UNIVERSITY OF ILLINOIS AT CHICAGO

Rendering Shadows

Lorenzo Di Tucci
Ettore Trainiti
Jed Siripun

Definition

Shadows are very important element if you want to create realistic images. The importance of shadows in also stated by Wagner in “The effect of shadow quality on the perception of spatial relationships in computer generated imagery”. Here, the author shows that is better have an inaccurate shadow than none at all. This, because the eye is fairly forgiving about the shape of the shadow.
The terminology related to shadows is as follow. An *occluder* is an object that cast a shadow into a *receiver*. The occluder is hit from a light, the *light source*, in other words, the point of origin of the light. The light source can be a *point light* or a *area/volume light source*. In the case of a point light, the shadow that we generate is an *hard shadow*, otherwise it is a *soft shadow*.
So, we can start with the definition of shadow. A **Shadow** is an area where the light, coming from a light source, is completely or partially obstructed by an object, the occluder. In Computer Graphics, the process of drawing shadows onto a scene on the screen of the computer is called **Shadow Rendering**.

Motivations

Rendering shadows helps to improve how realistic an image is.


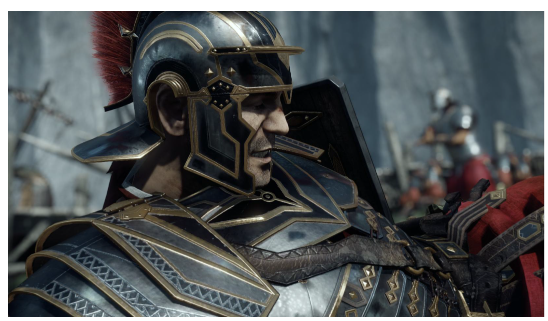

Between the two images, the soldier with the shadows looks much more realistic. The shadows help to hide the details that are not completely touched up. Without shadows, it gives him a well lit face despite his helmet which should be blocking the sun.
It also gives the placement of the object within the frame as well as gives a relative shape and size to an object.

Between these three robots, the shadow helps to indicate how far up from the ground that the robot is hanging. The first one infers that the robot is touching the ground with its feet, the second hanging slightly above the ground, and the third hanging high into the air.

Shadows are also used to identify surface characteristics and details regarding the light source used. As instance, in the pictures below, we can notice that we can't infer about the geometry of the surface of the first picture but, thanks to the shadow of the robot in the second figure, we can notice that the surface is wavy.

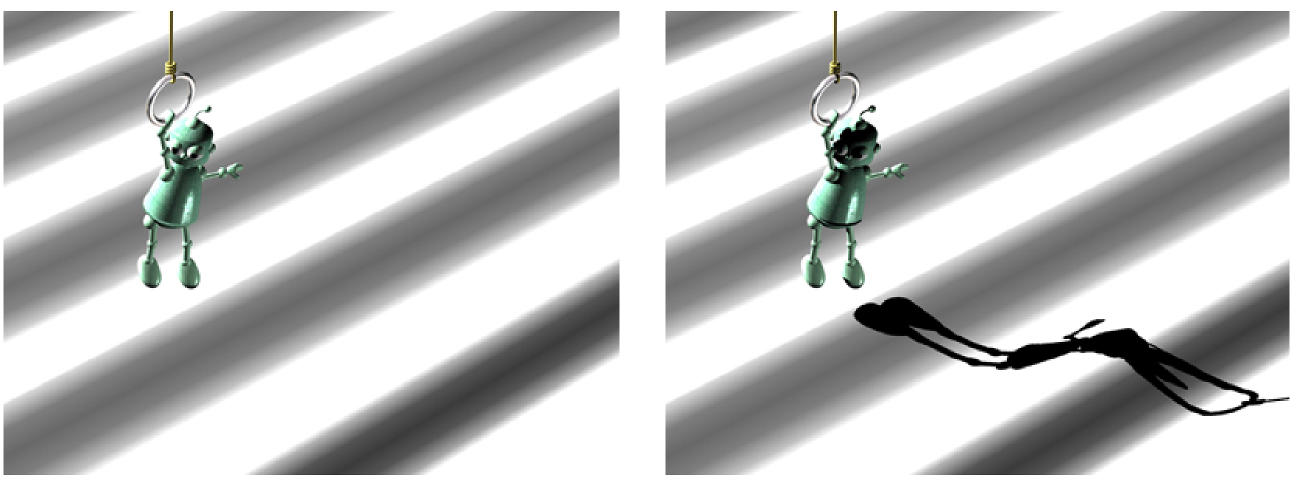

Shadows, indeed, provide valuable visual informations.

