

Compression of 3D Mesh Geometry and Vertex Attributes for Mobile Graphics

Jongseok Lee

Department of Computer Science and Engineering
Pohang University of Science and Technology (POSTECH)
Pohang, Republic of Korea
thirdeye@postech.ac.kr

Sungyul Choe

Samsung Electronics
Suwon, Republic of Korea
sungyul.choe@samsung.com

Seungyong Lee

Department of Computer Science and Engineering
Pohang University of Science and Technology (POSTECH)
Pohang, Republic of Korea
leesy@postech.ac.kr

Received 9 August 2010; Accepted 26 August 2010

This paper presents a compression scheme for mesh geometry, which is suitable for mobile graphics. The main focus is to enable real-time decoding of compressed vertex positions while providing reasonable compression ratios. Our scheme is based on local quantization of vertex positions with mesh partitioning. To prevent visual seams along the partitioning boundaries, we constrain the locally quantized cells of all mesh partitions to have the same size and aligned local axes. We propose a mesh partitioning algorithm to minimize the size of locally quantized cells, which relates to the distortion of a restored mesh. Vertex coordinates are stored in main memory and transmitted to graphics hardware for rendering in the quantized form, saving memory space and system bus bandwidth. Decoding operation is combined with model geometry transformation, and the only overhead to restore vertex positions is one matrix multiplication for each mesh partition. In our experiments, a 32-bit floating point vertex coordinate is quantized into an 8-bit integer, which is the smallest data size supported in a mobile graphics

Copyright(c)2010 by The Korean Institute of Information Scientists and Engineers (KIISE). Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Permission to post author-prepared versions of the work on author's personal web pages or on the noncommercial servers of their employer is granted without fee provided that the KIISE citation and notice of the copyright are included. Copyrights for components of this work owned by authors other than KIISE must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires an explicit prior permission and/or a fee. Request permission to republish from: JCSE Editorial Office, KIISE. FAX +82 2 521 1352 or email office@kiise.org. The Office must receive a signed hard copy of the Copyright form.

library. With this setting, the distortions of the restored meshes are comparable to 11-bit global quantization of vertex coordinates. We also apply the proposed approach to compression of vertex attributes, such as vertex normals and texture coordinates, and show that gains similar to vertex geometry can be obtained through local quantization with mesh partitioning.

Categories and Subject Descriptors: Computer Graphics [**Modeling**]: Mesh Compression

General Terms: Mesh compression, Geometry encoding

Additional Key Words and Phrases: Mobile graphics, Local quantization, Mesh partitioning

1. INTRODUCTION

Recent mobile devices, such as mobile phones, PDAs, and hand-held game players, are equipped with quite decent features of computing power, storage, graphics, and network. Especially, for mobile graphics, API standards such as OpenGL ES and JSR-184 have been proposed [Pulli et al. 2005], and the graphics performance has been drastically improved with development of mobile GPUs. With these advances, 3D graphics is becoming more popular in mobile applications of games, user interfaces, screen savers, map services, etc. This trend invokes increasing demands for various graphics techniques that reflect unique characteristics of mobile devices.

Despite the recent improvements of mobile hardware, compared to desktop PC graphics, mobile graphics is still constrained by limited resources of low computing powers, small amounts of memory, and low bandwidths. In addition, while mobile devices are usually powered by batteries, battery technology has not caught up with the energy requirement of 3D graphics processing.

Data compression can provide an excellent solution to help overcome these limitations. Obviously, compressed graphics data enable us to better utilize storage space and network bandwidth. Furthermore, by decoding compressed data on-the-fly on graphics hardware [Ström and Akenine-Möller 2005], we can reduce memory access and data transmission through the system bus, which saves power consumption. However, to be suitable for mobile graphics, a compression technique should be simple enough not to impose a big overhead on the low computing power.

In this paper, we propose a simple and effective compression technique for mesh geometry, which can be used for mobile graphics in practice. The main design objective is real-time on-the-fly decoding of vertex positions on graphics hardware. Such property enables the compressed form to be used for both storing the vertex positions of a mesh in main memory and transmitting to graphics hardware for rendering. As a result, main memory space and system bus bandwidth required for mesh geometry can be reduced in a rendering system.

To fulfill the objective with reasonable compression ratios, we use local quantization of vertex positions with mesh partitioning. With local quantization, vertex positions can be restored independently from each other by simple transformations, and compression rate-distortion can be controlled by manipulating the local bounding cubes through mesh partitioning. However, straightforward local quantization that processes each mesh partition separately suffers from visual seams along the partitioning boundaries of the restored mesh. To prevent the problem, we quantize all mesh partitions in a coherent way to have the same-sized and axis-aligned quantized cells. Then, the distortion of the restored mesh relates with the size of the quantized

cells and we propose a mesh partitioning algorithm to minimize the size.

The setting of our experiments is to quantize a 32-bit floating point vertex coordinate into an 8-bit integer, which is the smallest data size supported in a mobile graphics library. Local quantization incurs encoding overhead for bounding cube information and duplicated vertices. However, experimental results show that the data reduction ratios are still over 70%. The distortions of the restored meshes are comparable to 11-bit global quantization of vertex coordinates, which is close to 12-bit global quantization that is considered enough for representing the geometry of moderate-sized meshes in mesh compression [Alliez and Gotsman 2005; Peng et al. 2005]. These results imply that we can achieve the accuracy of 11-bit global quantization, which is little worse than 12-bit global quantization conventionally used in mesh compression experiments [Alliez and Gotsman 2005; Peng et al. 2005], with 8-bit integer representation of vertex coordinates, which can be efficiently processed by mobile graphics APIs.

