Visual Acceptance Evaluation of Soft Shadow Algorithms for Virtual TV Studios

Hüseyin Ayten and Jens Herder
Department of Media
FH Düsseldorf, University of Applied Sciences
Josef-Gockeln-Str. 9, D-40474 Düsseldorf, Germany
voice/fax: [+49](211) 4351-810/803
www: vsvr.medien.fh-duesseldorf.de
{hueseyin.ayten,herder}@fh-duesseldorf.de

Wolfgang Vonolfen
Fraunhofer Institute for Intelligent Analysis and Information Systems IAIS
Schloss Birlinghoven, 53757 Sankt Augustin, Germany
voice: [+49] (2241) 14-3425
www: www.iais.fraunhofer.de
wolfgang.vonolfen@iais.fraunhofer.de

ABSTRACT

Shadows in computer graphics are an important rendering aspect for spatial objects. For realtime computer applications such as games, it is essential to represent shadows as accurate as possible. Also, various TV stations work with virtual studio systems instead of real studio sets. Especially for those systems, a realistic impression of the rendered and mixed scene is important. One challenge, hence, is the creation of a natural shadow impression. This paper presents the results of an empirical study to compare the performance and quality of different shadow mapping methods. For this test, a prototype studio renderer was developed. A percentage closer filter (PCF) with a number of specific resolutions is used to minimize the aliasing issue. More advanced algorithms which generate smooth shadows like the percentage closer soft shadow (PCSS) method as well as the variance shadow maps (VSM) method are analysed. Different open source APIs are used to develop the virtual studio renderer, giving the benefit of permanent enhancement. The Ogre 3D graphic engine is used to implement the rendering system, benefiting from various functions and plugins. The transmission of the tracking data is accomplished with the VRPN server/ client and the Intersense API. The different shadow algorithms are compared in a virtual studio environment which also casts real shadows and thus gives a chance for a direct comparison throughout the empirical user study. The performance is measured in frames per second.

Categories and Subject Descriptors

I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*; H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*augmented virtualities*

HC'10 , Dec. 8 – 10, 2010, Aizu-Wakamatsu, Japan.

General Terms

Algorithms, Experimentation, Verification

Keywords

Virtual Studio, Soft Shadow, Realtime Rendering, Augmented Reality

1. INTRODUCTION

The application of virtual studio technology is increasingly important in today's TV and video productions and becomes widely accepted. In an earlier article [7], we demonstrated our interactive virtual studio (TV) for post production, using the camera tracking system IS-900 SCT [16] from Intersense and the 3dk rendering system. In this study we use the same tracking, but developed a new rendering system based on Ogre3d for studying realtime shadows for virtual sets (TV).

The interaction control as well as multiple interfaces of a virtual set environment have been discussed and classified in [4]. An overview about virtual set environments and an introduction to virtual studios can be found in [3]. The requirements for virtual studios usually differ from those for post productions, a constant real-time frame rate of e.g. 50 half images per second for the PAL video format (60 Hz for NTSC) and accurate image quality is needed [10]. This is because, in virtual studios, the composition of real objects and computer generated graphics is supposed to take place immediately during the production of the real shots, thus making post production tasks superfluous.

A general discussion of virtual and real light/shadow interaction for virtual studio applications can be found in [15]. We discussed the light matching problems and requirements for virtual studio application in [5] before. In [6] we experimented with shadow map rendering with HDR-based lighting estimation. Real light in the virtual studio controlled the virtual shadow. Soft shadows were achieved using many light sources implying many shadow maps. The prototype at that time suffered from low resolution shadow maps and low frame rate.

2. SHADOW ALGORITHMS

In the following sections we briefly describe four widely used shadow map rendering algorithm and show the rendering for an example set.

