

Improving Rendering Performance by Texture-Map-Based Triangle Strips

Yu Yang, Tulika Mitra and Huang Zhiyong

(Department of Computer Science, School of Computing, National University of Singapore, Singapore 117543)

Abstract

Improving the rendering performance is a basic problem for computer graphics system. In this paper, we are aiming to investigate the impact on the rendering performance of some geometry compression methods. These methods are devised to optimize the use of the vertex cache. We will study how they interact with the on-chip texture cache. Based on the study, one simple method of improving rendering performance by texture-map-based triangle strips is proposed to balance the use of the vertex and texture caches. We have conducted the experiments to show the effectiveness of this method.

Keywords: rendering performance; geometry compression; texture mapping; caching

中图法分类号 TP391

1. Introduction

The development of computer graphics applications is quickly increasing and it is common to find systems which will result in complex models of millions polygons with texture mapping. In many cases, the full detailed geometry and texture maps must be sent down to graphics hardware for rendering. As the users demand ever larger and more realistic 3D models, the transmission time, rendering time and storage requirements grow rapidly. Thus, real-time graphics hardware is increasingly facing a memory bus bandwidth bottleneck in which a large amount of data cannot be sent fast enough to the graphics pipeline for rendering [13].

Limited memory bandwidth is a barrier to increasing PC graphics performance. The memory interface gets inundated with multiple, continuous, high bandwidth demands such as pixel writes, pixel reads, display refresh, AGP bus transactions, and texture reads. Unfortunately, end users notice a slowdown in graphics performance when one of their multiple demands gets bottlenecked by the memory interface.

In this paper, we are aiming to investigate the impact on the rendering performance of some geometry compression methods that are devised to optimize the use of the vertex cache. We will study how they interact with the on-chip texture cache [4, 8]. Based on the study, one simple method of improving rendering by texture-map-based triangle strips is proposed to balance the use of the vertex and texture caches. We have conducted the experiments to show the effectiveness of this method.

2. Background and Related Work

To reduce the bottleneck effect which appears as an obstacle to the increasing needs for fast rendering, the rendering process must be carefully examined. The

traditional OpenGL polygon-rendering pipeline consists of geometry processing, rasterization and image composition [7]. In the geometry processing stage, input data will go through transformation, shading, primitive assembly, visibility culling and projection. In this stage, the information needs to be processed is the geometry data, which includes the vertex position, face information, and normal vectors. In the rasterization stage, the raster images (e.g. texture, bump and environment maps) will be transmitted from system memory to graphics processor, with the use of a texture image cache to reduce the texture image bandwidth. In the image composition stage, it will process the z-buffer of each pixel and then write them to memory.

2.1. Geometry Compression

One solution for reducing the bandwidth is to compress the static geometry as an offline pre-process [5]. This strategy exploits the idea of using vertex cache. By using a relative small size of cache to hold frequently referenced vertices, savings are generated because to render the same object, fewer vertices are needed to pass through to the graphics subsystem. Some basic and popular geometry compression methods are taken for experiment. One is generalized triangle strip; the other is generalized triangle mesh.

Triangle strip is a very special way of organizing triangles. Considering the following triangulation shown in Figure 1(a), to maximize the use of the available data bandwidth, it is desirable to order the triangles so that consecutive triangles share an edge. Using such an ordering, only the incremental change of one vertex per triangle need to be specified, potentially reducing the rendering time by a factor of three. Besides, such an approach also has obvious benefits in compression for storing and transmitting models. To allow greater freedom in the creation of triangle strips, a “swap”

command permits one to alter the FIFO queuing discipline in a triangle strip as shown in Figure 1(b), the triangle strip can extend further by taking the sequence of (1 2 3 SWAP 4 5 6). The swap command gives greater freedom in the creation of triangle strips at the cost of one bit per vertex. This form of a triangle strip that includes swap command is referred to as a generalized triangle strip [6].




Figure 1. Triangle strips

The concept of a generalized triangle strip structure allows for a compact representation of geometry while maintaining a linear data structure. By confining itself to the linear strips, the generalized triangle strip leaves a potential factor of two in the space occupied.




