

Making 3D Textures Practical

Chandrajit Bajaj
Department of Computer Sciences
The University of Texas at Austin
U.S.A.
bajaj@cs.utexas.edu

Insung Ihm Sanghun Park
Department of Computer Science
Sogang University
Korea
{ihm, hun}@graphlab.sogang.ac.kr

Abstract

While 2D texture mapping is one of the most powerful rendering techniques that make 3D objects appear visually interesting, it suffers from visual artifacts produced when 2D image patterns are wrapped onto the surface of objects with arbitrary shapes. On the other hand, 3D texture mapping generates highly natural visual effects in which objects appear carved from lumps of materials rather than laminated with thin sheets as in 2D texture mapping. Storing 3D texture images in a table for fast mapping computations, instead of evaluating procedures on the fly, however, has been considered impractical due to the extremely high memory requirements. In this paper, we present a new effective method for 3D texture mapping designed for real-time rendering of polygonal models. Our scheme attempts to resolve the potential texture memory problem arising from the very large size of 3D images by compressing them using a wavelet-based encoding method. The experimental results on various non-trivial 3D textures and polygonal models show that high compression rates are achieved with few visual artifacts in the rendered image and a small impact on rendering time. The simplicity of our compression-based scheme will make it possible to implement practical 3D texture mapping in software/hardware rendering systems including the real-time 3D graphics APIs like OpenGL and Direct3D.

Keywords: Texture mapping, Solid texture, 3D compression, Wavelet, Real-time rendering, OpenGL

1. Introduction

Texture mapping is one of the most powerful rendering techniques that make three-dimensional objects appear visually more complex and realistic [6]. Two-dimensional texture mapping has been popular in creating many interesting visual effects by projecting or wrapping 2D image patterns onto the surface of solid objects. While it has proved

very useful in adding realism in rendering, 2D texture mapping suffers from the limitation that it is often difficult to wrap 2D patterns, without visual artifacts, onto the surface of objects having complicated shapes. As an attempt to alleviate the computational complications of wrapping as well as to resolve the visual artifacts, Peachey [12] and Perlin [13] presented the use of space filling 3D texture images, called *solid textures*. Many of the textures found in nature such as wood, marble, and gases, are easily simulated with solid textures that map three-dimensional object space to color space [3]. Unlike 2D textures, they exist not only on the surface of objects but also inside the objects. Texture colors are assigned to any point of the entire solid object simply by evaluating the specified functions or codes according to their positions in 3D space. The 3D solid texture mapping can be viewed as immersing geometric objects in virtual volumes associated with 3D textures, and getting necessary texture colors from the solid textures. This 3D texture mapping produces highly natural visual effects in which objects appear carved from lumps of materials rather than laminated on the surfaces as in 2D texture mapping. The difference between 2D and 3D mapping is prominent particularly when objects have complicated geometry and topology since 3D textures are not visually affected by the distortions that exist in object parameter space.

Many useful 3D textures are generally synthesized procedurally instead of painting or digitizing them (Refer to [3] for several interesting examples.). They are based on mathematical functions or programs that take 3D coordinates of points as input, and compute their corresponding texture values. The evaluation is usually carried out on the fly during the rendering computation. While procedural texture models provide a very compact representation, evaluating procedural textures as necessary during texture mapping leads to slower rendering than accessing pre-sampled textures stored in simple arrays.

While using sampled 3D texture maps in 3D volumetric form is faster, they tend to take up a large amount of texture memory. For example, when a 3D RGB texture

with resolution $256 \times 256 \times 256$ is represented in one byte per color channel, it requires 48 Mbytes of texture memory. Although some recent graphics systems allow the use of main memory for textures, such texture memory costs are an impossible burden on most current graphics systems. Storing several elaborate textures with high resolution, say, $512 \times 512 \times 512$ would be prohibitive even to the most advanced rendering systems. Obviously, there is a tradeoff between the size of texture memory and the computation time. Explicitly storing sampled textures in dedicated memory, and fetching texture colors as necessary, as in the current graphics accelerator supporting real-time texture mapping, can generate images faster than evaluating them on the fly. To make this feasible for 3D texture mapping, however, efficient solutions for manipulating potentially huge textures effectively need to be devised.

This paper presents a new and practical scheme for real-time 3D texture mapping which is easily implemented. Our technique relies on 3D volume compression and efficient processing of compressed solid textures. The idea of rendering directly from compressed textures has been presented in the past in [2], where they used vector quantization to compress 2D textures in simple or mip-map form. To compress 3D textures, we use a wavelet-based compression method that provides fast decoding to random data access, as well as fairly high compression rates [1]. This compression technique exploits the power of wavelet theory and naturally provides multi-resolution representations of 3D RGB volumes. With this compression method, we can store mip-maps for 3D textures of non-trivial resolutions very compactly in texture memory. Its fast random access decoding ability also results in only a small impact on rendering time. The simplicity of our new 3D texture mapping scheme makes it easy to implement in software/hardware rendering systems. Furthermore, 3D real-time graphics APIs like OpenGL and Direct3D can be extended with little effort to include 3D texture mapping without heavy demand for very large texture memories.

The rest of this paper is organized as follows: In Section 2, we provide a detailed description of the new compression-based 3D texture mapping technique. Experimental results on various 3D textures and polygonal objects are reported in Section 3. Finally, we present conclusions and directions for further research in Section 4.