Figure 1: Standard shadow map

2.1 Shadow Maps

The standard shadow mapping algorithm [14] contains three steps: 1. the scene rendered from the light perspective and the z-value (depth) are stored in the shadow map. 2. the scene rendered from the camera perspective and the z-values are transformed in the light space. 3. both z-values are compared and if $z_k > z_l$, the pixel will be drawn with the shadow color. If the resolution of a shadow map is low, aliasing artefacts like stairs become visible at the border of the shadow. Also the edges of a shadow are sharp like using a point light source. Figure 1 shows the standard shadow map rendering together with a person casting a natural shadow in a virtual set.

2.2 Percentage closer filter (PCF)

The difference to the percentage closer filtering [12] is that the result of the depth comparison will be filtered. The reason is to minimize artefacts and to get softer shadow edges. Depending on the size of the filter the neighbour depths around each pixel are compared with the actual depth-value. The result determines the percentage of shadow for the respective pixel. Figure 2 shows the rendering of a virtual set including the respective virtual PCF shadow together with a real person casting a natural shadow.

2.3 Variance shadow maps (VSM)

Variance Shadow Maps [1] store the depth and squared depth values in the shadow map. One advantage is that the shadow map can be pre-filtered with standard CG filtering (mipmap, anisotropic filtering). The variance of the depth values is calculated to determine the shadow degree for a pixel. Figure 3 shows the rendering of a virtual set including the respective virtual VSM shadow together with a real person casting a natural shadow.

2.4 Percentage closer soft shadow (PCSS)

The percentage closer soft shadow [2] generates shadows with a core (umbra) and a semi intensive border (penumbra). To generate the penumbra, the algorithm assumes that the shadow receiver is placed in parallel to the shadow caster. The light size as well as the distance between shadow receiver and shadow caster is used to calculate correct um-

Figure 2: PCF - percentage closer filter

Figure 3: VSM - variance shadow maps

Figure 4: PCSS - percentage closer soft shadow

bra and penumbra sizes. At first, the algorithm searches in a calculated region of the shadow map for shadow-blockers and then stores the mean value of all found depths. Second, it calculates the penumbra and finally does the PCF filtering. Figure 4 shows the rendering of a virtual set including the respective virtual PCSS shadow together with a real person casting a natural shadow.

3. IMPLEMENTATION

For our testing, we use the following system components: The perspective of the studio camera is tracked by the Intersense tracking system IS-900 SCT. A VRPN server [13] distributes the tracking data to the other components of the virtual studio system. The virtual studio render engine receives the camera parameters over Ethernet, using a VRPN client as part of the render loop. The graphics engine is based on Ogre3D [9]. Directshow API is used for the integration of video textures into the virtual set. The video output is generated by an NVidia Quadro FX 5600 SDI graphics/video card. The graphic workstation is based on an Intel Xeon x5472, 3 Ghz, 8 GB.

4. EVALUATION

The evaluation comprises a comparison of the different approaches with respect to visual quality as well as the performance. It is, however, mainly based on the user study as discussed in the third part of this section.

4.1 Visual quality

The visual quality of the four different shadow rendering approaches can be seen in figure 5. Each shadow procedure was configured in order to deliver satisfactory results (i.e. without artefacts and with acceptable performance). During the tests, the following settings were used: The PCF shadow method (see top left sample) based on a shadow texture of 2048x2048 pixels and a 16-step filtering. The PCSS shadow method (see top right sample) based on a shadow texture of 1024x1024 pixels, a blocker search over 16 steps and 16 steps filtering. The standard shadow map (see bottom left sample) based on a shadow texture of 4096x4096 pixels without any optimizations. The VSM shadow method (see bottom right sample) based on a shadow texture of 1024x1024 pixels and an 8x8 filter.