Figure 2. Generalized triangle mesh

While the geometry in Figure 2 can be represented by one triangle strip, many of the interior vertices appear twice in the strip. This is inherent in any approach wishing to avoid references to the old data. A generalized technique is employed to address this problem. The old vertices are explicitly pushed into a queue, and then implicitly referenced from the queue in the future when the old vertex is desired again. This queue is referred to as the mesh buffer. The combination of generalized triangle strips and mesh buffer references is referred to as a generalized triangle mesh [3].

2.2. Texture Mapping

Texture mapping can substantially enhance the realism and visual complexity of computer generated imagery [2, 9]. Two characteristics of texture mapping are: (1) texture images often require large amounts of memory, and (2) it requires many calculations and texture lookups. These characteristics cause it to be the main performance bottleneck in graphics pipeline. Since the number of pixels that are texture mapped can be quite large (typically tens to hundreds of millions per second), and each one requires multiple texture lookups (usually 8), the memory bandwidth requirements to texture memory can be very large (typically several gigabytes per second). For example, the approximate bandwidth requirement for a professional application running full-screen at a

resolution of 1,280×1,024, and drawing a complex ~~trilinear textured~~ scene filling the graphics window is $1,280 \times 1,024 \times (16 \text{ bytes} + 32 \text{ bytes}) \times 60 \text{ fps} \times 3 = 11.32 \text{ GB/sec}$. Assuming that 3 out of every 4 texel fetches can be satisfied from the texture cache, the bytes transferred from memory to GPU arising from texture fetches would be reduced by about 75%. This may appear somewhat aggressive. However, considering that the neighboring pixels can easily share a significant number of the same texels and a texture surface also typically covers a reasonable screen area in terms of pixels, if the texture cache is large enough, texel reuse will be significantly increased. Its impact on the bandwidth requirement is significant. Using the above illustration, the bytes transferred from memory become: $1,280 \times 1,024 \times (16 \text{ bytes} + 8 \text{ bytes}) \times 60 \text{ fps} \times 3 = 5.66 \text{ GB/sec}$. The bandwidth requirement is reduced to nearly a third and goes from being beyond the limit of traditional memory controller architectures to something actually achievable.

One proposed approach is to use an SRAM cache with each fragment (screen pixel) generator [1]. The premise is that there is a substantial amount of locality of reference in texture mapped scene. By using the small size of SRAM texture cache, lower latency of access to texture memory and higher bandwidth can be achieved. There are three factors important to texture cache behavior (1) the representation of texture images in memory, (2) the cache organization, and (3) the rasterization order on screen.

2.3. The Problem of Geometry Compression

Different triangles traverse orders will definitely affect the access patterns of the texture images. In Figure 3 for example, because each triangle in geometry will be mapped to a certain part of the texture images, the traversal order of T1-T2-T3-T4 will generate different access sequence to the texture images compared to that of T1-T2-T4-T3. The texture cache is used to store a small amount of texture image data for further references. If the triangle traverse orders are rather random, it may generate a large amount of cache misses and thereby worsen the rendering performance.




Figure 3. Traverse orders for triangles

The triangle traversal orders are determined by the geometry compression scheme because these compression schemes exploit a special way of arranging triangles to be sent for rendering. This order is important to different caches used on the graphics chip, so the compression scheme should be aware of the utilization of those caches.

The geometry compression schemes usually ignore the importance of the texture cache; they focus only on the vertex cache. For example, if a triangle strip happens to be mapped to different texture images or many distant parts of one image, the triangle rendering order will not give a good rendering performance.

3. Texture-Map-Based Triangle Strips and Implementation

Based on the above analysis, we present a simple solution by introducing the texture maps into the meshes: the texture-map-based triangle strips (Figure 4).




Figure 4. A texture-map-based triangle strip

The idea is simple. It introduces one level above any existing geometry compression scheme. Now, the triangles are processed in more than one list. First, they will be organized in groups by the texture maps applied to them. In each group, the triangles that are texture mapped by only one image are ordered by a data representation of the geometry compression scheme, e.g., the triangle strips. This way, we can balance the use of the texture caching and vertex caching.

We have implemented the method using C++ and OpenGL and running the experiments on a Pentium IV PC machine with 256MB memory and 1.6GB harddisk.