In addition to positions, we apply our compression technique to other vertex attributes, such as normals and texture coordinates. Experimental results demonstrate that our technique can produce satisfactory results for encoding the attributes.

The remainder of this paper is organized as follows. In Sec. 2, we review the related work. Sec. 3 introduces the basic idea of our technique. Sec. 4 and Sec. 5 explain the encoding and decoding processes, respectively. We give experimental results not only for geometry in Sec. 6 but also for other vertex attributes in Sec. 7. Finally, we have conclusion and discussion in Sec. 8.

2. RELATED WORK

Mesh compression has been an active research area in graphics and excellent techniques have been developed for connectivity and geometry encoding. A comprehensive review of mesh compression techniques, which is not within the scope of this paper, can be found in excellent survey papers [Alliez and Gotsman 2005; Peng et al. 2005].

Although previous mesh compression algorithms offer good performance in compression ratio, their encoding and decoding processes are rather complicated and can impose much overhead on the relatively low computing power of mobile hardware. More importantly, most of them sequentially traverse mesh elements for encoding, which prohibits parallel decoding of encoded elements on graphics hardware. Consequently, previous algorithms cannot be directly adopted for mesh compression on mobile devices.

Calver [Calver 2002] described the basic principles for various quantization methods of vertex geometry which allow real-time decoding with programmable vertex shaders. He mentioned that local quantization can be used for vertex geometry to obtain better compression rate-distortion. However, he presented no solution for the problem of visual seams and no specific technique to partition a given mesh for local quantization.

Purnomo et al. [Purnomo et al. 2005] refine the idea of Calver and allocate different numbers of bits for the components of vertex attributes, such as positions, normals, and texture coordinates, to minimize the image-space error. At rendering time, compressed vertex data, which have been packed into fixed 96 bits, are transmitted

to a programmable vertex shader and restored on-the-fly. However, in terms of vertex geometry compression, their technique uses global quantization, which provides worse compression rate-distortion than local quantization.

In contrast to [Calver 2002], in this paper, we provide effective solutions to handle visual seams and mesh partitioning for local quantization. Compared to [Purnomo et al. 2005], in addition to the difference of using local quantization, our technique requires no programmable vertex shader because the decoding operation can be incorporated into the standard geometry transformation, which makes it applicable to a wide range of mobile graphics platforms.

3. BASIC IDEA

In this paper, we only consider compression of mesh geometry. We assume that mesh connectivity is represented in an adequate form (e.g., triangle strips), which is supported in mobile graphics API standards, such as JSR-184 [Pulli et al. 2005].

To restore vertex geometry in real time on a mobile device, a compression algorithm should satisfy the following properties;

- The decoding process is simple enough for real-time performance.
- The decoding operations of vertex positions have no dependency among each other to allow parallel processing.
- Encoded data are represented in a fixed-size format, which matches with the data channel size between CPU and graphics hardware.

A simple solution to satisfy these requirements is the quantization of vertex positions. Original vertex positions, represented by 32-bit floating point numbers, can be encoded with fewer bits by truncating least significant bits. We can easily restore the original positions from the quantized ones by simple additions and multiplications.

In addition, there is no dependency among quantized vertex positions. In compression of irregular meshes with reasonable sizes, 12-bit global quantization is assumed to introduce no strong visual artifacts and used as a preprocessing step before geometry encoding [Alliez and Gotsman 2005; Peng et al. 2005]. However, the use of 12-bit quantized positions can cause unnecessary bitwise operations or waste of bandwidth because most graphics APIs support only fixed-sized data channels (e.g., 8, 16, and 32




Figure 1. Comparison between global and local quantization results with 8 bits. (a) Global quantization; (b) local quantization with 128 partitions.

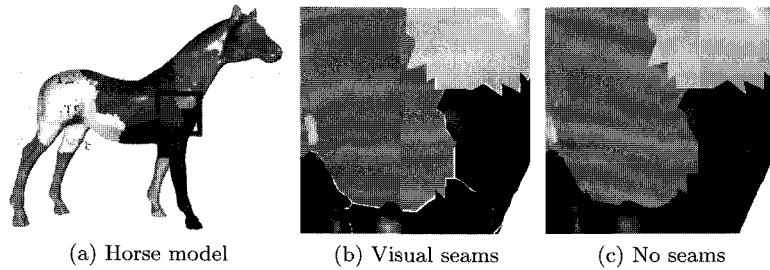


Figure 2. (a) Rendering of a horse model. (b) and (c) show zoom-ins of the red rectangle in (a); (b) visual seams from straightforward local quantization; (c) visual seams have been resolved by our method.