History

The history of the shadow volume technique began with Frank Crow in 1977 however, it did not become popular until the game Doom 3 was released in 2005. The use of this shadow volume technique within Doom 3 became known as Carmack's Reverse.

Shadow mapping was invented a year later than shadow volume in 1978. Lance Williams created the use of shadow comparisons through the depth buffer. Shadow mapping is currently the more popular option within gaming today as it supports soft shadows.

Projection shadows is an older technique that can be used for shadow rendering. It is the process of projecting the shadow of the object upon the plane/floor. This technique is only practical when the shadow is being casted upon a flat surface. Therefore this technique is outdated in comparison to the other two techniques cited before.

Composition of Shadow

Shadow is composed of three different parts, **umbra, penumbra** and **antumbra.** All this parts should be drawn in order the let the object be more realistic.

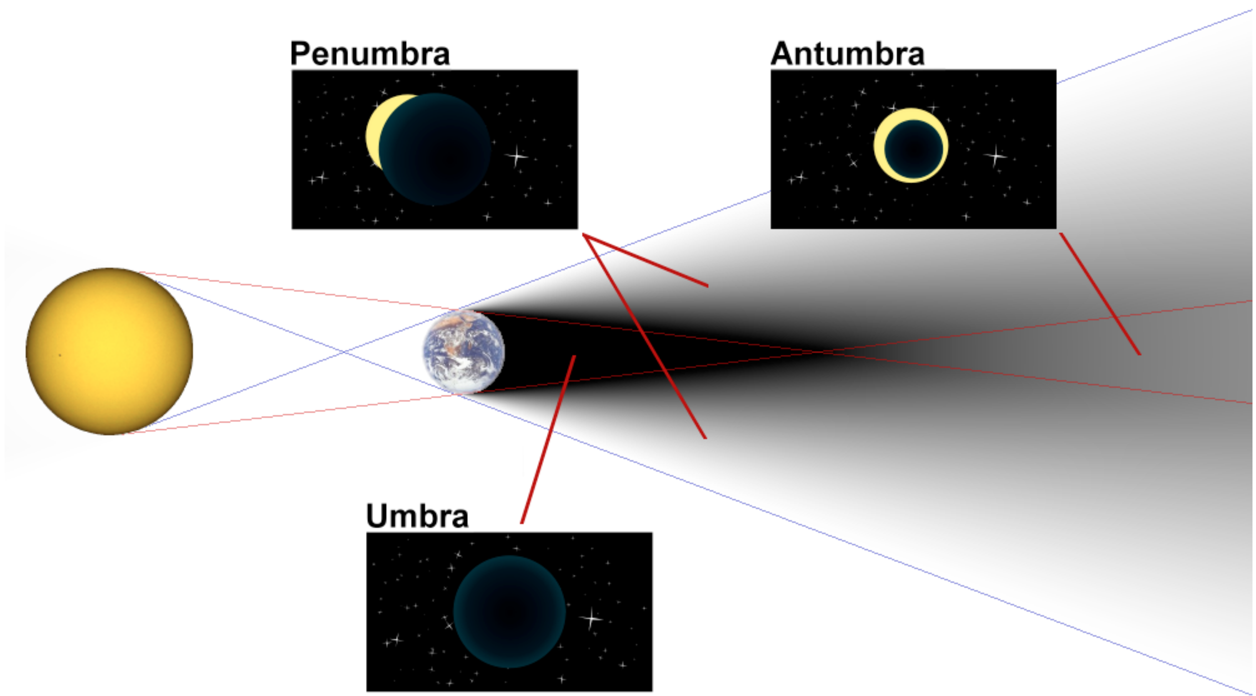

If we consider an environment where we have a light that hits an object, as in the picture above, we can define the umbra as the region where the light is **completely** blocked by the occluding body. This is the darkest shadow, as no light hits that region. A user, that place itself inside this area, experience a total eclipse.
Then, we can observe that, at the sides of the umbra, there is a lighter region. This region is called Penumbra. It is characterized by the fact that only a **portion** of the light is occluded.

So, some light passes and highlights the area. This is also called partial eclipse, as an observer in that region would partially see the light.
The last part of a shadow is called Antumbra. This region, is where the occluding body appears completely contained within the disk of the light source. Here, a user can experience an annular eclipse. If we move from the antumbra region, directly into the umbra region, the apparent size of the occluding body would increase until we would arrive to the umbra region.