First, we implemented a 3-D polygonal graphics pipeline. This is responsible for geometry, clipping, and lighting of vertices, rasterization, shading, texture mapping and z-buffering. The pipeline is similar to the one described in [1]. Specifically, the texture mapping implementation is based on mipmap method [14] using OpenGL [12]. Since the pipeline is implemented in software, it is easy to experiment with different rasterization order of triangles.

Second, we implemented a function to trace the runtime Open GL calls that are made by a graphics application. This was done using the Mesa (<http://www.mesa3d.org/>), an open source 3-D graphics library with an API which is similar to that of OpenGL. It is easy to explore into all graphics application function calls and discover the process of how rendering is taking place.

Third, we implemented a trace-driven cache simulator that can model different cache sizes and queuing

disciplines. Whenever the software based fragment generator accesses texels from the memory, all the accessing records will be kept and later passed to the cache simulator. The cache simulator runs after the graphics pipeline.

4. Experiment Study

Experiment study is conducted as follows: First, the geometry compression is carried out on the 3D objects as an offline pre-process. Next, the visiting order of triangles is extracted from the compressed geometry data for future analysis of the traversal order effects. Then the compressed data are sent to the 3-D graphics pipeline of the simulation environment discussed above.

4.1. Geometry Compression

The geometry compression methods are generalized triangle strip and generalized triangle mesh. For generalized triangle strip, the optimized algorithm by Evans et al. is adapted [6]. The algorithm ideally will result in the cost of one vertex per triangle. However, in our experiment, the cost of around 1.1 to 1.25 vertex per triangle is achieved. For the generalized triangle mesh, the implementation from Chow is adapted [3]. By increasing the buffer size from a capacity of 2 vertices to 16 vertices, the cost will go down to a theoretical minimum of 0.5 vertex per triangle. In practice the result turns out to be an average of 0.67 vertex per triangle with an average of 43% vertex reuse. The comparison of the cost per triangle between two algorithms is illustrated in Figure 5.

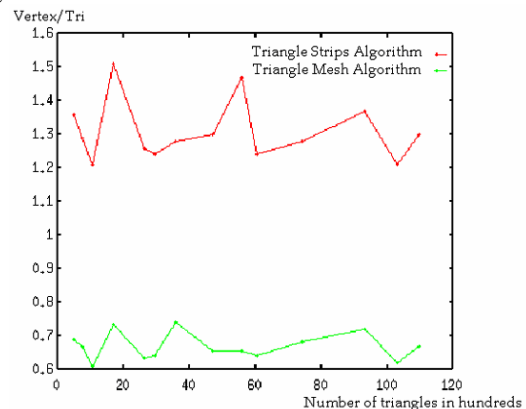


Figure 5. Comparison between two algorithms in term of vertex to triangle ratio

4.2. Texture Cache Size

Working set size is a measure of the amount of data that is actively in use at a particular time. Most applications have a hierarchy of working sets [11]. Suppose the screen resolution is R and the depth complexity is d which represent the average number of

pixels that are rendered for each pixel location. Thus, the number of pixel N_{pix} generated during rasterization is $N_{pix} = R * d$. In our experiment, the MIP level is chosen to be 1:1, so $N_{pix} = N_{tex}$. The working size W is given by $N_{tex} * \text{texel size}$. Figure 6 shows how the cache behaves as the cache size changes. A larger cache implies that more of the textures that are previously transmitted will be reused. But there is also a cost-efficient issue to be considered. Implementing very large texture cache on graphics chip is much more expensive than memory on board. Taking this into consideration, a texture cache size which results in about 75% hit ratio is adapted in NVIDIA's graphics card design [10].

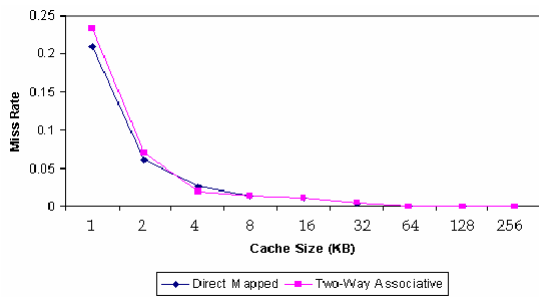


Figure 6. Texture cache miss rates

4.3. Cache Queuing Disciplines

Experiment studies were carried out on two types of queuing disciplines for maintaining the buffer:

- (1) First-in, first-out (FIFO) - this implies that there is no rearrangement of the textures in the buffer. FIFO is the easiest to implement in hardware.
- (2) Least recently used (LRU) – LRU dynamically rearranges the texture in the buffer. The least recently used texture is eliminated when a new texture is added to the queue.