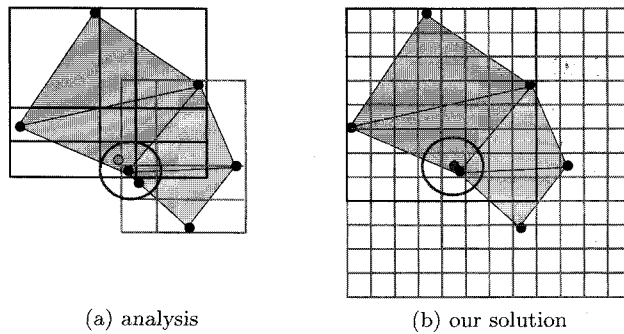


Figure 3. (a) A black vertex shared by two sub-meshes is restored onto the green and red positions due to the difference of locally quantized cells. (b) The restored positions from two sub-meshes have become the same due to the alignment of locally quantized cells.

bits). On the other hand, if we globally quantize vertex positions with 8 bits, we may have severe distortions of mesh geometry (see Figure 1(a)).

To resolve such a problem, Calver [Calver 2002] described a local quantization method, which can take full advantage of graphics hardware. A mesh is divided into several sub-meshes and each sub-mesh is independently quantized. Then, the distortion of the restored mesh is reduced because the size of a bounding cube, which is the range of quantization, decreases. As a result, 8-bit local quantization of mesh geometry can provide better accuracy for restored vertex positions than global quantization (see Figure 9(b)). However, in this case, the restored mesh contains a critical problem, i.e., visual seams along the sub-mesh boundaries, as shown in Figure 2(b). Calver [Calver 2002] did not provide any solution to handle the problem.

The main cause of the visual seam problem is that shared boundary vertices of sub-meshes should be encoded more than once. Sub-meshes are disjoint subsets of mesh faces and a boundary vertex of a sub-mesh also belongs to one or more other sub-meshes, except at the mesh boundary. When sub-meshes are separately encoded, a boundary vertex v shared among k sub-meshes is encoded k times, once for each sub-mesh, with different local quantization. At decoding time, a restored vertex position is the center of a quantized cell. The restored positions of v in the k sub-meshes do

not match each other if the corresponding locally quantized cells of the sub-meshes are not aligned. See Figure 3(a) for illustration.

The basic idea of our solution for the visual seam problem is to align the locally quantized cells of sub-meshes sharing boundary vertices. As illustrated in Figure 3(b), if the locally quantized cells are aligned, the restored positions of a shared vertex will be the same among the sub-meshes. This constraint of aligned quantized cells should hold for every pair of sub-meshes sharing a boundary vertex. Consequently, the constraint propagates through the adjacency of sub-meshes, and the locally quantized cells of all sub-meshes should have the same size and aligned axes. In this paper, we use the largest bounding cube of sub-meshes to determine the cell size and axes of the aligned local quantization, which is then applied to all sub-meshes. Figure 2(c) shows that the rendering result from our local quantization method does not contain any visual seams.

4. ENCODING PROCESS

The encoding process consists of two parts; mesh partitioning, which divides an input mesh into several sub-meshes, and local quantization of vertex positions, which uses the mesh partitioning result.

4.1 Mesh Partitioning

With quantization of vertex positions, the distortion of a restored mesh from the original is dominated by the size of quantized cells. As described in Sec. 3, in our approach, the size of locally quantized cells is determined by the largest bounding cube of mesh partitions. Hence, to minimize the distortion, the largest bounding cube should be made as small as possible.

We adapt the mesh partitioning framework based on *Lloyd's algorithm* [Lloyd 1982], which has been successfully used for mesh segmentation [Sander et al. 2003; Choe et al. 2006]. After initial partitioning has been obtained with a given number of partitions, the partitioning result is iteratively updated by repeating two steps; *seed re-computation* and *region growing*. In the seed re-computation step, the seed triangles are repositioned at the centers of partitions for the next update. In the region growing step, faces are added to partitions in the increasing order of distances from the nearest seed faces.

For initial partitioning, we use the *hierarchical face cluster merging* approach [Garland et al. 2001]. At the beginning, each face is a single partition. We iteratively merge two neighbor partitions into one until the number of partitions is reduced to the desired number. To select the merged pair of partitions at each iteration, we use the cost function defined by

$$F(C_i, C_j) = \max\{x_{ij}, y_{ij}, z_{ij}\}, \quad (1)$$

where x_{ij} , y_{ij} , and z_{ij} are respectively the x , y , and z sizes of the axis aligned bounding cube containing the two partitions, C_i and C_j . By minimizing $F(C_i, C_j)$, we can reduce the bounding cube sizes of resulting mesh partitions.

For seed-recomputation and region growing, we use L_∞ metric to define the distance function. To reposition the seed of a mesh partition, we first compute the center of the




Figure 4. Mesh partitioning results.

bounding cube of the partition, which can be considered as the L_∞ center of the partition. Although we can use the center as the new seed of the partition, we select the nearest face to the center as the new seed face because we use the adjacency of faces in region growing for efficiency. In region growing, the distance of a face f from the seed face of a partition C is defined by

$$F(f, C) = \max\{\delta_x, \delta_y, \delta_z\}, \quad (2)$$

where δ_x is the maximum of the x -distances of three vertices of f from the center of the seed face of C . δ_y and δ_z are defined similarly.

In mesh segmentation techniques [Sander et al. 2003; Choe et al. 2006], it has been demonstrated that the result of Lloyd's algorithm partitions a given mesh into almost equal-sized sub-meshes, where the size is measured in the adopted distance function. Similarly, with L_∞ metric, we can expect that the sub-meshes resulting from Lloyd's algorithm have almost same sizes of bounding cubes because a L_∞ sphere in 3D is a cube. As a result, the largest bounding cube of mesh partitions can be made as small as possible for a given number of partitions. Figure 4 shows examples of the final partitioning result.