2. A New 3D Texture Mapping Scheme

In this section, we describe our new 3D texture mapping method suitable for real-time rendering of polygonal models. The idea presented here can also be used effectively in other rendering systems such as RenderMan [14] to enhance the texture mapping speed. The key point in our texture mapping scheme is to extract only the neces-

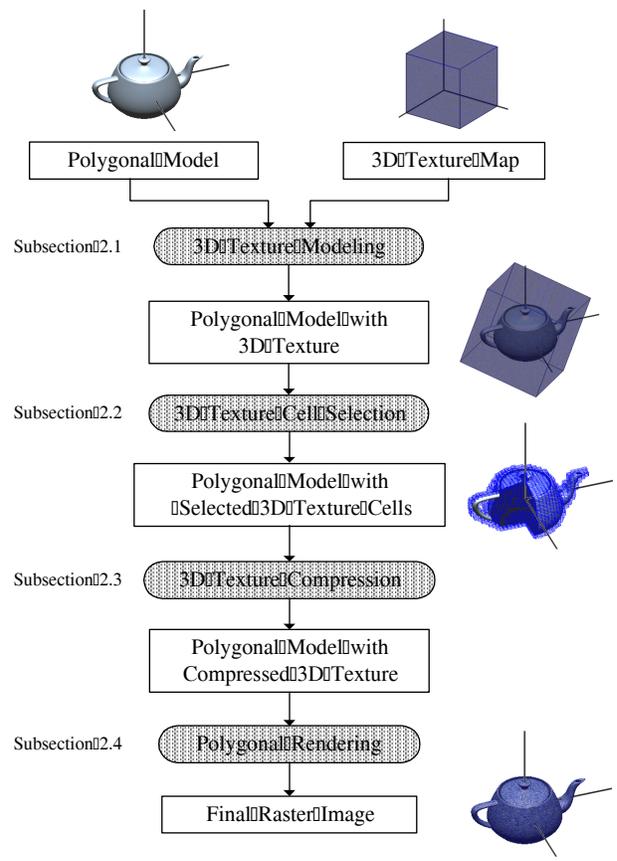


Figure 1. 3D Texture Mapping Pipeline

sary portion from the full 3D texture map, then compress it in compact form where fast run-time decoding for random access to texel values is possible. In particular, the compression method we apply is based on wavelet theory, and naturally supports multi-resolution representations of 3D textures. This capability of the compression method makes it easy to construct a 3D texture mip-map using a small amount of texture memory. Figure 1 illustrates the 3D texture mapping pipeline in which the first three steps, *3D Texture Modeling*, *3D Texture Cell Selection*, and *3D Texture Compression* comprise the necessary pre-processing stages. In the following subsections, we provide detailed explanations of the various stages in the pipeline.

2.1. 3D Texture Modeling

Our scheme assumes, as input texture, a sampled 3D RGB texture stored in a 3D array. It is generated by sampling texel values from a three-dimensional texture field that is usually described procedurally. The storage requirements are very high for uncompressed 3D texture maps at reasonable resolution: 256^3 and 512^3 RGB textures need

48 Mbytes and 384 Mbytes, respectively. This is one of the reasons which make fast 3D texture mapping with stored textures appear impractical.

In the texture modeling stage, a polygonal object in its object space $R_{os} = \{(x, y, z) \mid x, y, \text{ and } z \text{ are real}\}$ is textured by putting it in a 3D texture defined in the texture space $R_{ts} = \{(s, t, r) \mid 0 \leq s, t, r \leq 1\}$, and finding the intersection of the object's surface and the solid texture. Texturing an object can be viewed as determining a function $f : R_{os} \rightarrow R_{ts}$. This function f can be chosen arbitrarily.

2.2. 3D Texture Cell Selection

Once a mapping between a polygonal object and a 3D texture map is fixed, the unnecessary texture data is eliminated to reduce storage space. Consider an $n_s \times n_t \times n_r$ texture. In our scheme, the texture data is subdivided into small subblocks of size $n_c \times n_c \times n_c$, called *texture cells* (In our current implementation, the resolution of texture cell is $4 \times 4 \times 4$). The texture cell is a basic unit for selecting texture data that is actually needed for rendering.

In this 3D texture cell selection stage, each polygon on the boundary of an object is 3D-scan-converted to find all the texture cells that intersect with the surface of the solid object. Notice that texels in the selected texture cells contain all the texture information necessary for rendering. The cells that are not chosen are replaced by null cells, that is, cells with black color. By keeping nearby texels surrounding the surface of an object in this intermediate stage, a large portion of texture data is removed to alleviate the potential prohibitive storage requirement. The selected texture cells take only a small percentage of the original texture data. The null cells still exist in the texture map in this stage, and the texture size remains the same. However, the spatial coherence created by null cells makes an encoding scheme efficiently compress the 3D texture in compact form in the next stage.

2.3. 3D Texture Compression

2.3.1. Choosing an Appropriate Compression Technique

There exist many data compression methods for efficient storage and transmission. It is very important to choose a compression technique which is most appropriate for this specific 3D texture mapping application. We have several issues to consider as similarly discussed in [2, 9]:

1. **High compression rate and visual fidelity.** Non-trivial 3D textures are often very large in size, ranging from a few dozen megabytes to several hundred megabytes. When a mip-map is used for a pre-filtered multi-resolution representation, the size gets even larger. Developing real-time applications with

such data assumes, implicitly or explicitly, that the entire data is loaded into main memory for efficient run-time processing. This places an enormous burden on storage space as well as transmission bandwidth. While lossless compression techniques preserve data without introducing reconstruction errors, they often fail to achieve compression rates high enough for practical implementation of 3D texture mapping. The loss of information associated with lossy compression methods, however, needs to be controlled properly as it is important to minimize the distortion in the reconstructed textures.