4.2 Performance

The approaches for the shadow calculation as used in the empirical studies for this paper require an additional rendering pass and thus consume considerable GPU and/or CPU) power on top of the usual rendering. Hence, selecting the best algorithm suitable for professional realtime production purposes also requires to consider performance aspects rather than only visual quality. The performance depends on various factors such as the scene complexity (i.e. the number of polygons), the usage of visual effects (like e.g. particle systems), the size of the rendered image (screen resolution), etc. An additional performance slow down is then caused by the shadow algorithm, the shadow map resolution, and the degree of edge smoothing. A virtual studio render engine, on the other hand, usually requires a constant 50 (PAL) or 60 (NTSC) Hz frame rate and might be even the double when rendering stereo. An easy way to compare the realtime ability of the various approaches is to measure the resulting frame rate when using the same virtual set. We measured the performance for different screen resolutions, the SD PAL format with a resolution of 720x576 pixels and the HD PAL format of 1920x1080 pixels. To determine the dependencies between set complexity and shadow approach, sets with different numbers of polygons were tested. Furthermore, different numbers of textures showing a life video feed (from zero up to two video textures) both read and decoded from the harddisk were used in the virtual scene. All these features were taken into consideration for the performance tests because different combinations represent typical situations in a professional broadcast TV production.

Shadow Algorithm	no video	one video	two videos
`Shadow Map`	989	532	132
`VSM` with 4x4 Filter	513	412	128
`PCF` with 16 tabs	449	357	125
`PCSS` with 16 tabs	416	332	121

Table 1: Performance in frames per second for the scene rendered in SD PAL resolution and different number of video textures

Table 1 shows the results of the performance tests when using a render engine with SD PAL resolution and a shadow map texture-size of 1024x1024 pixels. When using one or two video feeds, the first video feed has a SD PAL resolution and in case of two videos the second video feed a HD 50 Hz resolution. The video streams were read in from the harddisk and decoded by the CPU and then transferred to the video texture space of the graphic board. Calculation of shadows is done by a shader on the graphic board. Without video texture or when using only a single video (SD PAL), a very high frame rate could be achieved with all rendering methods. Among those, the highest frame rate was reached when using the standard shadow map algorithm. This is because expensive processing (e.g. for blurring) is not part of this approach. The PCSS method, on the other hand, shows the lowest frame rates due to the more complex calculation on the GPU. All four algorithms show very similar frame rates when two videos are played back. This is because the shadow calculation via shaders on the GPU run in parallel to