4.2 Local Quantization

In this paper, we use 8-bit quantization for vertex coordinates. The steps for local quantization can be summarized as follows;

- (1) Calculate the bounding cube for each mesh partition and find the largest x , y , and z bounding cube sizes among all partitions.
- (2) Calculate (C_x, C_y, C_z) , the x , y , and z sizes of the quantized cell, by dividing the largest x , y , and z bounding cube sizes by $(2^8 - 1)$, respectively.
- (3) Quantize all vertex positions by dividing original vertex coordinates by the quantized cell size (C_x, C_y, C_z) and truncating the fractional values.
- (4) For each partition, find the minimum quantized coordinates for the x , y , and z axes, and keep the values as the offsets (O_x, O_y, O_z) . Then, subtract (O_x, O_y, O_z) from the quantized coordinates of vertices, obtaining the final quantized coordinates (P_x, P_y, P_z) in the range of $[0, 255]$.

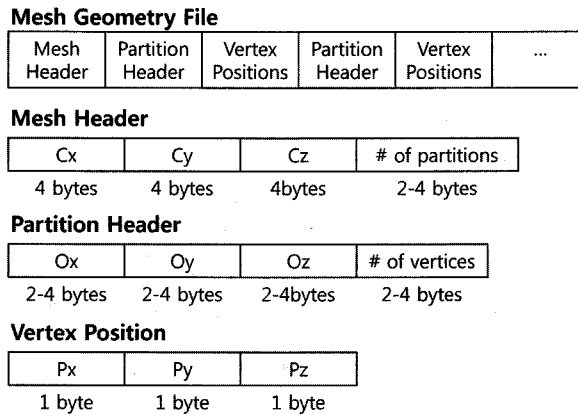


Figure 5. File structure for encoded mesh geometry. C, O, and P denote the quantized cell size, offset value, and vertex coordinate, respectively.

- (5) Save the size of the quantized cell, offsets for partitions, and quantized vertex positions.

Note that for calculating the size of the quantized cell, we use (2^8-1) , instead of 2^8 , to guarantee that the quantized vertex coordinates after offsetting fall in the range of $[0, 255]$.

4.3 File Structure

After local quantization, to restore vertex positions, we need the size of the quantized cell, offsets for partitions, and quantized vertex positions. The size of the quantized cell is stored only once for a mesh and represented by three 32-bit floating point numbers. The offsets for a mesh partition correspond to the quantized coordinates of the origin of the bounding cube of the partition. To represent the offsets, we use 16 to 32-bit integers depending on the mesh size. Since we are using 8-bit local quantization, a quantized vertex position requires 24 bits, i.e., three 8-bit unsigned integers. In addition, we should record the number of partitions and the number of vertices in each partition. These numbers are represented by 16 to 32-bit integers depending on the mesh size. Figure 5 shows the file structure that contains encoded mesh geometry.

There are two kinds of storage overhead incurred by encoding with local quantization. One is the header overhead, which corresponds to the data in the mesh header and partition headers in Figure 5. The other is the vertex overhead. As mentioned in Sec. 3, the boundary vertices shared among sub-meshes should be duplicated when they are locally quantized with different offsets. Hence, in the encoded mesh geometry, there are more quantized vertex positions than the number of vertices.

5. DECODING PROCESS

When we render a mesh with encoded geometry, we first restore the quantized cell

size (C_x, C_y, C_z) from the mesh header. Then, each partition of the mesh is rendered with the following procedure;

- (1) Restore the offset values (O_x, O_y, O_z) from the partition header.
- (2) Calculate a 4×4 decoding matrix

$$M_d = \begin{pmatrix} C_x & 0 & 0 & C_x \cdot O_x \\ 0 & C_y & 0 & C_y \cdot O_y \\ 0 & 0 & C_z & C_z \cdot O_z \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3)$$

- (3) Multiply matrix M_d by the geometry transformation matrix, e.g., the modelview matrix in OpenGL.
- (4) Render the polygons in the partition.

In the decoding process, vertex geometry data are transmitted from main memory to graphics hardware in the encoded form, i.e., an 8-bit integer for a vertex coordinate. When we render a polygon, the vertex positions are automatically restored through the geometry transformation matrix which has been modified by the decoding matrix M_d . Note that M_d is the composition of two matrices, one for adding offsets to locally quantized vertex coordinates in $[0, 255]$ and the other for multiplying the quantized cell size by the quantized coordinates after offsetting to obtain 32-bit floating point coordinates.

Our decoding process is nicely incorporated into the standard rendering pipeline. As a result, at rendering time, the only major overhead is one matrix multiplication per a partition, which can be considered negligible. Although boundary vertices are duplicated among partitions, the number of rendered polygons does not change and the number of vertices processed for rendering remains the same.




Figure 6. Models used in our experiments. From the top left in the clockwise order, children, dancer, feline, filigree, horse, and squirrel.