2. **Fast decoding for random access.** The general concern of most lossy compression schemes is achieving the best compression rate with minimal distortion in the reconstructed images [5, 15]. Such compression methods, however, often impose constraints on the random access decoding ability, which makes them inappropriate for interactive texture mapping applications especially where it is difficult to predict data access patterns in advance. For instance, variable-bitrate or differential encoding schemes such as Huffman or arithmetic coders coupled to block JPEG or MPEG schemes, do not lend themselves to efficiently decode individual texels that are accessed in a random pattern during run-time.
3. **Multi-resolution representation.** Mip-mapping is the most commonly used anti-aliasing technique for 2D texture mapping [18]. A mip-map of a 2D texture is a pyramid of pre-filtered images obtained by averaging down the original image to successively lower resolutions. Mip-mapping with this level-of-detail representations of textures offers fast and constant filtering of texels, and its simplicity lends itself to an efficient hardware implementation. The idea naturally extends to 3D textures although they have regarded it as impractical due to the memory requirements. It is highly recommended to choose a compression technique that provides a multi-resolution representation in its compression scheme.
4. **Exploitation of 3D data redundancy.** 3D textures are three-dimensional data that exhibit redundancy in all three dimensions. A compression scheme devised for 2D images could be applied to compress each slice in 3D textures, however, a good compression technique must be able to fully exploit data coherence in all three dimensions to maximize the compression performance.
5. **Selective block-wise compression.** In some applications like ours, it is more efficient to selectively

compress a certain portion of data rather than the entire dataset. It is very desirable that a compression scheme includes this selective compression capability in its encoding algorithm for effective compression.

2.3.2. The Zerobit Encoding Scheme

The above five desirable characteristics are common to most real-time applications that must handle discrete sampled data of very large sizes. Vector quantization has been popular in developing such applications mainly because it supports fast random decoding through table lookups [4]. Some recent applications of vector quantization in the computer graphics field, include compression of CT/MRI datasets [10], light fields [9], and 2D textures [2].

Recently, a new compression scheme for 3D RGB image has been developed as an alternative to vector quantization [1]. This technique, called *zerobit encoding*, is suitable for applications wherein data are accessed in an unpredictable manner, and real-time performance of decoding is required. It extends the compression scheme [7, 8] for 3D gray-scale volume data to compression of 3D RGB images, and its new encoding structure significantly improves decompression speeds. Unlike vector quantization, the zerobit encoding scheme, based on the wavelet theory, naturally offers a multi-resolution representation for 3D images. Experimental results on test data sets show that this compression scheme provides fast random access to compressed data in addition to achieving fairly high compression rates. The basic compression unit for the technique is a $4 \times 4 \times 4$ subimage, called *cell*, to which the 3D Haar transform is applied twice to exploit data coherence in all of the three dimensions. Once compressed, three levels of detail are represented in compressed form. Notice that the texture cell in our 3D texture mapping scheme naturally corresponds to the cell in this compression technique.

Figure 2 shows sample statistics on the performance of the zerobit encoding and vector quantization used in [9] for two representative light field datasets *buddha* and *dragon* with resolution $32 \times 32 \times 256 \times 256$ (192MBytes) [1]. To apply the zerobit encoding technique, 4D sampled light field datasets were reformulated into 3D images, then compressed. While the vector quantization yielded compression rates 21.79 and 20.18 for *buddha* and *dragon*, the zerobit encoding method produced higher rates of 44.51 to 91.11 and 38.21 to 83.03 at the selected four target ratios, respectively (Figure 2(a))¹. The PSNR results show that the qualities of reconstructed images are about the same when about 2% and 5% of coefficients are used in zerobit encoding for the *buddha* and *dragon* datasets, respectively (See Figure 2(c) and (d) for the portion of two sample *buddha*

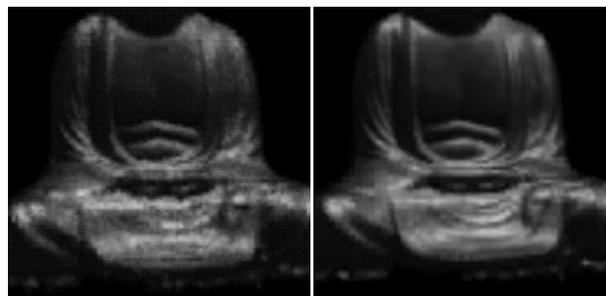
¹These rates exclude the gzip compression, that could follow both compression methods for efficient storage as in [9].

		Vector Quantization	Zerobit Encoding			
			2.0%	3.0%	4.0%	5.0%
buddha	Size (MB)	8.81	2.11	2.90	3.63	4.31
	Comp. Rate	21.79	91.11	66.26	52.89	44.51
	PSNR (dB)	38.00	39.26	41.70	43.63	45.18
dragon	Size (MB)	9.52	2.31	3.15	4.09	5.02
	Comp. Rate	20.18	83.03	60.87	46.99	38.21
	PSNR (dB)	35.58	31.00	32.17	33.37	34.40

(a) Compression Rate and Quality

		Vector Quantization	Zerobit Encoding			
			2.0%	3.0%	4.0%	5.0%
buddha	st-lerp	9.46	13.60	13.60	13.60	13.60
	uvst-lerp	2.68	2.99	2.98	2.98	2.98
dragon	st-lerp	17.55	24.60	24.44	24.20	23.97
	uvst-lerp	5.66	5.74	5.71	5.66	5.62

(b) Rendering Time (Frames Per Second)



(c) Vector Quantization

(d) Zerobit Encoding: 3% of Wavelet Coefficients Used

Figure 2. Comparisons with Vector Quantization on Light Field Datasets [1]

images.).