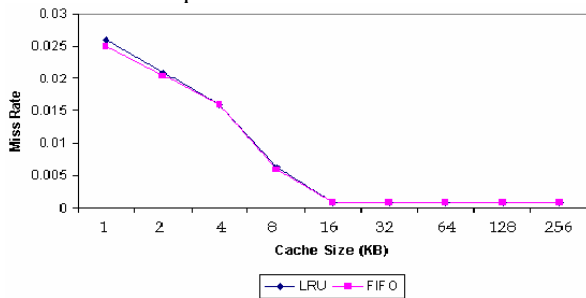


Figure 7. Texture cache miss rate for LRU and FIFO

The results of running test on two queuing disciplines with different cache sizes are presented in Figure 7. It shows the cost of the LRU and FIFO queuing disciplines versus different cache sizes. As can be seen the advantages to be gained from larger buffer sizes starts

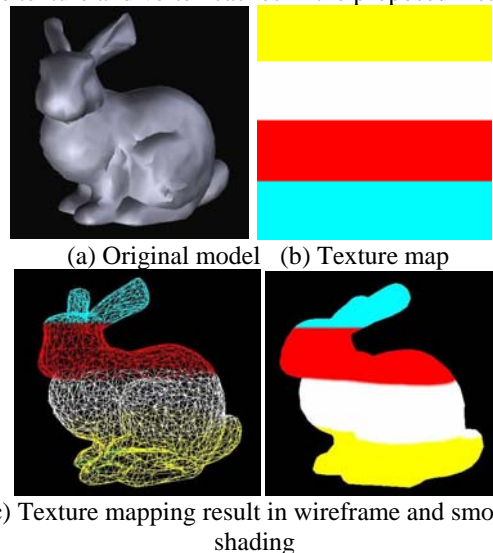
diminishing beyond a buffer size of about 16KB. For buffer sizes less than 16KB, LRU performs better than FIFO scheme by a factor of about 5%.

4.4. Triangle Strip vs. Texture-Based Method

We carried out the follow-up experiment on the texture cache size range from 1KB to 64KB using LRU replacement policy. We compared our proposed algorithm and the triangle strip algorithm. We omit the generalized triangle mesh algorithm because it is just a special type of triangle strip and they will have similar result.

In the comparative study, we looked into two types of objects. One type of objects is constructed with one connected component, like the Stanford bunny. The other type of objects is constructed with some separated components. For example, the foot object we used in the experiment, the skeletons and bones of the foot are naturally separated.

Figure 8 shows an example of texture mapping on the Stanford bunny model. The original model (Figure 8 (a)) is texture mapped with a four color image (Figure 8 (b)). The rendering result is shown in wireframe and smooth shading respectively (Figure 8(c)). The performance study result using the texture cache miss rate as the metric is shown in Figure 9. We compared the proposed method, where the triangle strips are organized in four bins respective to the four colors in the texture map, with the strip compression method, where the triangle strips are organized as only one sequence. We can see the cache miss rate of the proposed method, indicated by “texture” in the figure, is lower than the strip compression method, indicated by “strip” in the figure, due to balance the use of the texture and vertex caches in the proposed method.



(c) Texture mapping result in wireframe and smooth shading

Figure 8. Texture mapping on Stanford bunny

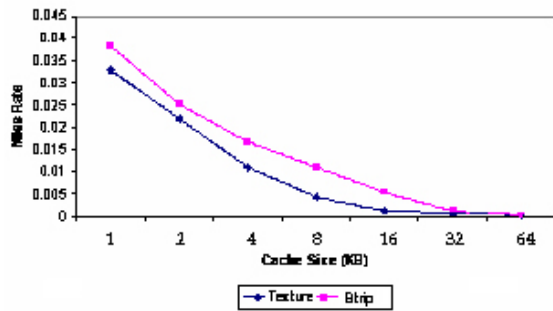


Figure 9. Texture cache miss rate of bunny object

Similarly, we have conducted another experiment on a foot skeleton model (Figure 10) and the performance study result using the texture cache miss rate is shown in Figure 11.