6. EXPERIMENTAL RESULTS

6.1 Experimental Settings

We have implemented and tested our algorithm using a variety versions of Khronos OpenGL graphics APIs [Pulli et al. 2005], such as OpenGL 1.x, OpenGL 2.x, and OpenGL ES 2.x emulator. In OpenGL 1.x and OpenGL 2.x implementations, the decoding matrix M_d is formed and multiplied by the model-view matrix before rendering each sub-mesh. For OpenGL ES 2.x emulator, M_d is formed and used when we set up a vertex shader for a sub-mesh. Experiments have been performed on a desktop PC with an Intel Pentium 4 3.6GHz CPU, 2GB RAM, and an NVIDIA GeForce 7800 GTX graphics card.

Figure 6 shows models used in our experiments. In the experiments, we measure the reduction ratio for vertex geometry data. We also count the number of duplicated vertices along the boundaries of mesh partitions. That is,

$$\text{Data Reduction Ratio} = \left(1 - \frac{\text{compressed data}}{\text{uncompressed data}}\right) \times 100$$

$$\text{Vertex Overhead} = \frac{\text{\# of duplicated vertices}}{\text{\# of vertices}} \times 100$$

The distortion of a mesh restored from encoded geometry is measured from the original mesh in the average L2 norm using the Metro tool [Cignoni et al. 1998].

6.2 Compression Results

Figure 7 shows the distortions, data reduction ratios, and vertex overheads for increasing numbers of partitions with the dancer model. When the number of partitions becomes larger, the data reduction ratio decreases because we need to store more partition headers and the vertex overhead increases. In contrast, the distortion of the restored mesh is reduced with more partitions because the bounding cubes of partitions become smaller, which increases the accuracy of quantization. As we can expect,

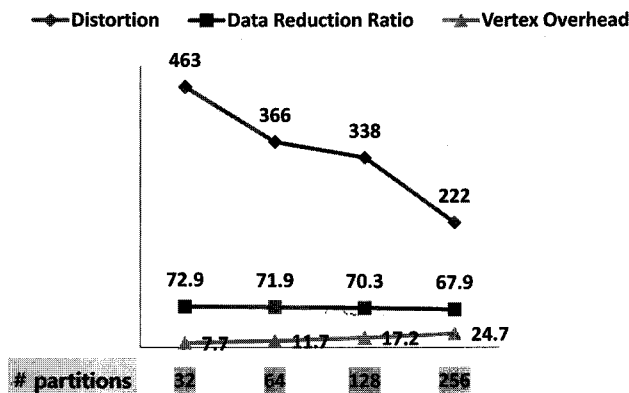


Figure 7. Compression results with different numbers of partitions for dancer model. Distortions are shown in the unit of 10^{-6} .

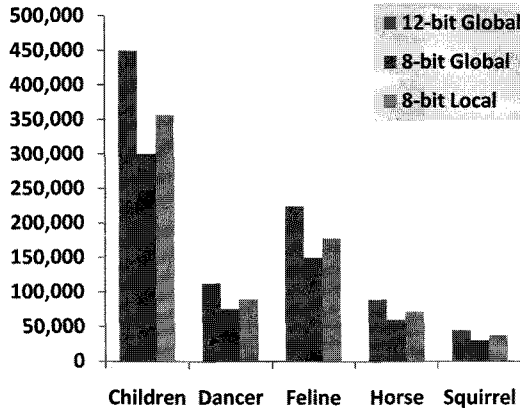


Figure 8. Compressed data sizes with different quantization methods.

Table I. Comparison of the distortions of restored meshes between 11-bit global, 8-bit global, and our 8-bit local quantization.

Model	# V	# P	11-bit Global	8-bit Global	8-bit Local
Children	100,000	564	0.000436	0.003429	0.000395
Dancer	24,998	135	0.000449	0.003622	0.000293
Feline	49,864	270	0.000442	0.003583	0.000373
Filigree	514,300	2480	0.000452	0.003642	0.000160
Horse	19,851	128	0.000430	0.003375	0.000462
Squirrel	9,995	52	0.000416	0.003360	0.000966

such tendency appeared for all models in our experiments.

Figure 8 compares compressed data sizes of vertex geometry with respect to the quantization methods. The compressed data size from our method for local quantization is variable according to the number of partitions, so we use the number of partitions which provides 70% (+0.1) data reduction ratio. Both 8-bit global and our 8-bit local quantization methods require the same number of bits to represent a vertex position, i.e., 8 bits for each coordinate. However, our 8-bit local quantization has additional data for partition headers and vertex overhead, and the data size corresponds to the amount between 9 and 10 bit global quantization results.

Table I shows the distortions of restored meshes with different quantization methods. As shown in Figure 7, the distortion from our local quantization varies with the number of partitions. In Table I, we use the number of partitions which provides 70% (+0.1) data reduction ratio. Table I shows that the distortion of 8-bit local quantization is quite less than that of 8-bit global quantization for every model. In addition, the restored quality of 8-bit local quantization is better than or similar to 11-bit global quantization in most cases, except the Squirrel model which is the smallest.

Figure 9 shows visual comparison of the restored meshes from global and local quantization with 8 bits. We can clearly see visual degradation in Figure 9(a), while




Figure 9. Visual quality comparison between 8-bit global and our 8-bit local quantization.

Table II. Comparison between spatial subdivision and our partitioning method. # P denotes the number of partitions and # V' denotes the number of vertices including duplicated vertices.