The image-based rendering time, spent on displaying 76 frames of 382×382 pixels with gradually varying viewing parameters, was measured on an SGI workstation with a 195 MHz MIPS R10000 CPU. Two cases of bilinear interpolation on the *st*-plane (*st-lerp*) and quadrilinear interpolation on both *uv*- and *st*-planes (*uvst-lerp*) were tested (Figure 2(b)). The timing results show the zerobit encoding scheme is faster for both datasets in most cases. Note that the reconstruction cost per data item for vector quantization is very cheap since decompression is performed through a simple codebook lookup, and is cheaper than zerobit encoding on the average. However, zerobit encoding decompresses several data items, 4 planes in this case, at the same time, and is very quick particularly when data in empty background regions are reconstructed, which results in overall faster rendering. When nearest samples without interpolation are taken during image-based rendering, the zerobit encoding method yielded frame rates 39 and 52 for

the two datasets, respectively.

While the empirical comparisons for a few applications cannot prove that the zerobit encoding is always superior to vector quantization, we find the former compares very favorably to the latter. In our 3D texture mapping technique, we use the zerobit encoding scheme to compress the selected texture cells. As will be explained in the next section, it also turns out to be very effective in compressing 3D textures.

2.4. Polygonal Rendering with Compressed Textures

2.4.1. A New Capability for OpenGL 1.2

When applying textures to geometric objects, the necessary texel values are repeatedly fetched from zerobit-encoded 3D textures using their texture coordinates. This compression-based 3D texture mapping can enhance the rendering speed in any rendering method including time-consuming photo-realistic rendering. In our implementation, we applied our scheme to real-time rendering and extended the OpenGL library to include the feature of 3D texture mapping with zerobit-encoded textures. Note that 3D texture mapping has been implemented as a feature of the OpenGL 1.1 extensions, and is now one of the core capabilities that must be supported by all OpenGL 1.2 implementations [16]. The `glTexImage1D()` and `glTexImage2D()` functions are extended for 3D texture mapping where the command for specifying a three-dimensional texture image is defined as

```
void glTexImage3D (GLenum target, GLint
  level, GLint internalformat, GLsizei width,
  GLsizei height, GLsizei depth, GLint
  border, GLenum format, GLenum type, void
  *data);
```

With *target* `GL_TEXTURE_3D`, this command reads a texture of size $width \times height \times depth$, that is stored in memory, pointed by *data*, in *internalformat*. For a compressed texture, our extension uses `GL_UNSIGNED_BYTE_COMPRESSED` for *type* to read a compressed texture, whose texels are stored in unsigned character, on levels *level*, *level+1*, and *level+2*.

When 3D texture mapping is enabled by calling `glEnable(GL_TEXTURE_3D)`, and a compressed 3D texture is specified, the texture is assumed to be in compressed form, and texels are fetched from the zerobit-encoded structure rather than a simple array. The extension is easy to implement since the new capability can be included simply by adding proper state variables and decoding functions. Other utility functions, such as creating encoded 3D textures with user-specified compression rates, could also be included in the OpenGL Utility Library (GLU).

2.4.2. Compact Representations of 3D Mip-Maps

A 3D mip-map is an ordered set of 3D arrays representing the same texture where each successive array has a resolution lower than its previous one. Currently, we are implementing the 3D mip-map capability so that mip-maps as well as single 3D textures are represented very compactly. Given a base 3D texture, the zerobit-encoded structure represents three levels of detail with level number 0, 1, and 2. The reduced images on the next three levels can be stored in another zerobit-encoded structure. An alternative is to store the texture images with lower resolutions except on level 0, 1, and 2, in simple 3D arrays. The images on the higher levels take up only a small amount of storage. For example, when a $256 \times 256 \times 256$ RGB texture image in unsigned character is loaded, the entire reduced images on levels 3, 4, \dots , 8 require only about 110 Kbytes in total.

2.5. Sharing of a 3D Texture between Multiple Objects

In our current implementation, only the 3D texture cells surrounding the surface of a polygonal object are selected and compressed. When a texture is compressed object by object, it could lead to a waste of texture memory. That is, if a 3D image is shared by multiple polygonal objects, the same 3D texture cells can be replicated for several objects. We are extending our method to support three types of compression modes: The first mode, called `zerobit_encoding_single_object` is one we have described in this article. The second mode `zerobit_encoding_multiple_objects` is for the case in which several polygonal objects share a common 3D texture image. In this mode, all the 3D texture cells that are used by at least one object are selected before encoding. The last mode `zerobit_encoding_entire_texture` handles the dynamic situation in which it is difficult or impossible to predict which texture cells shall be used for rendering. For instance, an interesting animation can be generated by making an object float in a texture field, dynamically binding texture coordinates. In this case, the first two compression modes are not appropriate. The third mode compresses the entire 3D texture and loads it for rendering. While it is the most expensive one, this mode provides a flexibility in texture mapping.

3. Experimental Results

We have implemented our new 3D texture mapping scheme by extending the MESA 3D Graphics Library which is a publicly available OpenGL implementation [11]. The current version 3.0 supports the 3D texture mapping feature where the entire texture image is stored in a simple array without any compression. We added the necessary

Object	# of Faces	# of Selected Cells	Ratio (%)
Teapot	1,152	6,761	2.6
Dragon	12,078	7,083	2.7
Bunny	69,451	14,551	5.6
Head	203,544	31,122	11.9

Table 1. Ratios of Selected Texture Cells

state variables and functions to handle zerobit-encoded 3D texture maps.