Computing all the elements of a shadow is computationally expensive and, as pointed out in the paper *"Between Umbra and Penumbra"* compute realistics shadows efficiently is a difficult problem, particularly in the case of non-point light sources.
In the paper said above, they also tried to estimate a complexity of boundaries between umbra, penumbra and regions that are completely enlightened. Their experiments has been conducted on a plane in a polyhedral scene consisting of k convex objects of total complexity n.
The results they achieved are that, if we reduce the light to a line, the umbra may have $\Omega(nk^2 + k^4)$ connected components and $O(nk^3)$ complexity. For what concerns a polygonal light source they discovered that it can generate an umbra with $\Omega(n^2k^3 + nk^5)$ components and with complexity $O(n^2k^3)$.
All this, to state that umbra is much more intrinsic than the boundaries between penumbra and areas completely enlightened. In fact, penumbra has a lower complexity.

Hard & Soft Shadows

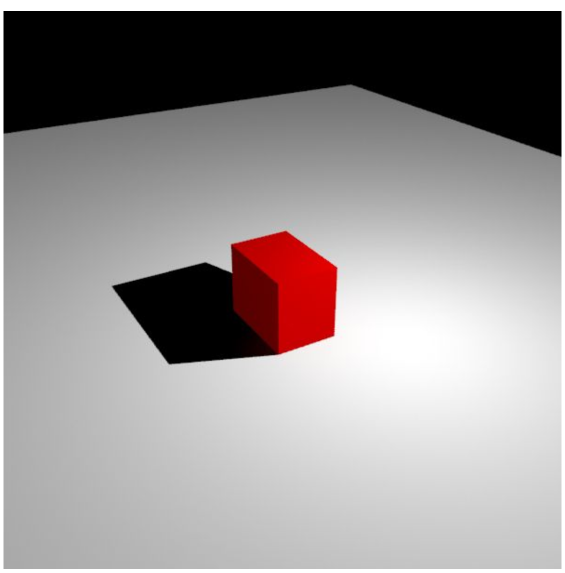

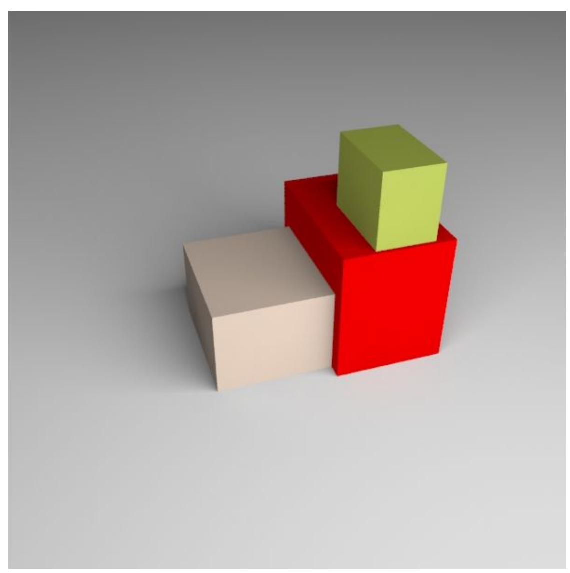

Hard shadows are shadows with defined edges. They are made from basic ray traced shadows. A ray is shot from the camera, hits the surface, and sends a shadow ray out to see if the light reached a certain point. If the shadow ray is obstructed, then it is within a shadow.

This kind of shadow is the easiest to compute but, as drawback is the less realistic.

Soft shadows are shadows that have softer edges that look more realistic. The light source is treated as a physical object making the shadow rays look more like cones than a straight line. As we can observe from the pictures below, hard shadows has only the umbra component while, considering the picture on the right, we can notice that soft shadows are composed of both umbra and penumbra and, for this reason, they look more realistic.