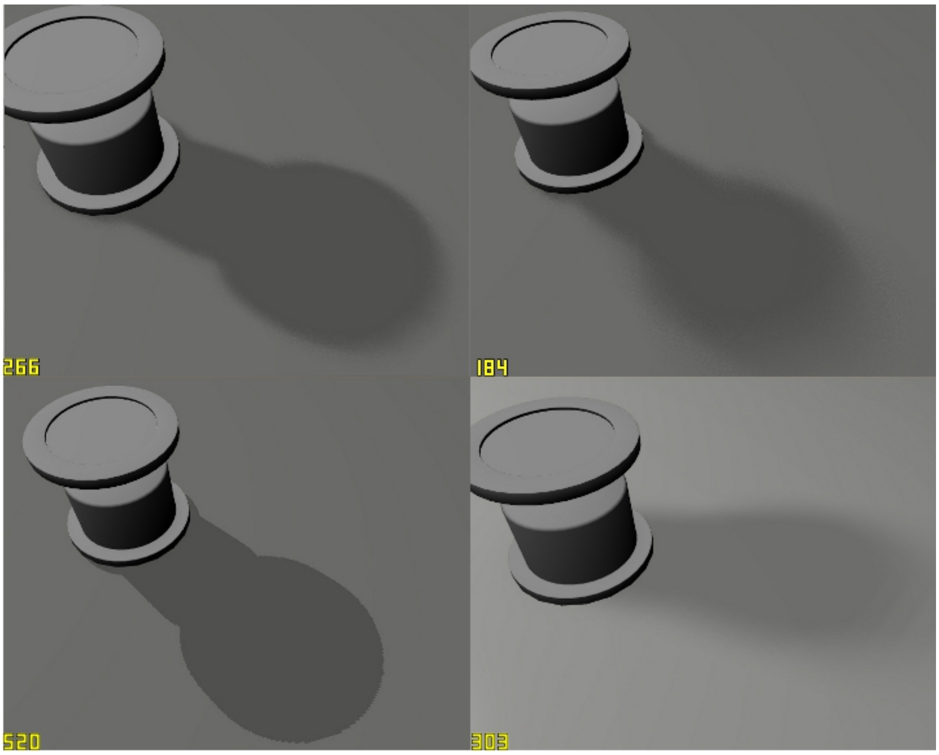

Figure 5: Comparison of four different realtime shadow rendering approaches: PCF **(top left),** PCSS **(top right), shadow map (bottom left), and** VSM **(bottom right)**

the video texture operations. I.e. they mainly have to wait for the completion of the video transfer.

Shadow Algorithm	no video	one video	two videos
`Shadow Map`	411	329	130
`VSM` with 4x4 Filter	236	205	128
`PCF` with 16 tabs	198	174	125
`PCSS` with 16 tabs	149	140	118

Table 2: Performance in frames per second for the scene rendered in HD 50 Hz **resolution and different number of video textures**

Table 2 summarizes the results from the tests with the virtual set rendered in HD 1080P 50 Hz resolution. The first column of both tables (see also table 1) represents frame rates for the scene without video textures and indicates a performance slow-down of more than 50% whereas the rightmost columns of both tables containing the values for two video textures still show quite similar values. As a consequence, a low usage of CPU power reveals the big differences between HD and SD introduced by the various shadow algorithms. This is occluded when computational efforts by the CPU increases due to a higher number of video feeds.

4.3 User study

In this section, we present the results from the empirical studies. The goal was to evaluate and compare the visual quality of the rendered shadows from the different approaches. Fifteen participants were interviewed to find out the acceptance of the soft shadows. (This is the minimum number of participants recommended by the ITU [8].) Four video clips with the different shadows in random order were shown to each proband. The presented content was created in a virtual studio to contain a virtual scene with virtual shadows from 3D objects together with a real shadows from a talent. The real shadow was mixed into the virtual scene using the usual chroma keyer. Then, the participants were asked the following four questions:

1. Do you like the shadow?
2. Which shadow would you prefer?
3. Does the shadow improve the realism?
4. Which does improve the realism best?

Description	Scale	PCF	VSM	PCSS	S. Map
Don't like	0	1	0	0	5
Not much	1	1	4	1	3
Like it	2	8	4	3	6
Like it much	3	4	5	6	1
Like it very much	4	1	2	5	0

Table 3: Questionnaire results for "Do you like the shadow?"

Table 3 shows the assessments for the first question. The first two columns depict the marking. The following columns