Model	# P	Subdivision		Our Partitioning	
		# V'	L^2 d-error	# V'	L^2 d-error
Children	264	112,439	0.000762	112,287	0.000537
Dancer	198	32,386	0.000498	30,433	0.000253
Feline	194	57,063	0.000522	57,364	0.000470
Filigree	232	558,864	0.000463	542,884	0.000352
Horse	159	23,679	0.000506	23,780	0.000438
Squirrel	254	13,756	0.000566	14,011	0.000498

Figure 9(b) is almost indistinguishable from the original in Figure 6. This example demonstrates that 8-bit local quantization has enough accuracy for representing the vertex geometry.

Table II compares compression results with our mesh partitioning method and spatial subdivision. For spatial subdivision, we divide the bounding cube of a given mesh into cells with 11-bit global quantization and cluster consecutive 256^3 cells into one partition. For fair comparison, we used the same number of partitions for our method as the spatial subdivision. In Table II, the vertex overheads of two methods are similar but the distortions of ours are less than spatial subdivision. Furthermore, spatial subdivision cannot freely control the number of partitions, which relates with the data reduction ratio and distortion, while our method can.

Table III shows the overhead in distortions of restored meshes induced for removing visual seams of mesh partitions. If we do not care about the visual seams, for each partition, we can compute a tight oriented bounding cube and perform more accurate local quantization than our globally aligned one. In Table III, as we can expect, the distortions from the tight bounding cubes are less than our method. However, the differences are not really big and visual seams are a more severe problem than a little more distortions in rendering a mesh.

By controlling the number of partitions for our 8-bit local quantization, we can

Table III. Comparison of the distortions of restored meshes between tight oriented and globally aligned bounding cubes. Recall that tight oriented bounding cubes result in visual seams. The data reduction ratio used for determining the number of partition for each model is 70% (+0:1).

Model	# P	Tight Cubes	Aligned Cubes
Children	564	0.000292	0.000395
Dancer	135	0.000246	0.000293
Feline	270	0.000279	0.000373
Filigree	2480	0.000136	0.000160
Horse	128	0.000416	0.000462
Squirrel	52	0.000824	0.000966

Table IV. Comparison between global and local quantization with similar amounts of distortions for the dancer model.

Global Quantization			Local Quantization		
b/v	L^2 d-error	size (KB)	# partitions	L^2 d-error	size (KB)
10	0.000884	92	10	0.000840	76
11	0.000449	101	40	0.000427	80
12	0.000228	110	256	0.000222	94

achieve a similar level of distortions to global quantization with a specified number of bits. Table IV shows experimental results with the dancer model. In Table IV, we can see that the compressed data size for vertex geometry with local quantization is always smaller than the corresponding global quantization. Table IV also implies that we can easily satisfy a requirement on the amount of distortions by controlling the number of partitions used for local quantization.

7. ENCODING OTHER VERTEX ATTRIBUTES

In addition to positions, our local quantization method can be applied to other vertex attributes, such as normals and texture coordinates. In this section, we present the encoding results of these attributes and discuss a desired change for the chartification method. In the experiments, the vertex coordinates of a mesh were represented by 8-bit local quantization.

7.1 Vertex Normal Encoding

Table V compares the quality of the rendered images between the global and local quantization of vertex normals. The quality is represented by Peak Signal-to-Noise Ratio (PSNR) from the image rendered with the original normals. With quantization, each component of a normal vector is encoded with 8 bits, similar to a vertex coordinate. For local quantization, we use the same mesh partitions used for vertex encoding in Sec. 6. In decoding time, a locally quantized vertex normal can also be restored by one matrix multiplication.

Table V. Quality comparison of the rendered images between local and global quantization of vertex normals.

Model	8-bit Global	8-bit Local
Children	57.00 dB	60.46 dB
Dancer	69.21 dB	72.40 dB
Feline	59.47 dB	63.60 dB
Filigree	61.60 dB	64.93 dB
Horse	60.97 dB	64.36 dB
Squirrel	61.58 dB	64.98 dB

In image processing domain, it is known that human vision system cannot recognize the difference of two images when PSNR is larger than 30-50 dB. Table V shows that in our experiments, both 8-bit global and local quantization for vertex normals provide satisfactory image quality. Nevertheless, 8-bit local quantization produces better results in Table V and will be beneficial when the rendered image is sensitive to vertex normals.

7.2 Texture Coordinate Encoding

Figure 10 compares texture mapping results of a mesh between local and global quantization of texture coordinates. The spaceship model is segmented into 128 partitions using the mesh partitioning method in Sec. 4.1. Texture coordinates can be encoded and restored in a similar way to vertex normals. The boundary of a circular

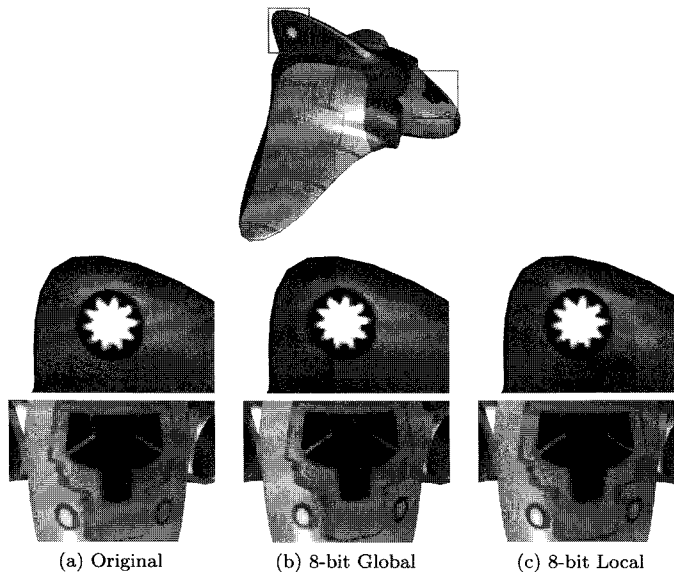


Figure 10. Comparison of texture mapping result between 8-bit global and 8-bit local quantization of texture coordinates.

mark in middle row is distorted in (b), while it preserves the shape in (c). Similarly, the patterns in the bottom row are slightly shifted in (b), while it is not in (c).