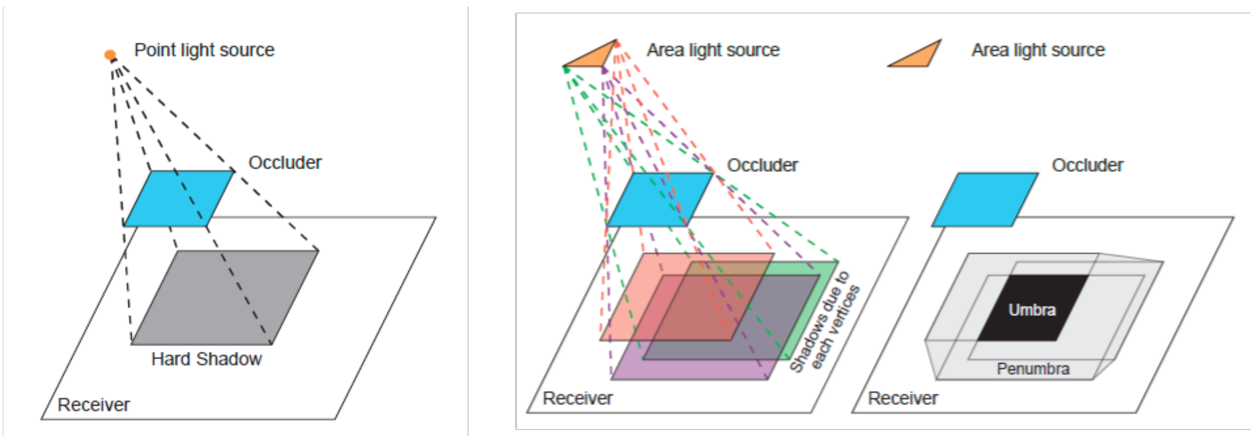

Soft Shadow are the more realistic than hard shadow, but they are computationally expensive and more complex to shape and render.

Techniques

In this section we are going to describe the two most commonly used techniques to render shadows in real time.

Shadow Mapping

Shadow Mapping is an Image-Based real-time shadow rendering technique. The basic way to compute shadows is by identifying which parts of the scene are hidden from the light source. Shadow Mapping achieves this result through the determination of visible surfaces from the light's point of view.
The rendering process starts by computing a view of the scene from the point of view of the light source. The z values from the z-buffer of the rendered scene are stored in an image referred as the Shadow Map.

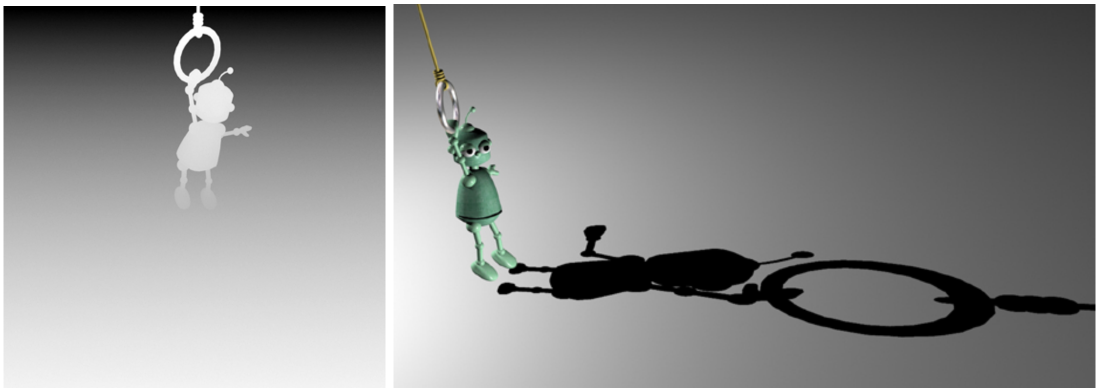

For each pixel in the scene we have its geometrical position. If the depth of this position, transposed to light's coordinates, is lower than the one stored in the Shadow Map the pixel is in shadow otherwise it is illuminated. The color of the objects is then modulated depending on whether they are in shadow or not.
Shadow mapping generates large aliasing problems if the light source is distant from the viewer. In this case the individual pixels from the shadow map are visible, resulting in the aliasing effect along the shadow outline. Several methods has been implemented to solve this problem and it is an active research topic.
For example, during the last i3D conference in San Francisco a technique called Sub-Pixel Shadow Maps (SPSM) has been proposed as an effective method that avoids both perspective and projection Aliasing at the cost of a small computational overhead. SPSM is based on the storage of a fixed-size partial representation of the scene geometry using conservative rasterization, combined with an original reconstruction of shadow edges.

