REAL-TIME SHADOW CASTING USING FAKE SOFT SHADOW VOLUME WITH STENCIL BUFFER

LEE KONG WENG

A thesis submitted in fulfilment of the
requirements for the award of the degree of
Master of Science (Computer Science)

Faculty of Computer Science and Information System
Universiti Teknologi Malaysia

MARCH 2006

To my mum
beloved Jess, family and friends.
Thanks for all the encouragement and loves …

ACKNOWLEDGEMENT

My first thanks goes to Associate Professor Daut Daman, my supervisor for encouragement, guidance, critics and motivation. Without his help this thesis would not be have been the same as presented here, if not almost impossible. Besides that, I would to thank him for making me to write and published paper, which in the end helps me a lot in this research thesis.

The second thanks go to the department of Computer Graphics and Multimedia, UTM staff and lecturer for their contribution towards this research. I'm also indebted to Universiti Teknologi Malaysia (UTM) and Ministry Of Science, Technology and Innovation (MOSTI) for funding my Master's research. Other thanks including to those researchers, academicians and practitioners who made available their research work on the internet that help a lot in my research.

The final thanks go to my beloved Jessica for the encouragement, motivation, critics, arguments and making my heart warm. The next thanks goes to my mum and family for the support and understanding. Thanks to Baldeve, Saandi, Yus and my other lab mate for the discussions which are related or non-related to my research. Not to forget the next lab boys and girls for providing food and entertainment such as network game during the time when my brain seems to be offline (if your supervisor happen to be my evaluator then you all in big trouble), then there are group of course mate who always providing relaxation with badminton games and water recreation at Kota Tinggi. Thanks and may God bless all of you.

ABSTRACT

Shadows are essential to realistic and visually appealing images. The important elements in shadows are the accuracy and the dynamics of hard shadows, which provide scene information and spatial cue. Soft shadows on the other hand provide the realism. The research goal is to create an accurate and real-time dynamic fake soft shadow using stencil shadow volume algorithm. The performance of the algorithm must maintain the quality at acceptable level, as it is the main criteria. In this research, a point light source is used and the occluder is an opaque object. Culling and particle system are excluded, and light source are not allowed to get close to the occluder. The methodology consists of two steps; first is to create the hard shadow using shadow volume with stencil buffer, and second is to add in the soft shadow effects. The first step involves silhouette determination of the occluder, construction of the shadow volume by extruding the silhouette to a very large value, and rendering the shadow volume with depth-pass technique in order to determine which objects are to be shadowed. The second step is to retrieve the original geometry of the hard shadows, and to produce samples of hard shadows with a sequentially enlarged and softened colour. All these samples are finally averaged and blended together to produce the soft shadow effects. The algorithm is implemented using C++ and OpenGL where the occluders are ranging from simple to complex objects. Comparison using proposed algorithm with an established algorithm revealed that the proposed algorithm is two to five times faster. Results from various tests reveal that the proposed algorithm not only achieves better frame per second performance but also maintains the quality of soft shadows. Therefore, the importance of soft shadows cannot be neglected as it increases realism in computer-generated scenes.