7.3 Modification of Mesh Partitioning

For vertex normal and texture coordinate encoding, we used the mesh partitions computed by considering vertex positions, and obtained satisfactory results. However, the results can be improved by modifying the mesh partitioning method to incorporate other vertex attributes than positions.

In our mesh partitioning method, the partitioning results are mainly determined by the cost function Eq. (2). We can modify the cost function by combining the differences of vertex normals and texture coordinates as well as positions. See [Sander et al. 2003; Choe et al. 2006] for an example of combining several terms in a cost function. As a result, the maximum bounding boxes of the vertex attributes can be reduced in local quantization, and the distortions of encoded attribute values can become smaller.

8. SUMMARY AND DISCUSSION

The compression rate of our geometry encoding is $24+\alpha$ bits per vertex (bpv), where α comes from the overhead of partition headers and duplicated vertices. This rate is far worse than the state-of-the-art techniques, which give about 10 to 13 bpv for geometry encoding [Alliez and Gotsman 2005; Peng et al. 2005]. However, our goal is to develop a mesh geometry compression technique suitable for mobile graphics, instead of achieving best compression rates.

Our technique has the following nice properties;

- *simplicity* It is simple enough for real-time decoding on a mobile device. The decoding operation is combined with the model geometry transformation, and the only overhead is one matrix multiplication for each mesh partition and the vertices in the same partition can be processed in parallel, exploiting the capability of graphics hardware.
- *on-the-fly decoding* At rendering time, vertex positions are restored on-the-fly through the model geometry transformation. Vertex data can be stored in main memory and transmitted to graphics hardware in the quantized form, saving memory space and system bus bandwidth.
- *better compression ratio* Our technique provides better compression ratios and better restoration quality of vertex geometry than the previous mobile compression technique [Purnomo et al. 2005], which uses global quantization.
- *fixed data size* Encoded vertex coordinates have the same fixed data size, which offers benefits in addressing and random access of vertices. If the data size is supported in a mobile graphics library, no bitwise operation is needed to extract encoded vertex coordinates.
- *compatibility* Since the decoding operation is compatible with the standard graphics pipeline, in addition to mobile graphics, our technique can be used for PC and game console graphics without modification.
- *flexibility* A requirement on the amount of distortions can easily be satisfied by controlling the number of partitions used for local quantization.

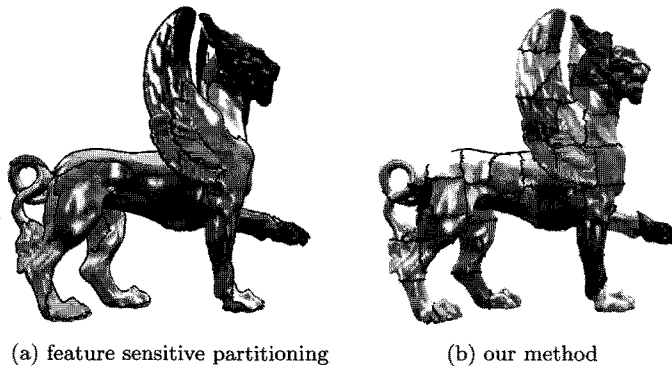


Figure 11. Comparison of partitioning results with 64 partitions between (a) feature sensitive partitioning and (b) our method.

– *usefulness* Our framework can be used for not only the vertex positions but also various vertex attributes, such as vertex normal and texture coordinates.

In mobile hardware, basic units for data processing are 8, 16, and 32 bits. If a different number of bits are used, data handling, especially transmission through the system bus, will become inefficient. As we mentioned in Sec. 3, the use of more than 12 bits for a vertex coordinate does not help a lot to increase the visual quality of rendering. Therefore, our choice of allocating 8 bits for each vertex coordinate seems best in terms of compatibility with mobile hardware and visual quality.

Our mesh partitioning method minimizes the bounding cubes of partitions. In Sec. 6, we showed that our method is better than spatial subdivision in terms of the distortions of restored meshes. There are other alternatives for mesh partitioning, including feature sensitive partitioning used for multi-chart geometry images [Sander et al. 2003]. In feature sensitive partitioning, a mesh is divided along feature lines and planar partitions are preferred. Figure 11 compares the partitioning results from both methods with the same number of partitions. In Figure 11, the volumes of some bounding cubes from feature sensitive partitioning could be smaller than those from ours. However, when we compute the largest x , y , and z sizes of the bounding cubes, our method will give smaller values than feature sensitive partitioning. The largest x , y , and z sizes are used for determining the sizes of aligned locally quantized cells, which are directly related to the distortions of a restored mesh. As a result, our method is more effective for reducing the distortions of a restored mesh than feature sensitive partitioning.