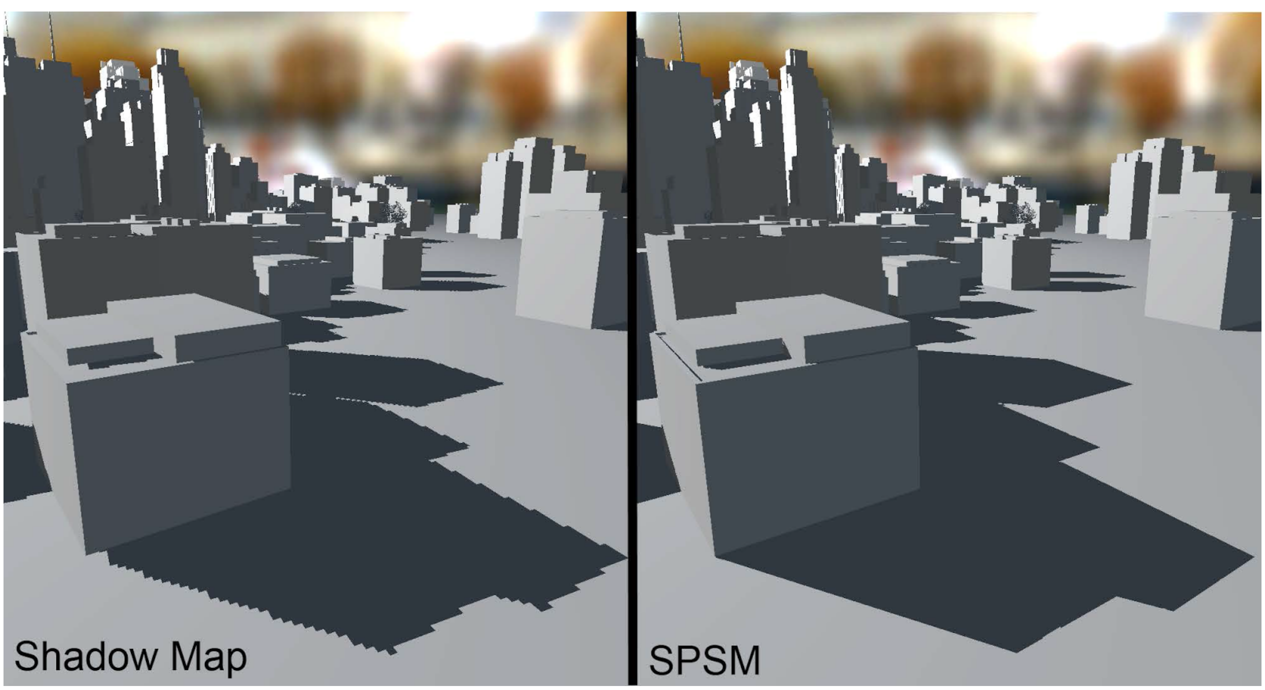

Shadow Mapping in summary :

Advantages	Disadvantages
Can be implemented entirely using graphics hardware	Subject to many sampling and aliasing problems
Creating the shadow map is relatively fast	Cannot handle omni-directional lights
Handles self-shadowing	At least two rendering passes required

Shadow Volume

Shadow volume is a Geometry-Based real-time shadow rendering technique.
The main idea behind Shadow volumes is to create volumes that “represent” darkness. The algorithm consists in finding the outline of objects occluding the light, extruding the outline to infinity along the light direction thus forming shadow volumes.

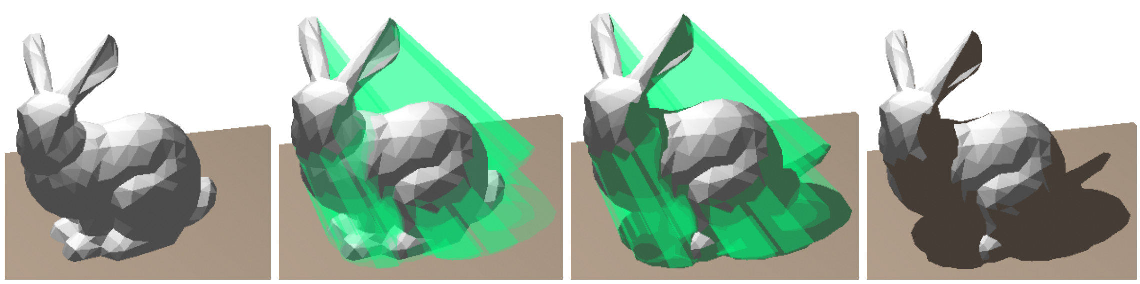

The shadow volume is calculated in two steps.
The first step consists in finding the silhouette of the occluder as viewed from the light source. Then we construct the shadow volume by extruding these edges along the direction of the light source. For each pixel in the image rendered, we count the number of faces of the shadow volume that we are crossing between the point and the object rendered.

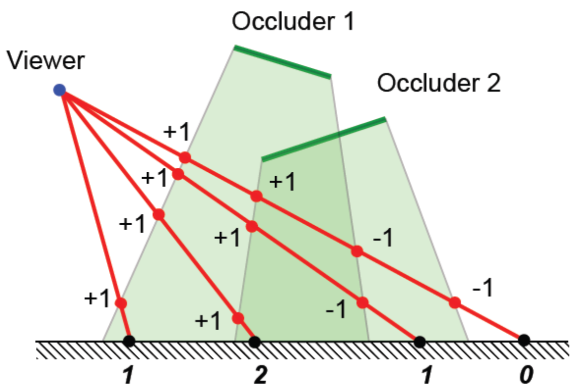

Front facing faces of the shadow volume increment the count, back-facing faces decrement the count. If the total number of faces is positive, then we are inside the shadow volume hence the pixel is in shadow.