ABSTRAK

Bayang-bayang penting kepada imej yang realistik dan menarik. Elemen utama dalam bayang-bayang adalah ketepatan dan kedinamikan bayang-bayang kasar kerana ia memberikan maklumat dan kue spatial. Bayang-bayang lembut pula memberikan kesan realistik. Matlamat penyelidikan ini adalah untuk menghasilkan bayang-bayang lembut palsu dinamik yang tepat dan masa nyata menggunakan algoritma isipadu bayang-bayang stensil. Pencapaian algoritma ini mestilah mengekalkan kualiti pada tahap yang boleh diterima kerana ia adalah kriteria utama. Dalam kajian ini sumber cahaya titik digunakan dan penghadang adalah objek yang tidak lutsinar. Sistem partikel dan penakaian dikecualikan dan sumber cahaya tidak dibenarkan menghampiri penghadang. Metodologi terbahagi kepada dua peringkat; pertama ialah untuk menghasilkan bayang-bayang kasar menggunakan isipadu bayang-bayang dengan penimbal stensil, dan kedua ialah untuk menambah kesan bayang-bayang lembut. Peringkat pertama melibatkan penentuan bayang-bayang bagi penghadang, pembinaan isipadu bayang-bayang dengan memanjangkan bayang-bayang ke satu nilai yang besar, dan menghasilkan isipadu bayang-bayang dengan teknik lulus-kedalaman untuk menentukan objek yang perlu menjana bayang-bayang. Peringkat kedua ialah mendapatkan geometri asal bayang-bayang kasar, dan menghasilkan sampel bayangan kasar dengan pembesaran berjujukan dan warna yang dilembutkan. Semua sampel ini akhirnya dipuratakan dan digabungkan untuk menghasilkan kesan bayang-bayang lembut. Algoritma ini dilaksanakan menggunakan C++ dan OpenGL di mana julat penghadang ialah dari objek yang mudah sehingga kompleks. Perbandingan antara algoritma cadangan dengan algoritma yang telah terbukti menunjukkan bahawa algoritma cadangan ialah dua hingga lima kali lebih laju. Keputusan dari beberapa ujian menunjukkan bahawa algoritma cadangan bukan sahaja mencapai kelajuan kerangka per saat yang lebih baik tetapi juga mengekalkan kualiti bayang-bayang lembut. Oleh itu, kepentingan bayang-bayang lembut tidak boleh dinafikan kerana ia menambah realistik dalam pemandangan janaan komputer.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

.asc	-	ASCII file format
.txt	-	Text file format
2D	-	Two dimensional
3D	-	Three dimensional
CPU	-	Central Processing Unit
FPS	-	Frame per seconds
GB	-	Gigabyte
GHz	-	Giga Hertz
GPU	-	Graphic Processing Unit
MHz	-	Mega Hertz
PC	-	Personal Computer
RAM	-	Random Access Memory
RAMDAC	-	Random Access Memory Digital to Analog Converter

LIST OF APPENDICES

CHAPTER 1

INTRODUCTION

1.1 Introduction

Computer graphics have become one of the most important entities in advertisements, virtual environment applications, simulations and games. Today's users have become smarter and more sensitive with highly sophisticated requirements towards what they view on computer screen, displays or console. This has been a rocket development since ten years back, when users were still indulged in the green screen DOS environment that is nearly wiped out by today's technologies.

Programmers nowadays write computer graphics applications and tools in high-level language such as C and C++ to run on hardware with gigantic memory, speedy access to secondary storage, powerful CPUs, high resolution displays and hardware graphics acceleration. This is however still far from today's user's demand and expectation, which have outstripped the increase in PC performance, forcing programmers to enhance graphics content and features while struggling to maintain high performance.

Computer graphics applications and tools such as virtual environments, simulations and games have utilised very sophisticated and complex algorithms to make them realistic and life-like. Components such as lighting, shading, texture mapping, shadows and others are the main elements in making a scene look realistic.

Hence, this research is aimed at developing an algorithm for generating soft shadows in real-time. Shadows, which play an important role in making a scene look realistic, are essential in games, virtual environment applications and simulation because they represent information for decision-making. The more realistic a shadow is the more informative it is for an analysis.

Shadows are essential to realistic and visually appealing images, but they are difficult to compute in most display environment especially in real-time games (Haller *et al.*, 2003). In computer games, shadows give the gamers the feel of dramatic and spooky effects. A game without the shadow effects is considered to lack realism. The scene tends to look flat and it is difficult to determine the size and spatial relation of the objects in the scene.

Besides being the main ingredient in adding realism to a scene, shadows also provide important spatial cues or information (Gibson *et al.*, 2003). This enables the viewer to determine whether an object is located on the plane or some distance away. Just as with lighting, the increasing level of realism will decrease rendering performance.

1.2 Research Background

The early work on shadows brings us back to 1977 when Frank Crow first published his paper on *Shadow Algorithm's for Computer Graphics* (Crow, 1977). His method explicitly clips shadow geometry to the view frustum, generating perfect caps where the volume crosses a clipping plane. Heidmann (1991) suggested the use of stencil buffer to implement Crow's original algorithm that gave the algorithm the name “stencil shadow volume” by which it is best known today. Stencil shadows belong to the group of volumetric shadow algorithms, as the shadowed volume in the scene is explicit in the algorithm.