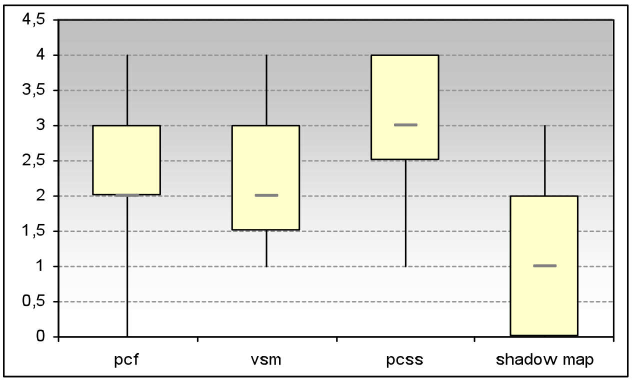

Figure 6: Rating for the appeal of the shadow impression

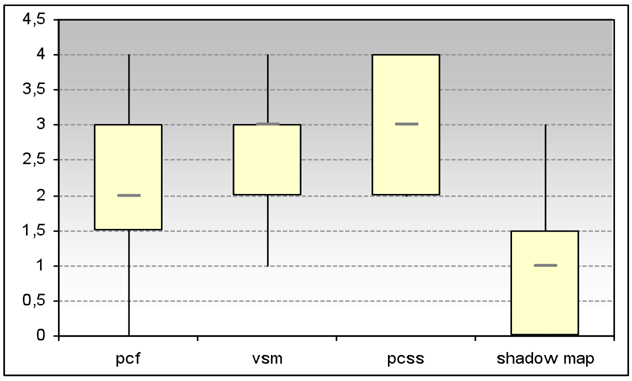

Figure 7: Realism rating of the shadow impression

show the counts for the different assessments, one for each shadow algorithm. An average mark assessment of 53% can be observed for the PCF algorithm. The other shadowmethods deliver an average assessment of VSM = 40%, PCSS = 62% and Shadow map = 28%. That is, according to this empirical study, the PCSS algorithm evaluates to be the shadow algorithm with the best acceptance. The VSM algorithm also creates soft shadow edges, but however, was valued quite low. The PCF method delivered average results and the standard shadow map method the worst.

Figure 6 and 7 each show a boxplot result of the interviews, which illustrate the distribution of the rating for each shadow method. In both cases, the rating for the appeal of the shadow impression and the rating for the realism of the shadow impression, the PCSS shadow method can be identified as the one with the highest rates and with small variation, too.

	PCF	VSM	PCSS	S.Map
Persons	1	4	10	0

Table 4: Questionnaire results for "Which would you prefer?"

Table 4 shows the results of a direct voting. The table entries are the number of participants that voted for the respective shadow method. Again, the PCSS method delivered the best results which is 62% of the participant's votes. The shadow algorithm VSM reached 25% of the votes, the PCF algorithm 12% and nobody preferred the Shadow Map algorithm.

Description	Scale	PCF	VSM	PCSS	S. Map
No	0	1	0	0	5
Not much	1	3	3	0	6
Yes	2	5	4	5	3
Yes much	3	4	5	4	1
Yes it does very much	4	2	3	6	0

Table 5: Questionnaire results for "Does the shadow improve the realism?"

Table 5 summarizes the answers to the third question where the participants should estimate the shadow method's contribution to the realistic or natural look of the scene. As a result, the PCSS algorithm reached the best ratings of 76%, followed by VSM with 63%, the PCF with 55%, and finally, the Shadow Map with 25%. Comparing the results with those from the first question shows similar results with respect to order and percentage. PCSS is evaluated best and Shadow Map worst. VSM and PCF changed order but did not vary significantly.

	PCF	VSM	PCSS	S.Map
Persons	2	4	9	0

Table 6: Questionnaire results for "Which does improve the realism best?"

Table 6 contains the results from a direct voting for a shadow method regarding the best realistic or natural look. The resulting order (PCSS with 60% of the votes, VSM with (26%), PCF with 13%, and no votes for Shadow Map) corresponds to the results of the third question. That is, we have the results of the first two questions affirming each other as well as the results of the last two questions affirming each other, too.

5. CONCLUSION

We implemented a prototype virtual studio rendering system to test different shadow rendering methods with respect to similarity to natural shadows and acceptance by observers. For this purpose, we made empirical studies with participants to assess the virtual shadow quality. The study was accompanied by performance tests because for professional virtual studio productions both aspects, the realtime performance as well as the image quality, play an important role and therefore must be considered in common. In our studies, it turned out that the shadow computation performance may be camouflaged when major parts of the overall scene computation is done by the CPU. This is because soft shadow algorithms can be implemented on the GPU of the graphics board, thus running in parallel. Then, it might be recommendable to choose the best soft shadow algorithm. When scene complexity or additional effects will slow down the rendering process on the graphics board, however, performance becomes more important. According to our empirical studies, the PCSS method delivers best quality results, but also consumes the most GPU power. If performance matters, the VSM and PCF methods might be an alternative both lying in the midfield of the participant's quality ratings and also slightly improving the GPU performance GPU as compared to the PCSS method. For mixed reality environments like our virtual studio setting, rendering quality might be more important than for virtual environments because direct comparison of real and virtual shadow are visible (compare with [11]). In any case, the Shadow Map method delivered worst ratings regarding the quality but also shows best rendering performances. As a consequence, Shadow Maps should be preferred on complex scenes whereas the PCSS method should be chosen, when enough GPU power is available.