We have generated four different 3D texture images of size $256 \times 256 \times 256$ (Figure 3). The texture images have three channel RGB colors, and their sizes amount to 48 Mbytes, respectively. The three textures **Bmarble**, **Wood**, and **Eroded** were created using the RenderMan surface shaders `blue_marble()`, `wood()`, and `eroded()`, respectively [17]. The surface shader `gmarbtpolish()` for the texture **Gmarbtpol** were written by Larry Gritz, and is available as a part of the Blue Moon Rendering Tools (BMRT). Our 3D texture mapping technique has been applied to several polygonal models with various shapes and sizes, including those listed in Table 1. The teapot model **Teapot** was polygonized from a parametric equation. The next two models **Dragon** and **Bunny** were obtained from Viewpoint and the Stanford 3D Scanning Repository, respectively. The last model **Head** was created by generating an iso-surface from the UNC CT scan of a human head. The table shows how many $4 \times 4 \times 4$ texture cells are selected from the entire $262,144 (= 64 \times 64 \times 64)$ cells in $256 \times 256 \times 256$ textures through the 3D texture cell selection stage after texture modeling. We observe that the ratios of selected cells are quite small.

To find out how compactly these 3D textures can be associated with the polygonal objects, we compressed selected texture cells for the entire 16 combinations as shown in Table 2. In the zerobit encoding scheme, a user specifies a ratio of wavelet coefficients to be used after truncation to control the degree of compression [1]. The number, shown in the “Target Ratio” field of Table 2, represents an approximate ratio of wavelet coefficients that are actually used in encoding. We compressed 3D textures at three target ratios 3%, 5%, and 10%, and rendered the polygonal objects from the compressed textures. In this table, we compare sizes and compression rates for various cases. Observe that it took less than 1 Mbytes of memory for all combinations, ranging from 172 Kbytes to 804 Kbytes. Considering that the size of the original textures is 48 Mbytes, we see that very high compression rates are indeed achieved through zerobit encoding.

Figure 4 shows sample images rendered with the linear filter `GL_LINEAR` from the compressed textures having a desired ratio of 10%. When the 3D textures are compressed with a desired ratio higher than 10%, the texture-mapped

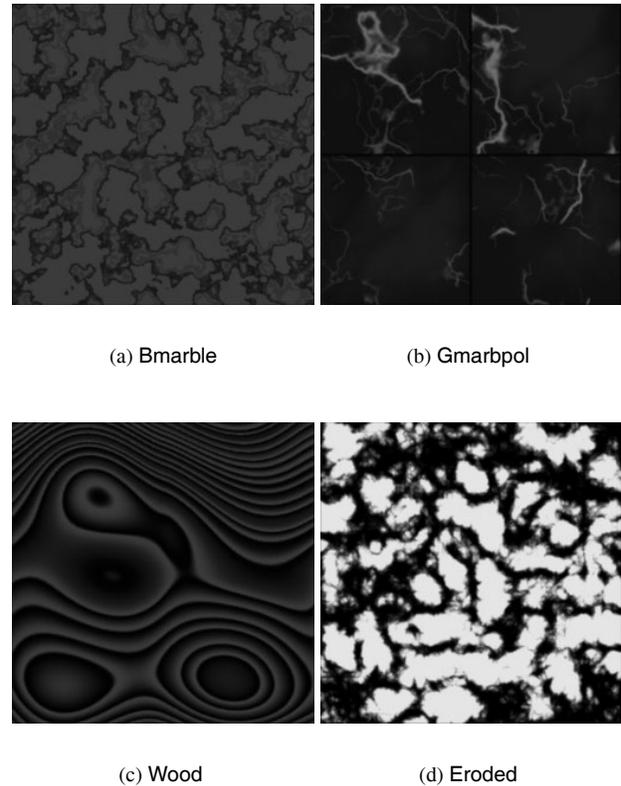


Figure 3. Sample Slices from the Four Example 3D Textures

images, produced with the linear filter, are almost free of aliasing artifacts which are often caused by the loss of information during lossy compression. In Figure 5, we enlarged a portion of the **Bunny** images to make the compression artifacts more visible. When the ratio is 3%, the blocky artifacts are clearly visible, but most features are still preserved well enough for many real-time applications such as 3D games and animation.

In order to check the timing performance, we measured the running time, spent on rendering 54 frames of 512×512 pixels with incrementally varying viewing parameters. They include all computations for rendering including 3D texture mapping, view parameter setting, and displaying the final images. The timings were measured on an SGI Octane workstation with a 195 MHz R10000 CPU and 256 Mbytes of memory without hardware graphics acceleration. Table 3 reports the average time per frame in seconds for three difference rendering modes. The “GSO” field in this table is the time taken for rendering the objects using Gouraud shading only, and indicates how complicated is the involved rendering. Then, our new compression-based texturing scheme was compared with texture map-

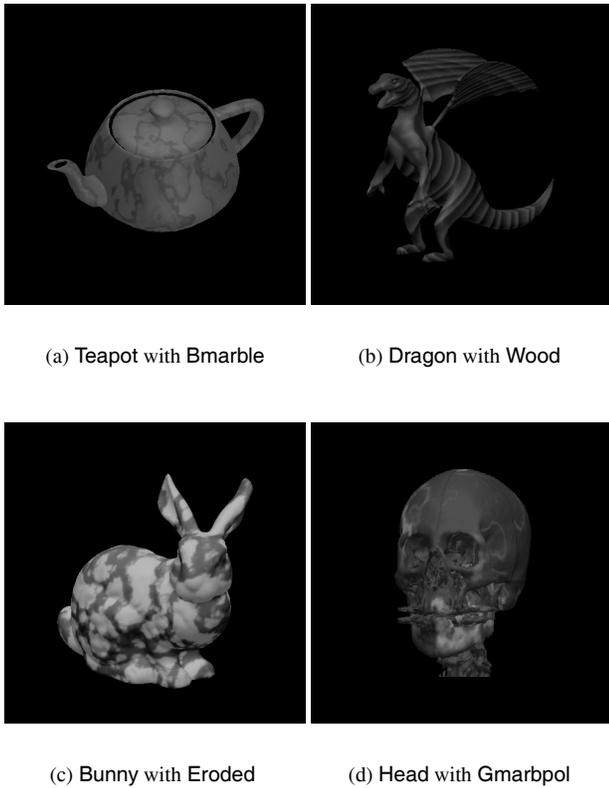


Figure 4. Images Rendered with `GL_LINEAR` from Compressed Textures (10%)