In year 2000, Carmack suggested a slightly different approach which entails that the view rays are traced from infinity towards the eye. The tracing ends when encountering the pixel on the geometry that is closest to the eye, which is discussed in (Carmack, 2000). This reversal of the view rays' direction has given the algorithm the name *Carmacks reverse*.

Lengyel (2002) created a hybrid algorithm that uses faster z-pass rendering when the view port is not shadowed and reverts to robust z-fail rendering when the view port is shadowed. Several new improvement methods based upon shadow volume are suggested. Assarson *et al.* (2002, 2003) describe how to create soft shadows using penumbra wedges rendered from shadow volume. Fauerby Kjaer (2003) used a technique for highly efficient coverage calculation for spherical light sources. This technique can avoid clipping operations in the pixel shader and let the texture sampler do the clipping for free. However, this technique has a setback being limited to spherical light sources.

Shadow mapping by Williams (1978) is an alternative, image-based method for shadow determination. It offers an advantage over the shadow volume technique for its performance is not directly dependent on the complexity of the shadowing object. It is essentially a technique involving the transformation of distances from one viewpoint, that of the viewer to distances from another viewpoint, that of the light source. Unlike shadow volume, shadow maps have the ability to shadow any object that can be rendered.

A lot of variants of shadow maps were studied and developed because of the problems in shadow maps such as aliasing and quantisation. The Deep Shadow proposed by Lokovic and Veach from Pixar Animation Studios, is a new and more evolved shadowing technique suited especially for producing fast and high-quality shadows of self-shadowing primitives such as hair, fur and smoke (Lokovic and Veach, 2000). Adaptive Shadow Maps by Fernando *et al.* (2001) which are an extension to the traditional shadow mapping technique, remove aliasing problems by resolving pixel size mismatches between the eye and light source views. An approach to adapt shadow map resolution to the current view is discussed in

(Stamminger and Drettakis, 2002). By using a non-uniform parameterisation, a single *perspective shadow map* can be generated, providing high resolution for near objects and decreased resolution for a viewer who moves away from the viewpoint.

Another interesting technique for shadow generation that has recently been demonstrated by McCool is a hybrid of the shadow volume technique and the shadow map technique. This technique, Shadow Volume Reconstruction by McCool (2000), is based on the observation that once a shadow map has been generated, the information in the shadow map (i.e. the depth, or *z* information for the scene as observed from the light source) defines a volume.

Last but not least is the planar/projection method which, although being inaccurate, is top in its class in terms of performance. An easy-to-implement type of shadow can be created using projection transforms. An object is simply projected onto a plane, and then rendered as a separate primitive. Computing the shadow involves applying an orthographic or perspective projection matrix to the model view transform, then rendering the projected object in the desired shadow colour. Blinn (1988) admits to having first attempted an approach to solve the problem in an unintuitive (but slightly quicker) way. He also presented his revised methodical approach (Blinn, 1996) and showed problems associated with it.

1.3 Current Issues

Currently most real-time shadow algorithms have been limited to *hard shadows*. Hard shadows are the result of light sources being modelled as a single point with no area and can be recognised by a very sharp transition from light into shadow. If light sources are modelled with an actual shape of an area or volume, *soft shadows* can be produced. Soft shadows can be recognised by their inclusion of a penumbra region: an area that is neither fully lit nor fully in shadow.

The visual quality of soft shadows compared to that of hard shadows is very high. It is thus highly desirable to be able to apply real-time soft shadows to computer games. Unfortunately, the computations required for soft shadows are much more complex than those required for hard shadows (Chan and Durand, 2003). Currently, no released game has utilised true, real-time and dynamic soft shadows.

Real world shadows are usually *soft* with a smooth transition from full light to full shadow. This effect is due to the indirect illumination and the volume light source. Light sources within a volume of soft shadows are points where *a part of* the light is visible. This is called the penumbra region. The areas where the geometry is in full shadow with no visible light are called the umbra region.