Advantages	Disadvantages
Works for omnidirectional Light sources	Computation time depends on objects complexity
Renders Eye-View Pixel precision shadows	Requires expensive computation of the occluder outline
handles self-shadowing	Rendering the shadow volume consumes the fillrate of the graphics card.

Soft Shadow Image based algorithms

Image based soft shadowing algorithms are based on the Shadow mapping technique previously described. Some methods commonly used consist in :

- Combination of several point based shadow images
- Layered attenuation Maps
- Quantitative information in the shadow map
- Single Sample Soft Shadows
- Convolution Technique


The images above show a multi sampled soft shadow. On the left a 4 sample soft shadow, on the right a 1024 sample soft shadow.

Soft Shadow Geometry based algorithms

Soft shadow Geometry based algorithms build upon the Shadow Volume rendering technique. Several methods can be used, like :

- Combining together several shadow volumes taken from point samples on the light source, similarly as done in image based soft shadowing
- Extending the shadow volume using a specific heuristic
- Computing a penumbra volume for each edge of the shadow outline

Comparison Table

The following comparison table, taken from “Hasenfratz et al.” ‘s paper briefly summarizes and compares the two approaches :

Method	Time	Quality	Tunable	Light	Scene	Required Hardware
Image-based						
Multi-samples[22, 25]	I	*	Y	Polygon	1 planar receiver	
Distributed Multi-samples[28]	RT	**	Y	Planar		ShadowMap
Single sample[9, 33]	RT	*	Y	Sphere		ShadowMap
Convolution[45]	I	**	Y	Polygon		2D Convol.
Visibility Channel[24, 54]	I	**	Y	Linear, Polygon		2D Convol.
Geometry-based						
Plateaus[19]	I	**	Y	Sphere	1 planar receiver	
Penumbra Map[53]	RT	**	Y	Sphere		Vertex & Frag. Programs
Smoothie[11]	RT	**	Y	Sphere		Vertex & Frag. Programs
Soft Shadow Volumes[2, 4, 5]	RT	***	Y	Sphere, Rect.		Fragment Programs

RT Stands for “Real Time” while I stands for “Interactive”.

Challenges and Directions

One of the biggest problem in rendering shadows is that they are computationally expensive. In fact, it is difficult to efficiently render precise shadows. Nowadays, thanks to more powerful graphic cards, it is possible to draw more precise shadows, and the best algorithms for scene with moderate complexity are the ones based on shadow volumes. The problem with this kind of rendering, is that the volumes generated have a very high complexity. One of the solutions, that has been found is the one explained in the paper “ Per-Triangle shadow volumes using a view-sample Cluster Hierarchy” . In this paper, they recall a method that try to solve the problem of the volume complexity by testing each individual triangle shadow volume against a hierarchical depth map. In this way, the algorithm allows volumes that are in front of, or behind, the rendered view samples to be efficiently culled. In the paper, they optimize the algorithm by building a full 3D acceleration structure over the view samples and testing per-triangle shadow volumes against that.

Another issue in shadow rendering is real time performance for applications using many lights. Examples of application like this are games. In the paper “Efficient Virtual Shadow Maps for Many Lights”, the authors explored the use of hardware-supported virtual cube-map shadows to efficiently implement high quality shadows from hundreds of light sources in real time and within a bounded memory footprint.
Actually, in recent years, many applications that enable real-time performance for games using lots of lights have been presented, but none of them was drawing shadows.
The conclusion of the paper cited above, states that, using hardware supported virtual cube shadow maps is a viable method if we want to achieve high quality real time shadow, with the

possibility to use hundreds of lights. They considered also memory requirements and, by using this method, it remains proportional to the number of shaded example.
The problem of real time multi-lights rendering is moving in the way of exploring different light distribution and combining virtual shadow maps with cascaded, or parallel split, shadow maps.

Aliasing is another big issue when it comes to shadow rendering and is due to a regular structure of the dots that create an image. It occurs most frequently along the silhouette of the shadow, small objects, and complicated maps. According to Kolivand and Sunar's Anti-Aliasing article, when using shadow mapping, aliasing occurs because of the mismatch of pixel size of the view point and the light point. In order to fix this problem, solutions are being researched such as sampling, filtering, and partitioning and each of these solutions has sub-solutions within them. There are different types of samplings that can be done such as super-sampling, stochastic sampling, and distorted shadow maps. Below is an example of the use of the distorted shadow maps by use of a redistributive mechanism in order to enlarge the silhouette of the shadow by making the parts that are not inside the penumbra regions. The left is before use and the right is after use.