ping without compression to evaluate overheads for fetching texels from compressed textures. Two filtering methods `GL_NEAREST` and `GL_LINEAR` were tested whose performances are shown in the “3DTMN” and “3DTML” fields, respectively. The running time is proportional to the number of pixels that objects are projected into. As indicated by the test results, the zerobit encoding method provides very fast decoding speeds. We observe only a 14 percent and a 15 percent impact on rendering time on average for the nearest and the linear filter, respectively. Notice that the linear filtering method takes, for instance, 0.37 second to render **Teapot** from its uncompressed texture of size 48 Mbytes. On the other hand, the same filtering takes 0.44 second to produce a **Teapot** image with few visual artifacts from its compressed texture of size 268 KBytes (target ratio = 10%). The benefit from our compression-based 3D texture mapping is evident, and is critical in particular when texture memory resource is rather limited.

We have also generated two more elaborate textures of $512 \times 512 \times 512$ whose sizes are 384 Mbytes, and tested our texture mapping scheme with these huge textures (Table 4). The experiments indicate that 500 Kbytes to 1.61 Mbytes

Texture	Object	Target Ratio	Size (KB)	Comp. Rate
Bmarble 256^3	Teapot	3%	188	261.5 : 1
		5%	224	219.4 : 1
		10%	268	183.4 : 1
	Dragon	3%	188	261.5 : 1
		5%	216	227.6 : 1
		10%	272	180.7 : 1
	Bunny	3%	272	180.7 : 1
		5%	340	144.6 : 1
		10%	468	105.0 : 1
	Head	3%	380	129.4 : 1
		5%	488	100.7 : 1
		10%	712	69.0 : 1
Gmarbpol 256^3	Teapot	3%	188	261.4 : 1
		5%	204	240.9 : 1
		10%	232	211.9 : 1
	Dragon	3%	172	285.8 : 1
		5%	192	256.0 : 1
		10%	228	215.6 : 1
	Bunny	3%	244	201.4 : 1
		5%	284	173.1 : 1
		10%	336	146.3 : 1
	Head	3%	332	148.1 : 1
		5%	420	117.0 : 1
		10%	540	91.0 : 1
Wood 256^3	Teapot	3%	192	256.0 : 1
		5%	224	219.4 : 1
		10%	304	161.7 : 1
	Dragon	3%	192	256.0 : 1
		5%	232	211.9 : 1
		10%	308	159.6 : 1
	Bunny	3%	288	170.7 : 1
		5%	372	132.1 : 1
		10%	524	93.8 : 1
	Head	3%	388	126.7 : 1
		5%	524	93.8 : 1
		10%	804	61.1 : 1
Eroded 256^3	Teapot	3%	188	261.5 : 1
		5%	224	219.4 : 1
		10%	288	170.7 : 1
	Dragon	3%	192	256.0 : 1
		5%	232	211.9 : 1
		10%	288	170.7 : 1
	Bunny	3%	280	175.5 : 1
		5%	356	138.1 : 1
		10%	492	99.9 : 1
	Head	3%	384	128.0 : 1
		5%	516	95.3 : 1
		10%	768	64.0 : 1

Table 2. Sizes of Compressed Textures

of memory are required to store the textures compressed at the target ratios 3%, 5%, and 10%, achieving compression rates of 238.0 : 1 to 786.4 : 1. Compared to the 256^3 textures, compression-based renderings take 1.26 (**Teapot** with **Bmarble**) and 1.09 (**Head** with **Gmarbpol**) times as long on the average for the 512^3 textures. We were not able to load the entire uncompressed textures for rendering onto our workstation with 256 Mbytes of main memory, but expect that the rendering times will also get slower at the same rate.

Figure 6 makes a comparison between renderings with four different texture mapping parameters. When **Teapot** is rendered from the 512^3 texture with a target ratio of 10% and the linear filter ((b)), the texture pattern on the sur-

Object & Texture	Target Ratio	GSO	3DTMN	3DTML
Teapot with Bmarble	uncomp.	0.05	0.13	0.37
	3%	–	0.14	0.42
	5%	–	0.15	0.43
	10%	–	0.16	0.44
Dragon with Wood	uncomp.	0.26	0.45	0.89
	3%	–	0.50	0.98
	5%	–	0.52	1.00
	10%	–	0.55	1.04
Bunny with Eroded	uncomp.	1.03	1.39	1.77
	3%	–	1.56	2.04
	5%	–	1.60	2.13
	10%	–	1.66	2.21
Head with Gmarbpol	uncomp.	2.87	3.68	4.51
	3%	–	3.89	4.85
	5%	–	3.94	4.90
	10%	–	4.00	5.03

Table 3. Rendering Time Per Frame (Seconds): GSO - Gouraud Shading Only, 3DTMN - 3D Texture Mapping (Nearest), 3DTML - 3D Texture Mapping (Linear)

face appears much clearer than in the image, produced from the uncompressed 256^3 texture with the same filter ((a)). When the faster but worse nearest filter is applied to the 512^3 texture with a target ratio 5% or 10% ((d)), consuming 0.20 second and 588 Kbytes (5%), and 0.23 second and 752 Kbytes (10%), the qualities are superior to the case in which the slower but better linear filter is applied to the 256^3 texture with a target ratio 10%, requiring 0.44 second and 268 Kbytes. Obviously, there is a tradeoff between rendering time, image quality, and memory requirement, and a choice of various texture mapping parameters should be made to optimize the application’s needs.