Currently, there are a few techniques that generate soft shadows. These techniques include Plateaus, Heckbert and Herf, technique by ATI and penumbra wedges in shadow volume. The latest technique is by Kasper and Carsten using fast coverage calculation for spherical light sources. However, unlike those hard shadow algorithms, all these soft shadow algorithms have never been implemented in any game engines or any real-time applications.

The next issue concerns hardware development. Nowadays, graphic cards are so significantly fast in rendering graphics. The two biggest leading companies, NVIDIA Corp and ATI Technologies Inc have produced graphics cards that give end-users graphics performance that once could only be run on supercomputers. Besides giving end-users high graphics performance, they also produce innovative tools for developing better graphics-related applications that maximise the usage of hardware. NVIDIA Corp in close collaboration with Microsoft Corporation has come up with Cg, a high-level C-like graphics programming language. The Cg environment consists of two components: the Cg Toolkit including the NVIDIA Cg Compiler Beta 1.0 optimised for DirectX® and OpenGL®; and the NVIDIA Cg Browser, a prototyping/visualisation environment with a large library of Cg shaders. Meanwhile ATI has developed Radeon SDK 2.0 to rival NVIDIA Cg. Radeon SDK 2.0 also takes the advantage of ATI® products using Microsoft DirectX®. In the next section, more will be discuss on the improvement in hardware particularly by these two mammoth companies.

Another change to the hardware development is the use of *two-sided stencil testing*. This new feature in recent GPU's allows the application programmer to set up and use different stencil states and operations for back facing and front facing triangles respectively. This reduces the CPU load of issuing render calls to the driver, reduces the amount of vertices that have to be transformed through a vertex shader, and reduces the amount of mesh data that has to be sent to the graphics card over the AGP bus. This development has led to significant gain in performance. These features however are only available in current graphic cards. Besides providing *two-sided stencil,* NVIDIA's high-end graphic cards also come with UltraShadow Technology for next generation games like Doom III from id Software in providing real-time cinematic effects of shadow. With the addition of programmable vertex shader functionality, another significant improvement to the algorithm is made possible. The calculation of the shadow mesh can be moved from the CPU to the GPU (McGuire and Hughes, 2004), which of course lifts a burden from the CPU but, more importantly, it also means that the shadow mesh data can be sent to the GPU one shot instead of every frame as was previously necessary.

1.4 Problem Statement

Real-time shadows remain a huge challenge to the developers of computer graphics applications and games. Fifteen frames per second are considered a standard for real-time animation, which gives mere 0.067 second for a developer to construct and render the whole scene (Govindaraju *et al.*, 2003). Most of the real-time shadows techniques reviewed here concentrate on giving realistic and accurate shadows. However, with limited time given, many methods have been developed to improve rendering speed at the expense of reduced quality.

Simple, inaccurate shadows (dark ellipse or rectangle below an object) have been used over the past. This is because some representation of shadows is better than none at all. However, this is not the solution for today's applications that demand more accurate, realistic and dynamic shadows. In 1991, Heidmann implemented accurate dynamic hard shadows that demonstrated the possibility of implementing shadow volume algorithm with stencil buffer. At that time, it only worked on fairly expensive high-end systems. Accurate dynamic shadows became commonly used around 1999 due to graphics hardware improvement and the generalisation of the algorithm presented by Carmack (2000).

However, there were later problems related to the realism of the shadows. The real-world shadows consist of soft shadows, the transitions of which from dark region to no shadow region are smooth. Though soft shadows are important in adding realism to a computer generated scenes but in general a fake soft shadow is enough for the current needs. This is because it does not matter whether shadows are soft or hard as the shadows still represent the same spatial cue and information.

High quality shadows come with a price of reduction in performance. So, to what extent of reduction in quality that can be tolerable and acceptable without compromising performance? The reduction of performance should not slow down the other functions such as rendering, lighting and others. A huge number of point lights are required to avoid obvious sampling artifacts, though this normally is done using graphics hardware.