In mesh partitioning, as a preprocessing, we can change the orientation of a given mesh so that the major axes of the mesh are aligned with the coordinate axes for vertex positions. We performed experiments for the preprocessing, where major axes of a mesh are determined by PCA (Principal Component Analysis) [Press et al. 1992] of vertex coordinates. For the models used in this paper, the major axes of a mesh are already almost aligned to the coordinate axes, and the distortions from local quantization with and without the preprocessing are not much different. So we did

not use the processing in the experiments in this paper. However, such preprocessing would be useful if the major mesh axes and the vertex coordinate axes are totally misaligned.


Our geometry compression scheme can be combined with any data structure (e.g., triangle strips) used for representing meshes in main memory. In our scheme, a given mesh is partitioned into sub-meshes. When each sub-mesh is stored in main memory using the selected data structure, our local quantization method can be used for representing the vertex coordinates with reduced memory space. Similarly, our scheme can be adapted for a mesh file format. For example, for an OBJ file, we can set each sub-mesh as a group and represent vertex coordinates in the quantized form, where additional information, similar to the partition header in Figure 5, is stored for each group.

Acknowledgments


The dancer, dancing children, filigree, and feline models are courtesy of IMATI, IMATI-GE, SensAble technologies, and Multi-Res Modeling Group at Caltech, respectively. The dancer, dancing children, filigree, horse, and squirrel models are provided by the AIM@SHAPE Shape Repository. This work was supported by the Industrial Strategic Technology Development Program of MKE/MCST/KEIT (KI001820, Development of Computational Photography Technologies for Image and Video Contents), the ERC Program of MEST/NRF (R11-2008-007-01002-3), and Samsung Electronics.

REFERENCES


- ALLIEZ, P. AND GOTSMAN, C. 2005. Recent advances in compression of 3D meshes. In *Advances in Multiresolution for Geometric Modelling*, N. A. Dodgson, M. S. Floater, and M. A. Sabin, Eds. Springer-Verlag, 3–26.
- CALVER, D. 2002. Vertex decompression in a shader. In *Direct3D ShaderX: Vertex and Pixel Shader Tips and Tricks*, W. F. Engel, Ed. Wordware, 172–187.
- CHOE, S., AHN, M., AND LEE, S. 2006. Feature sensitive out-of-core chartification of large polygonal meshes. In *Proc. of Computer Graphics International*, 518–529.
- CIGNONI, P., ROCCHINI, C., AND SCOPIGNO, R. 1998. Metro: Measuring error on simplified surfaces. *Computer Graphics Forum* 17, 2, 167–174.
- GARLAND, M., WILLMOTT, A., AND HECKBERT, P. S. 2001. Hierarchical face clustering on polygonal surfaces. In *Proc. 2001 ACM Symposium on Interactive 3D Graphics*, 49–58.
- LLOYD, S. 1982. Least square quantization in pcm. *IEEE Transaction on Information Theory* 28, 129–137.
- PENG, J., KIM, C.-S., AND KUO, C.-C. J. 2005. Technologies for 3D mesh compression: a survey.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, 456–495.
- PULLI, K., AARNIO, T., ROIMELA, K., AND VAARALA, J. 2005. Designing graphics programming interfaces for mobile devices. *IEEE Computer Graphics and Applications* 25, 8, 66–75.
- PURNOMO, B., BILODEAU, J., COHEN, J. D., AND KUMAR, S. 2005. Hardware-compatible vertex compression using quantization and simplification. In *Proc. Graphics Hardware 2005*, 53–61.
- SANDER, P. V., WOOD, Z. J., GORTLER, S. J., SNYDER, J., AND HOPPE, H. 2003. Multi-chart geometry images. In *Proc. Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, 146–155.
- STRÖM, J. AND AKENINE-MÖLLER, T. 2005. iPACKMAN: High-quality, low-complexity texture compression for mobile phones. In *Proc. Graphics Hardware*, 63–70.



Jongseok Lee received the BS and MS degrees in computer science and engineering from Konkuk University and Pohang University of Science and Technology (POSTECH), respectively. After finishing his MS degree, he has been working at LG Electronics as a Junior Research Engineer since 2008. His research interests include computer graphics, mesh compression, and graphics for mobile devices.



Sungyul Choe received the BS, MS, and Ph.D degrees in computer science and engineering from Pohang University of Science and Technology (POSTECH). After finishing his Ph.D degree, he has been working at Samsung Electronics as a Senior Engineer since 2008. His research interests include computer graphics, mesh processing, mesh compression, and graphics for mobile devices.



Seungyong Lee received the BS degree in computer science and statistics from Seoul National University in 1988 and the MS and PhD degrees in computer science from Korea Advanced Institute of Science and Technology (KAIST) in 1990 and 1995, respectively. He is now a professor of computer science and engineering at Pohang University of Science and Technology (POSTECH), Korea. From 1995 to 1996, he worked at the City College of New York as a postdoctoral research associate. Since 1996, he has been a faculty member and leading the Computer Graphics Group at POSTECH. From 2003 to 2004, he spent a sabbatical year at MPI Informatik, Germany, as a visiting senior researcher. His current research interests include image and video processing, nonphotorealistic rendering, 3D mesh processing, 3D surface reconstruction, and graphics applications.