Acknowledgements

Discussion and encouragement were given by Rainer Maguhn from vr3 virtual production oHG. Thanks belong also to Antje Müller for keeping the studio equipment running.

6. REFERENCES

[1] W. Donnelly and A. Lauritzen. Variance shadow maps. In *I3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games*, pages 161–165, New York, NY, USA, 2006. ACM.

[2] R. Fernando. Percentage-closer soft shadows. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Sketches*, page 35, New York, NY, USA, 2005. ACM.

[3] S. Gibbs, C. Arapis, C. Breiteneder, V. Lalioti, S. Mostafawy, and J. Speier. Virtual studios: An overview. *IEEE Multimedia*, 5(1):18–35, March 1998.

[4] J. Herder. Interactive content creation with virtual set environments. *Journal of the 3D-Forum Society, Japan*, 15(4):53–56, Dec. 2001.

[5] J. Herder. Matching light for virtual studio tv productions. In *9th Int. Conf. on Human and Computer*, pages 158–162, Aizu-Wakamatsu, Japan, Sept. 2006. University of Aizu.

[6] J. Herder, C. Neider, and S. Kinuwaki. Hdr-based lighting estimation for virtual studio (tv) environments. In *10th Int. Conf. on Human and Computer*, pages 111–117, Düsseldorf, Germany, Aizu-Wakamatsu, Japan, Dec. 2007. FH Düsseldorf, University of Aizu.

[7] J. Herder, W. Vonolfen, A. Griesert, S. Heuer, A. Hoffmann, and B. Höppner. Interactive virtual set applications for post production. In *Seventh Int. Conf. on Human and Computer*, Aizu-Wakamatsu, Japan, Sept. 2004. University of Aizu.

[8] ITU. Methodology for the subjective assessment of the quality of television pictures. In *ITU-R BT.500-11, Recommendation*. 2002. http://www.itu.int.

[9] G. Junker. *Pro Ogre 3D Programming*. Apress, USA, 2006.

[10] M. Moshkovitz. *The Virtual Studio Technologie and Techniques*. Focal Press, Butterworth-Heineman, 225 Wildwood Avenue Woburn, MA 01801-2041, 2000. ISBN 0-240-80425-2.

[11] L. Raya, S. Mata, and O. D. Robles. Perceptual evaluation of illumination effects in virtual environments. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, APGV '10, pages 167–167, New York, NY, USA, 2010. ACM.

[12] W. T. Reeves, D. H. Salesin, and R. L. Cook. Rendering antialiased shadows with depth maps. *SIGGRAPH Comput. Graph.*, 21(4):283–291, 1987.

[13] R. M. Taylor, II, T. C. Hudson, A. Seeger, H. Weber, J. Juliano, and A. T. Helser. Vrpn: a device-independent, network-transparent vr peripheral system. In *VRST '01: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 55–61, New York, NY, USA, 2001. ACM.

[14] L. Williams. Casting curved shadows on curved surfaces. *SIGGRAPH Comput. Graph.*, 12(3):270–274, 1978.

[15] A. Wojdala, M. Gruszewski, and R. Olech. Real-time shadow casting in virtual studio. *MGV (Machine Graphics and Vision*, 9(1/2):315–329, 2000.

[16] D. Wormel and M. Read. Unified Camera, Content and Talent Tracking in Digital Television and Movie Production. In NAB *2000, Las Vegas*, 73 Second Avenue Burlington, Massachusetts 01803 USA, 2000. InterSense Inc.