Many of the algorithms developed are too computationally intensive for real-time purpose. It is assumed that the target system is a PC or similar devices equipped with the common hardware acceleration for graphics. A common method to enhance the performance of an algorithm is to utilise graphics hardware. Triangle rendering can be implemented efficiently using hardware function. With the current states of development in graphics hardware, there is a hope that lot of shadow rendering algorithms can be transferred to GPU's rather than CPU's in the near future.

1.5 Research Goal

The ultimate goal of this research is to develop an algorithm for generating dynamic and realistic shadows using the shadow volume technique. The fake soft shadows generated which could achieve a real-time rate, can eventually be used in games, simulations and virtual environments.

1.6 Research Objectives

[1] To develop a shadow technique that creates an accurate shadow.

[2] To develop algorithm that are able to generate fake soft shadows which are capable of increasing realism in a scene.

[3] To develop a real-time shadow technique for scene rendering while maintaining high quality of shadows and frame-rate.

1.7 Research Scopes

The scopes of research are described as follows:

[1] The light source used in this research is a point light source, which does not have the effect of an ambient light spreading out.

[2] The generation of shadow excludes objects that use particle systems such as hair, grass, smoke and fog. Such objects shadow each other, sometimes itself, and the quantity of the polygons involved are enormous. With this, it is too complex and involves a lot of computation, making real-time implementation impossible.

[3] This research involves occluders that are opaque objects. Light penetrable objects such as semi-transparent and glasses are excluded. Such object will

permit some light to pass through them, which may alter the colour of shadows.

[4] The scenes consist of one light source at a time. Rendering more than one light source will require additional complex calculation and increase CPU usage; this in return will slow down the rendering process.

[5] Culling is not implemented in this research as in real game engine. Only the area of shadow that is viewed by eye/camera is rendered. Shadows are independent of the eye/camera position. But if culling is implemented it will reduce the number of polygons rendered.

[6] The situation where the light source gets very close to the occluder is also not included in this research. This is because shadow volumes eventually collapse into an increasingly marginal volume when the light gets extremely close.

[7] The improvement will focus on algorithm efficiency, not on hardware related enhancement. However, basic standard hardware acceleration function that is available on most major graphics acceleration will be utilised in implementing the algorithm.

[8] This research will only consider the outer penumbra with the assumption that the shadow will always have umbra.

1.8 Research Contributions

The main contribution of this research is the development of a shadow technique that enables the creation of reasonably accurate shadows resembling the shadow caster geometry. The algorithm use in this research is the well known shadow volume, the algorithm generates umbra. The umbra is then use to produce the penumbra. Hence, the soft shadow effects produce is a high quality dynamic soft shadow using existing software library and current graphics hardware.

The research algorithm is suitable for applications where the need to achieve both real-time performance and accurate soft shadow are essential. The research algorithm sacrifice the quality of the soft shadow effects to the extent of tolerate able to achieve real-time performance. With the enhancement of graphics hardware in the future, this algorithm will be able to cater for far more complex scenes.

The algorithm will be suitable for applications that emphasise on speed, realism and accuracy. Stated in the following are the applications that will benefit from this algorithm:

- **Film Industry**
  Special effects in movie require such an algorithm to help create realistic shadows for capturing and sustaining viewers' attention. It adds realism in cartoons and science-fiction movies.

- **Game**
  With the enhancement of scenes, especially those created by computers, the gamers will feel as though they are immersed in the game itself. This will give a feeling of pure satisfaction and real-life experience.

- **Simulation/Virtual Environment**
  The shadow effects provide the users of simulation applications a realistic and real-life experience.

- **Geographic Information System (GIS) Application**
  VIRTUAL GIS that has become an important tool for GIS scientists to make realistic and real-life applications possible with the shadow effects.

1.9 Organisation of Thesis

This thesis consists of six chapters as described:

- **Chapter 1** briefly states the introduction of the research and the research background. The current issues are discussed and problem statements defined. The goal, objectives and the scopes of the research were defined. In addition, the contributions made by this research are also stated.

- **Chapter 2** discusses in details the various categories of techniques available in generating shadows. The pros and cons of the techniques were reviewed, discussed and compared in the literature.

- **Chapter 3** explains the research methodologies that are applied in this research. This chapter shows all the steps taken to develop the research algorithm in detail.