4. Concluding Remarks

In this paper, we have presented a very effective method for 3D texture mapping, designed for real-time rendering of polygonal models. Our scheme attempts to resolve the potential texture memory problem arising from the very large sizes of 3D images by compressing them using the zero-bit encoding scheme. This compression scheme not only provides fairly high compression rates but also offers very fast random access to individual texels. The experimental results on various non-trivial 3D textures and polygonal objects show that high compression rates are achieved with a small impact on rendering time and few visual artifacts in the rendered images. The simplicity of our compression-based 3D texture mapping scheme will make it easy to implement in software/hardware rendering systems. Cur-

Object & Texture	Target Ratio	Size (KB)	Comp. Rate
Teapot with Bmarble	3%	500	786.4 : 1
	5%	588	668.7 : 1
	10%	752	522.9 : 1
Head with Gmarbpol	3%	1092	360.1 : 1
	5%	1316	298.8 : 1
	10%	1652	238.0 : 1

(a) Size of Compressed Texture

Object & Texture	Target Ratio	GSO	3DTMN	3DTML
Teapot with Bluemarble	uncomp.	0.05	–	–
	3%	–	0.18	0.48
	5%	–	0.20	0.50
	10%	–	0.23	0.54
Head with Gmarbpol	uncomp.	2.87	–	–
	3%	–	4.13	5.26
	5%	–	4.23	5.30
	10%	–	4.37	5.62

(b) Rendering Time Per Frame (Seconds): GSO - Gouraud Shading Only, 3DTMN - 3D Texture Mapping (Nearest), 3DTML - 3D Texture Mapping (Linear)

Table 4. Experimental Results on $512 \times 512 \times 512$ Textures

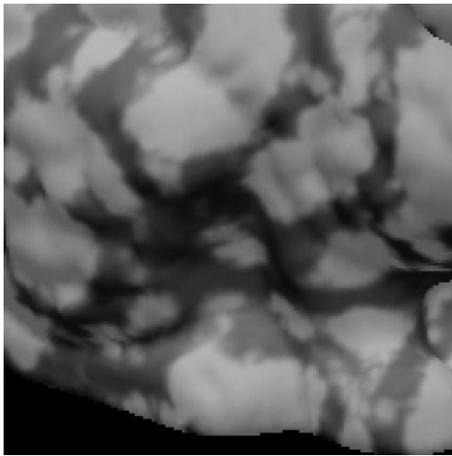
rently, we are implementing a three-dimensional version of mip-mapping for anti-aliasing of solid textures. Once an effective 3D mip-mapping technique is developed, 3D real-time graphics APIs like OpenGL and Direct3D will be extended with little effort to include 3D texture mapping without heavy demand for texture memories.

Acknowledgements

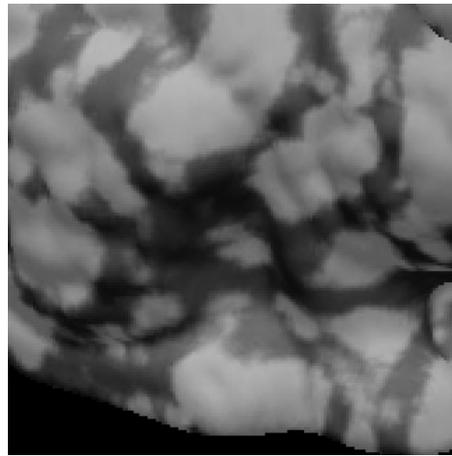
We would like to thank Kiju Park and Joongyeon Lee for their help with experiments. The MESA 3D Graphics Library is an OpenGL implementation written by Brian Paul. We wish to thank the Stanford Graphics Lab., the UNC Graphics Lab., Viewpoint, and Larry Gritz for their data and codes.

References

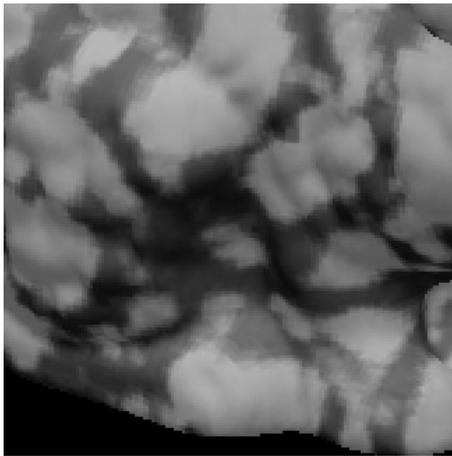
- [1] C. Bajaj, I. Ihm, and S. Park. 3D RGB image compression for interactive applications. *Submitted for publication*, March 1999.