Figure 10. The foot skeleton

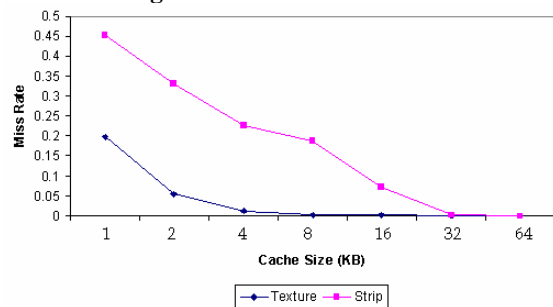


Figure 11. Texture cache miss rate of foot object

From the above results, we can see that for the Stanford bunny model, the proposed method has about 10%~15% less texture cache miss rate than the triangle strip algorithm. For the foot skeleton model, it has about 100%~300% less. To gain the better performance in texture caching performance, we need to pay a cost of a lower vertex caching performance. For the first model, the proposed method has about 1% more in the vertex caching cost than the triangle strip algorithm. For the second model, it has about 10% more. Since the texture information contributes a higher weight to the traffic transferred on the memory interface, overall the proposed method is still better than the triangle strip algorithm.

5. Conclusion and Future Work

We have studied the rendering performance of two geometry compression methods in texture mapping. The analysis and experiment results showed that the

compression methods alone are not optimal. A simple method of improving rendering by texture-map-based triangle strips was proposed and implemented.

Texture cache is very useful in reducing the memory bandwidth and improves rendering performance. Besides the rasterization order, there are two more factors that will affect the texture cache behavior for us to explore in future:

- (1) The cache organization: Study the effect of adapting a single texture cache to two-level caching strategies.
- (2) The representation of texture in storage place: Study what kind of addressing scheme will be most effective for fetching of textures.

6. References

- [1] K. Akeley, "Reality Engine Graphics", *Proc. SIGGRAPH'93*, Anaheim, California, pages 109-116.
- [2] J. F. Blinn, "The Truth about Texture Mapping", *IEEE Computer Graphics and Applications*, March 1990, 10(2), pages 78-83.
- [3] M. Chow, "Optimized geometry compression for real-time rendering", *Proc. Visualization '97*, pages 415-421.
- [4] M. Cox, N. Bhandari, M. Shantz, "Multi-level texture caching for 3D graphics hardware", *ACM SIGARCH Computer Architecture News*, 26(3), 1998, pages 86-97.
- [5] M. Deering, "Geometry Compression", *Proc. SIGGRAPH'95*, Los Angeles, California, pages 13-20.
- [6] F. Evans, S. Skiena, A. Vashney, "An Optimizing Triangle Stripes for Fast Rendering", *Proc. Visualization '96*, San Francisco, California, 1996, pages 319-326.
- [7] J. Foley, A. van Dam, S. Feiner, J. Hughes, "Computer Graphics Principles and Practice", Second Edition, Addison-Wesley Publishing Company, Inc, Boston, Massachusetts, 1990.
- [8] Z. S. Hakura, A. Gupta, "The design and analysis of a cache architecture for texture mapping", *ACM SIGARCH Computer Architecture News*, 25(2), 1997, pages 108-120.
- [9] P. S. Heckbert, "Fundamentals of Texture Mapping and Image Warping", Master's thesis, UCB/CSD 89/516, University of California at Berkeley, 1989.
- [10] NVIDIA® Corporation, "Lightspeed Memory Architecture II Tech Brief", Santa Clara, California, April 2002.
- [11] E. Rothberg, J. P. Singh, and A. Gupta, "Working Sets, Cache Sizes, and Node Granularity Issues for Large-Scale Multiprocessors", *ACM SIGARCH Computer Architecture News*, 21(2), 1993, Pages: 14 - 26.
- [12] M. Segal, and K. Akeley, "The OpenGL Graphics System: A Specification", Version 1.2.1, Silicon Graphics, Inc., Mountain View, California, April 1999.
- [13] Sun Microsystems, Inc. "Creator Graphics Technology, White Paper", Santa Clara, California, November 1995.
- [14] L. Williams, "Pyramidal Parametrics", *Proc. of SIGGRAPH '83*, pages 1-11.