- **Chapter 4** addresses the implementation of the algorithm developed. This includes the implementation of the methodology mentioned in Chapter 3 and the required considerations for the implementation to succeed.

- **Chapter 5** reveals and discusses the results from the implemented prototype and its testing.

- **Chapter 6** wraps up the research with concluding remarks and states the benefits gained from the research. Also stated are suggestions about future improvement work that can be further explored.

The recent developments of graphics hardware bring hopes to graphics applications developers and researchers. This is because most of the future development of soft shadow volume algorithm depends on graphics hardware features and speed. As graphics hardware become more affordable and powerful, there is no reason for developers and researchers not to further develop soft shadow volume algorithm. Future developers or researchers must take advantage of the graphics hardware to develop realistic and real-time soft shadow for graphics applications and games.

REFERENCES

Agrawala, M., Ramamoorthi, R., Heirich, A. and Moll, L. (2000). Efficient Image-Based Methods for Rendering Soft Shadows. *Proceeding of the 27th Annual Conference on Computer Graphics and Interactive Techniques.* July 23-28. New York (USA): ACM Press. 375-384.

Ahokas, T. (2002). Shadow Maps. *Tik-110.500 Seminar on Computer Graphics, HUT TML 2002.* April 23-24.

Aldridge, G. and Woods, E. (2004). Visualisation II: Robust, Geometry-Independent Shadow Volumes. *Proceeding of the 2nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia.* June 15-18. New York (USA): ACM Press. 250-253.

Amanatides, J. and Mitchell, D. (1990). Some Regularization Problems in Ray Tracing. *Proceedings of the Annual Conference on Graphics Interface.* May 14-18. Wellesley (USA): A K Peters, Ltd. 221-228.

Assarsson, U., Dougherty, M., Mounier, M. and Moller, T. A. (2003). An Optimized Soft Shadow Volume Algorithm with Real-time Performance. July 26-27. *Proceedings of the 2003 ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware.* Switzerland: Eurographics Association. 33-40.

Assarsson, U. and Moller, T. A (2003). A Geometry-Based Soft Shadow Volume Algorithm using Graphics Hardware. *Proc. ACM Transactions on Graphics (TOG).* 22(3): 511-520.

ATI Technologies Inc. (2000). *Multitexture Soft Shadows,* ATI Tecnologies Inc: ATI Technical Report.

Blinn, J. (1988). Me and My (Fake) Shadow. In: CG&A. *Computer Graphics and Applications.* Washington (USA): IEEE. 8(1): 82-86.

Blinn, J. (1996). *Jim Blinn's Corner: A Trip Down the Graphics Pipeline*, San Francisco (USA): Morgan Kaufmann Publishers, Inc.

Carmack, J. (2000). Email to private list. Unpublished.

Chan, E. and Durand, F. (2003). Rendering Fake Soft Shadows with Smoothies. *Proceeding of the 14th Eurographics Workshop on Rendering.* June 25-27. Switzerland: Eurographics Association. 208-218.

Chattopadhyay, S. and Fujimoto, A. (1987). Bidirectional Ray Tracing. *Proceeding of CG International '87 on Computer Graphics.* New York (USA): Springer-Verlag. 335-343.

Chin, N. and Feiner, S. (1989). Near Real-Time Shadow Generation Using BSP Trees. *Proceeding of the 16th Annual Conference on Computer Graphics and Interactive Techniques.* July 31-August 4. New York (USA): ACM Press. 99-106.

Chin, N. and Feiner, S. (1992). Fast Object-Precision Shadow Generation for Area Light Source. *Proc. of the 1992 Symposium on Interactive 3D Graphics.* 21-30.

Chrysanthou, Y. and Slater, M. (1995). Shadow Volume BSP Trees for Computation of Shadows in Dynamic Scenes. *Proceeding of the 1995 Symposium on Interactive 3D Graphics.* April 9-12. New York (USA): ACM Press. 45-50.

Cohen, M. and Greenberg, D. (1985). The Hemi-Cube: A Radiosity Solution for Complex Environments. *Proceeding of the 12th Annual Conference on Computer Graphics and Interactive Techniques*. July 22-26. New York (USA): ACM Press. 19(19): 31-40.

