANGER, L. C., FERWERDA, J. A., AND GREENBERG, D. P. 1992. Perceiving spatial relationships in computer-generated images. *IEEE Comput. Graph. Appl. 12*, 3, 44–51, 54–58.

WANGER, L. 1992. The effect of shadow quality on the perception of spatial relationships in computer generated imagery. In *SI3D '92: Proceedings of the 1992 symposium on Interactive 3D graphics*, ACM Press, New York, NY, USA, 39–42.

WEBER, E. H. 1834. *De Pulsu, Resorptione, Auditu et Tactu - Annotationes Anatomicae Et Physiologicae*. Koehler.

WILLIAMS, L. 1978. Casting curved shadows on curved surfaces. In *SIGGRAPH '78: Proceedings of the 5th annual conference on Computer graphics and interactive techniques*, ACM Press, 270–274.

WYMAN, C., AND HANSEN, C. 2003. Penumbra maps: Approximate soft shadows in real time. In *Proceedings of the 2003 Eurographics Symposium on Rendering*, Eurographics Association, 202–207.

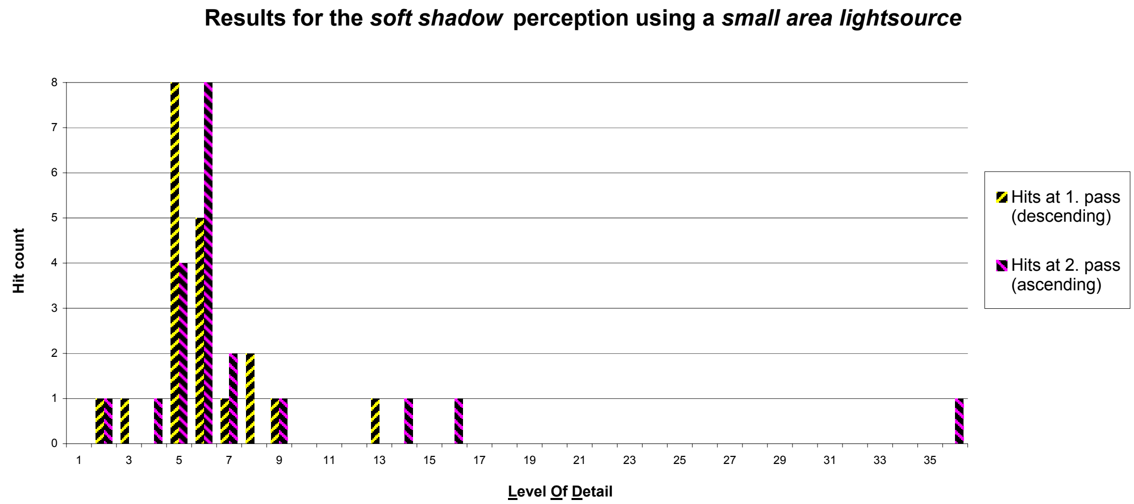

Figure 4: Results of the first two parts of the experiment for an area light source casting soft shadows.

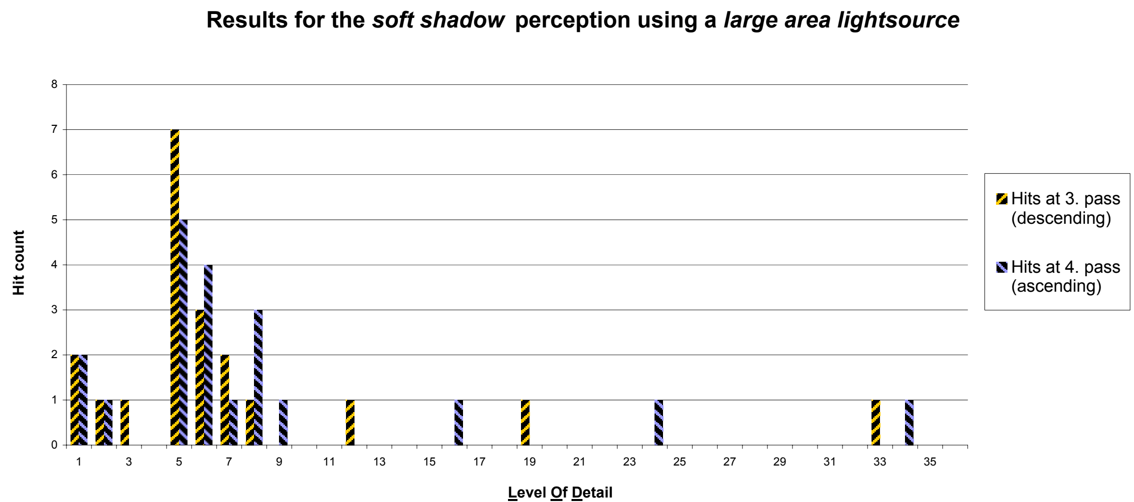

Figure 5: Results for an increased size of the area light source.

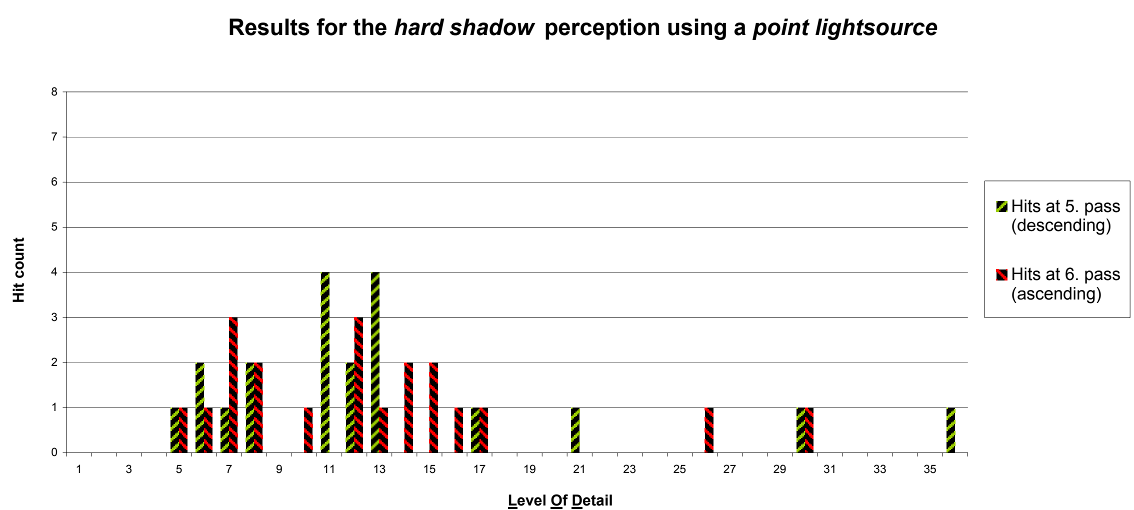

Figure 6: Results for hard shadows cast by a point light source.