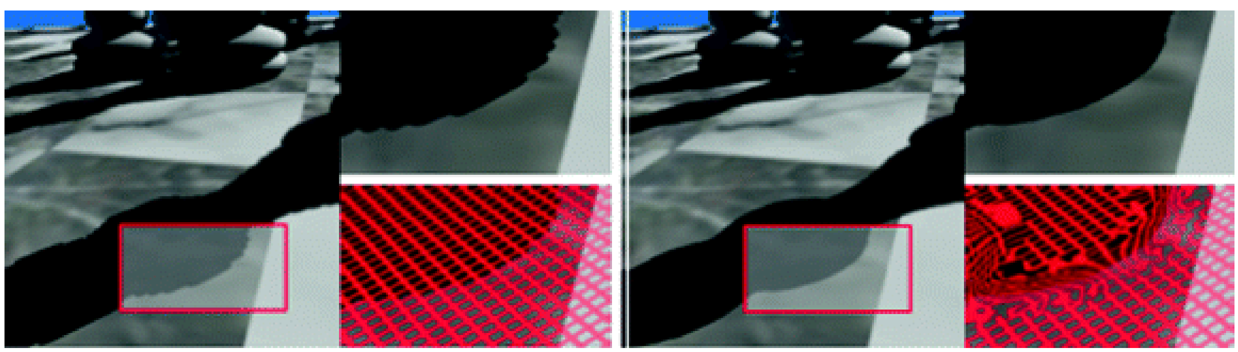

Another technique used for improving shadow quality is filtering. Filtering is typically used for soft shadows or semi-soft shadows. The difference between the two is that semi-soft shadows are better for outdoor rendering where there is typically only one light source while soft shadows are better for indoor rendering where there is usually multiple light sources.

Partitioning is the most practical approach for anti-aliasing, especially for larger scenes. It is the use of algorithms to subdivide the frustum along the z-axis and changing properties of each subdivision to reduce the time it takes to render the scene as a whole. An improvement to z-partitioning is to reduce the number of renderings to the number of partitions. Another improvement made was to find the exact position that was best for each split plane. That equation is $Ci=zn(zf/zn)(i/m)$ where zn and zf are the near and far plane distances in the view frustum and m is the number of partitions. Below is an example of the z-partitioning that captures four different shadow maps based upon the partition. The right hand side is the result of the four partitioned shadow map.

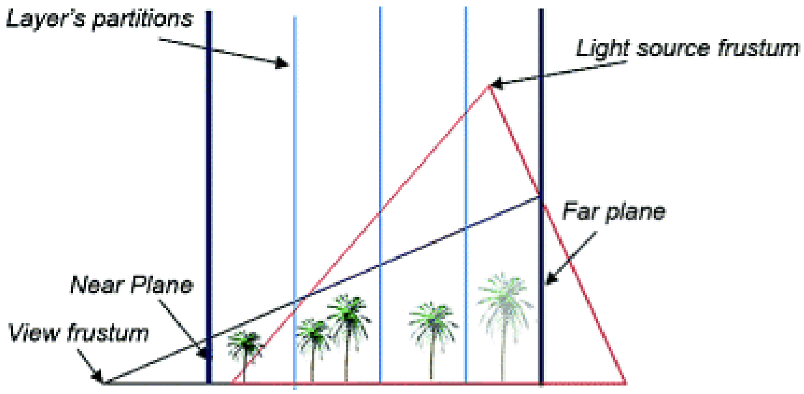

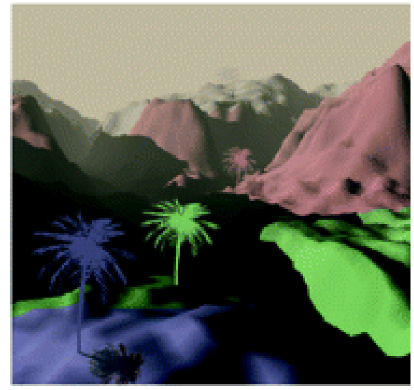

Based upon “Robust Real-Time Shadows for Dynamic 3D Scenes on the Web” by Eicke, T. , Jung, Y. , and Kuijper, A. one of the directions that shadow rendering is taking is looking more into the use of shadow rendering 3D objects within web browsers. The open-source JavaScript framework X3DOM is an important aspect to the development of the web-based 3D because it provides a way to integrate 3D into HTML5.