Cook, R., Porter, T. and Carpenter, L. (1984). Distributed Ray Tracing. *Proceeding of the 11th Annual Conference on Computer Graphics and Interactive Techniques*. July 23-27. New York (USA): ACM Press. 18(3): 137-145.

Cook, R., Reeves, W. and Salesin, D. (1987). Rendering Antialiased Shadows with Depth Maps. *Proceeding of the 14th Annual Conference on Computer Graphics and Interactive Techniques.* July 27-31. New York (USA): ACM Press. 21(4): 283-291.

Crow, F. C. (1977). Shadow Algorithms for Computer Graphics. *Proceeding of the 4th Annual Conference on Computer Graphics and Interactive Techniques.* July 20-22. New York (USA): ACM Press. 11(3): 242-248.

Dachsbacher, C. and Stamminger, M. (2003). Translucent Shadow Maps. *Proceeding of the 14th Eurographics Workshop on Rendering.* June 25-27. Switzerland: Eurographics Association. 197-201.

Drettakis, G. and Fiume, E. (1994). A Fast Shadow Algorithm for Area Light Sources Using Backprojection. *Proceeding of the 21st Annual Conference on*

Computer Graphics and Interactive Techniques. July 24-29. New York (USA): ACM Press. 223-230.

Everitt, C. and Kilgard, M. J. (2002). Practical and Robust Stenciled Shadow Volume for Hardware Accelerated Rendering. NVIDIA® Corporation: NVIDIA Technical Report.

Fauerby, K. and Kjaer, C. (2003). *Real-time Soft Shadows in a Game Engine*. University of Aarhus: Master's Thesis.

Fernando, R., Fernandez, S., Bala, K. and Greenberg, D. (2001). Adaptive Shadow Maps. *Proceeding of the 28th Annual Conference on Computer Graphics and Interactive Techniques.* August 12-17. New York (USA): ACM Press. 387-390.

Gibson, S., Cook, J., Howard, T. and Hubbold, R. (2003). Rapid Shadow Generation in Real-World Lighting Environments. *Proceeding of the 14th Eurographics Workshop on Rendering.* June 25-27. Switzerland: Eurographics Association. 219-229.

Govindaraju, N. K., Llyod, B, Yoon, S. E., Sud, E. and Manocha, D. (2003). Interactive Shadow Generation in Complex Environment. *Proc. ACM Transactions on Graphics (TOG).* 22(3): 501-510.

Haines, E. (2001). Soft Planar Shadows Using Plateaus. *Journal of Graphics Tools*. 6(1): 19-27.

Haller, M., Drab, S. and Hartmann, W. (2003). A Real-Time Shadow Approach for An Augmented Reality Application Using Shadow Volumes. *Proceeding of the ACM Symposium on Virtual Reality Software and Technology.* October 1-3. New York (USA): ACM Press. 56-65.

Hasenfratz, J. M., Lapierre, M., Holzschuch, N. and Sillion, F.X. (2003). A Survey of Real-time Soft Shadows Algorithm. *Proceeding of the EUROGRAPHICS 2003 Annual Conference*. 1-6 September 2003. Granada (Spain): Eurographics Association. State-of-the-Art Reports 1 (STAR1).

Heckbert, P. S. and Herf, M. (1997). *Simulating Soft Shadows with Graphics Hardware*. Carnegie Mellon University: Technical Report CMU-CS-97-104.

Heidmann, T. (1991). Real Shadows, Real Time. *Iris Universe, Silicon Graphics Inc*. 18: 23-31.

Herf, M. (1997). *Efficient Generation of Soft Shadow Texture.* Carnegie Mellon University: Technical Report CMU-CS-97-138.

Nguyen H, Huu. (1999). Casting Shadows on Volumes. In: Game Developer. *Game Developer*. San Francisco (USA): United Business Media. 6(3): 44-53.

Hun Y, Kwoon. (2002). *The Theory of Stencil Shadow Volumes*. GameDev.net.