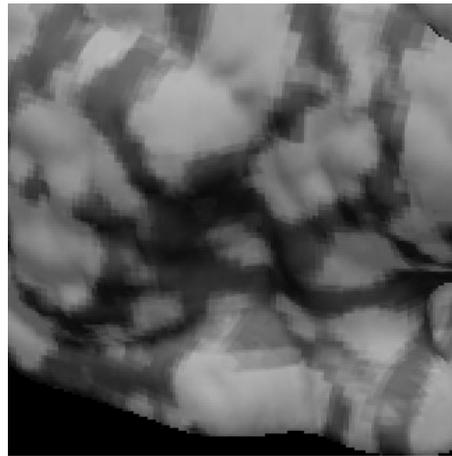
(a) Uncompressed



(b) Compressed (10%)



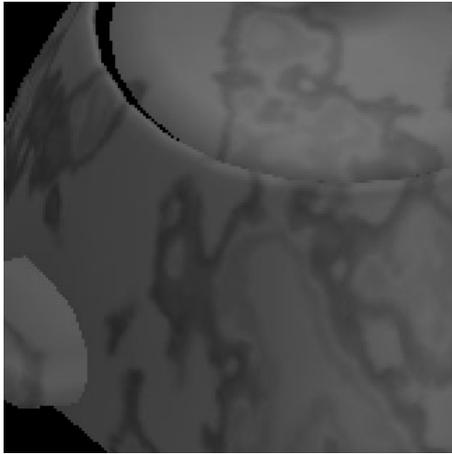
(c) Compressed (5%)



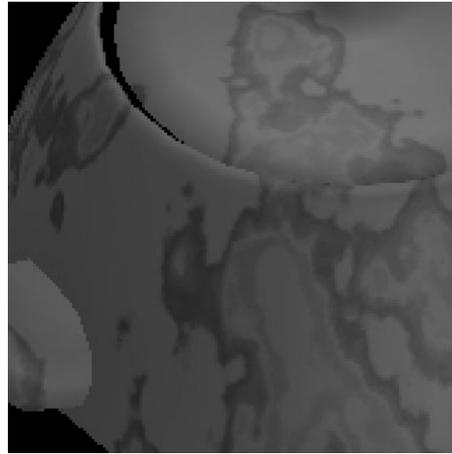
(d) Compressed (3%)

Figure 5. Aliasing Artifacts of Compression-Based 3D Texture Mapping

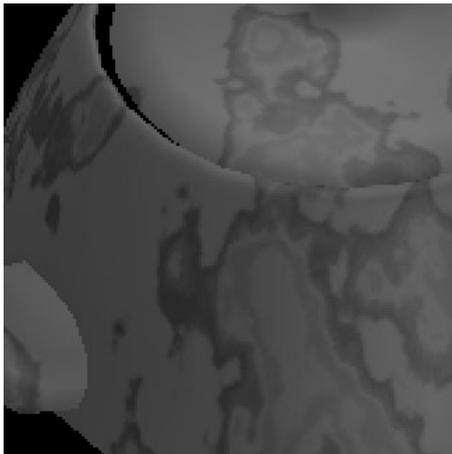
- [2] A. Beers, M. Agrawala, and N. Chaddha. Rendering from compressed texture. *Computer Graphics (Proc. SIGGRAPH '96)*, pages 373–378, 1996.
- [3] D. Ebert (Editor), F. Musgrave, D. Peachey, K. Perlin, and S. Worley. *Texturing and Modeling: A Procedural Approach*. AP Professional, 1994.
- [4] A. Gersho and R. Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, 1992.
- [5] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison-Wesley, 1993.
- [6] P. Heckbert. Survey of texture mapping. *IEEE Computer Graphics and Applications*, 6(11):56–67, 1986.
- [7] I. Ihm and S. Park. Wavelet-based 3D compression scheme for very large volume data. In *Proceedings of Graphics Interface '98*, pages 107–116, Vancouver, Canada, June 1998.
- [8] I. Ihm and S. Park. Wavelet-based 3D compression scheme for interactive visualization of very large volume data. *Computer Graphics Forum*, 18(1):3–15, 1999.
- [9] M. Levoy and P. Hanrahan. Light field rendering. *Computer Graphics (Proc. SIGGRAPH '96)*, pages 31–42, 1996.
- [10] P. Ning and L. Hesselink. Fast volume rendering of compressed data. In *Proceedings of Visualization '93*, pages 11–18, San Jose, October 1993.
- [11] B. Paul. *The Mesa 3D Graphics Library*. [http : //www.mesa3d.org](http://www.mesa3d.org), 1999.
- [12] D. Peachey. Solid texturing of complex surfaces. *Computer Graphics (Proc. SIGGRAPH '85)*, 19(3):279–286, 1985.
- [13] K. Perlin. An image synthesizer. *Computer Graphics (Proc. SIGGRAPH '85)*, 19(3):287–296, 1985.
- [14] Pixar. *The RenderMan[®] Interface (Version 3.1)*, Sept. 1989.
- [15] K. Sayood. *Introduction to Data Compression*. Morgan Kaufmann Publishers, Inc., 1996.



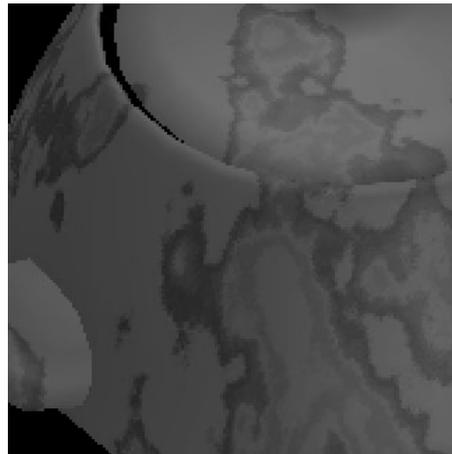
(a) 256^3 (Uncompressed & Linear)



(b) 512^3 (10% & Linear)



(c) 512^3 (3% & Linear)



(d) 512^3 (10% & Nearest)

Figure 6. Comparison between Renderings with 256^3 and 512^3 Textures

- [16] M. Segal and K. Akeley. *The OpenGL[®] Graphics System: A Specification (Version 1.2)*. Silicon Graphics, Inc., March 1998.
- [17] S. Upstill. *The RenderMan[™] Companion*. Addison-Wesley, 1990.
- [18] L. Williams. Pyramidal parametrics. *Computer Graphics (Proc. SIGGRAPH '83)*, 17(3):1–11, 1983.