Another upcoming research topic within shadow rendering is layered reflective shadow maps. It is voxel-based indirect light that supports dynamic lights and scenes with memory bound consumption. A voxel is the cross word between volume and pixel. It represents a value on a 3D grid. In the case of the picture below, one voxel is one box. The shaded cube is considered to be one voxel. The idea of this shadow rendering is to interactively simulate both diffuse and glossy single-bounce indirect illumination, according to the article “Layered Reflective Shadow Maps for Voxel-based Indirect Illumination”. Below are some examples of the layered reflective shadow mapping.

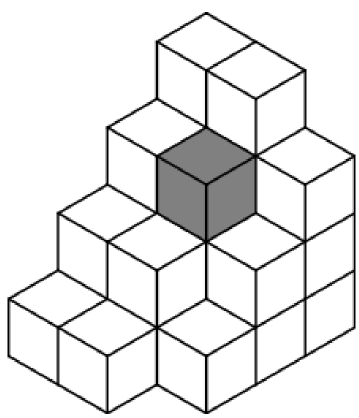

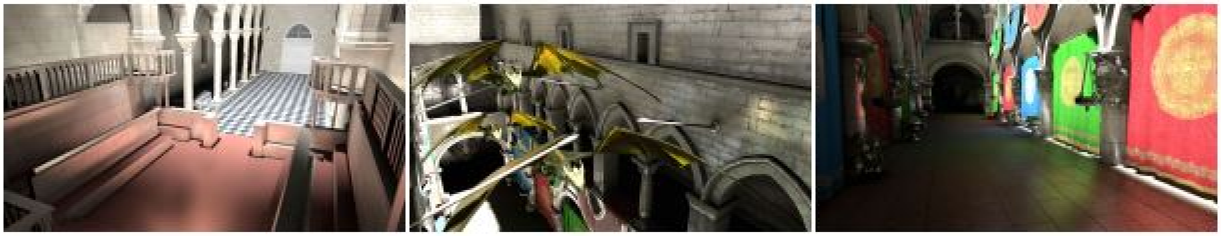

Conclusions

From this report it is evident that rendering shadows in the past has been limited by the computational power of graphics cards. Nowadays, thanks to the constants efforts of the research community, shadows are included in a variety of applications. Faster GPUs and up to date shadow rendering algorithms made shadows a pleasant reality. Researchers keep on finding new ways to address and solve the problems addressed in this report to achieve higher quality shadows and lower computational complexity. It is then evident that rendering shadows is a key point to make computer graphics look more realistic.

References

Woo A, Poulin P, Fournier A. A survey of shadow algorithms.
IEEE Computer Graphics and Applications, 1990.

Hasenfratz J. M. , Lapierre M. , Holzschuch N. , et al. A survey of real-time soft shadows algorithms. Computer Graphics Forum, 2003.

Liu N. , Pang M. , A survey of Shadow Rendering Algorithms : Projection Shadows and Shadow Volumes. Computer Science and Engineering, 2009. WCSE '09.

Kasyan N. , Wimmer M., Schwarz M. , Assarsson U., Eisemann E. .Real Time Shadows Siggraph 2013 Course.

Cutler B. , NordHauser R. , CSCI 4530/6530 Advanced Computer Graphics. Rensselaer Polytechnic Institute, Spring 2014

Lecocq P. , Marvie J. , Sourimant G. , Gautron P. . Sub-Pixel Shadow Mapping. i3D 2014

Ollson O., Sintorn E., Kampe V., Billeter M., Arssarsson U,. Efficient Virtual Shadow Maps for Many Lights i3D 2014

Sintorn E., Kampe V., Ollson O., Arssarsson U,. Per-Triangle Shadow Volumes Using a View-Sample Cluster Hierarchy

Kolivand, H. , Sunar M. S., Anti-aliasing in image based shadow generation techniques: a comprehensive survey, April 2014

Redelinghuys F., Argyriou V., Petrou M., Cast shadows estimation and syinthesis using the Walsh transform

Eicke, T. , Jung, Y. , Kuijper, A, Robust Real-Time Shadows for Dynamic 3D Scenes on the Web, 2014

Sugihara, M. , Rauwendaal, R. , Salvi, M., Layered Reflective Shadow Maps for Voxel-based Indirect Illumination, 2014

Akenine-Moller T., Haines E., Hoffman N., Real-Time Rendering - Third Edition - CRC Press

Wanger L. The effect of shadow quality on the perception of spatial relationships in computer generated imagery. i3D 1992

Demouth J., Devillers O.,Everett, H. , Glisse M. ,Lazard S. , Seidel R. Between umbra and penumbra. SOGC2007