Lengyel, E. (2002). *The Mechanics of Robust Stencil Shadows,* Gamasutra.com.

Lokovic, T. and Veach, E. (2000). *Deep Shadow Maps*. *Proceeding of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. July 23-28. New York (USA): ACM Press/ Addison-Wesley. 385-392.

Meaney, D. and O'Sullivan C. (1999). *Heuristical Real-time Shadows.* Trinity College Dublin: Technical Report TCD-CS-1999-19.

McCool, M. (2000). Shadow Volume Reconstruction from Depth Maps. *Proc. ACM Transactions on Graphics (TOG).* 19(1): 1-26.

McGuire, M., Hughes, J. F., Egan, K. T, Kilgard, J. M. and Everitt, C. (2003). Fast, Practical and Robust Shadows. NVIDIA® Corporation: NVIDIA Technical Report.

McGuire, M. and Hughes, J. F. (2004). Hardware-Determined Feature Edges. *Proceeding of the 3rd International Symposium on Non-Photorealistic Animation and Rendering.* June 7-9. New York (USA): ACM Press. 35-147.

Moeller, T. A and Haines, E. (2002). *Real-Time Rendering*. 2nd ed. Wellesley (USA): A K Peters, Ltd.

Moller, T. A. and Assarsson, U. (2002). Approximate Soft Shadows on Arbitrary Surfaces using Penumbra Wedges. June 26-28. *Proceedings of the 13th Eurographics workshop on Rendering.* Switzerland: Eurographics Association. 297-306.

Ng, R., Ramamoorthi, R and Hanrahan, P. (2003). All-Frequency Shadows Using Non-Linear Wavelet Lighting Approximation. *Proc. ACM Transactions on Graphics (TOG).* 22(3): 376-381.

Nishita, T. and Nakamae, E. (1985). Continuous Tone Representation of Three Dimensional Objects Taking Account of Shadows and Interreflection. *Proceeding of the 12th Annual Conference on Computer Graphics and Interactive Techniques*. July 22-26. New York (USA): ACM Press. 19(3): 23-30.

Pop, M., Duncan, C., Barequet, G., Goodrich, M., Huang, W. and Kumar, S. (2001). Efficient Perspective-Accurate Silhouette Computation and Applications. *Proceeding of the 17th Annual Symposium on Computational Geometry.* June 3-5. New York (USA): ACM Press. 60-68.

Sen, P., Cammarano, M. and Hanrahan, P. (2003). Shadow Silhouette Maps. *Proc. ACM Transactions on Graphics (TOG).* 22(3): 521-526.

Soler, C. and Sillion, F. X. (1988). Fast Calculation of Soft Shadow Textures Using Convolution. *Proceeding of the 25th Annual Conference on Computer Graphics and Interactive Techniques.* July 19-24. New York (USA): ACM Press. 321-332.

Stamminger, M. and Drettakis, G. (2002). Perspective Shadow Maps. *Proceeding of the 29th Annual Conference on Computer Graphics and Interactive Techniques.* July 21-26. New York (USA): ACM Press. 557-562.

Whitted, T. (1980). An Illumintaion Model for Shaded Display. *Proceeding of the 6th Annual Conference on Computer Graphics and Interactive Techniques*. August 8-10. New York: ACM Press. 23(6): 343-349.

Williams, L. (1978) Casting Curved Shadows on Curved Surfaces. *Proceeding of the 5th Annual Conference on Computer Graphics and Interactive Techniques*. August 23-25. New York (USA): ACM Press. 12(3): 270-274.

Woo, A., Poulin, P. and Fournier, A. (1990). A Survey of Shadow Algorithm. *Proc. IEEE Computer Graphics and Applications*. 10(6): 13-32.

Wymen, C. and Hansen, C. (2003). Penumbra Maps: Approximate Soft Shadows in Real-Time. *Proceeding of the 14th Eurographics Workshop on Rendering.* June 25-27. Switzerland: Eurographics Association. 202-207.

Yonas, A. (1979). Attached and Cast Shadows. In: Nodine, C. F. and Fisher, D. F (Eds.). *In Perception and Pictural Representation*. New York (USA): Praeger. 100-109.