

Rendering Complex Scenes Based on Spatial Subdivision, Object-Based Depth Mesh, and Occlusion Culling

Bo-Yin Lee Chih-Chun Chen Jung-Hong Chuang Wei-Wen Feng Ting Chiou
Jin-Bey Yu

Department of Computer Science and Information Engineering
National Chiao Tung University
1001 Ta Hsueh Road, Hsinchu, Taiwan 30050, ROC.
{bylee, chihchun, jhchuang, fengww, tchiou, jbyu}@csie.nctu.edu.tw

Abstract

In this paper, we combine geometry-based and image-based rendering techniques to develop a VR navigation system that aims to have efficiency relatively independent of the scene complexity. The system has two stages. In the preprocessing stage, the x - y plane of a 3D scene is partitioned into equal-sized hexagonal navigation cells. Then for each navigation cell, we associate each object outside the cell with either a LOD mesh or an object-based depth mesh depending on its self-occlusion error. The object with the error larger than a user-specified threshold, is associated with a LOD mesh of an appropriate resolution. For the object with error smaller than the threshold, we associated it with a depth mesh that is reduced from its original mesh based on the silhouette and depth information of its image rendered from the cell center. All LOD meshes are then culled by a conservative back-face computation, and then all LOD and depth meshes are culled by a conservative visibility computation, both aim to remove polygons that are invisible from any point inside the cell. At run-time stage, LOD meshes are rendered normally while depth meshes are rendered by texture mapping with their cached images. Techniques for run-time back-face culling and occlusion culling can be easily included. Our experimental results have depicted fast frame rates for complex environments with acceptable quality-loss.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing Algorithms;

Keywords: Virtual Reality, Visibility, Image Caching, Image-Based Rendering, Hybrid Rendering

1 INTRODUCTION

In order to achieve an immersive visual effect during the VR navigation, rendering with photo-realistic scene images and high frame rate has been our ultimate goal. However, in the traditional geometry-based rendering, complex scenes always require numerous polygons, and, therefore, can be rendered at an unacceptably low frame rate even using a state-of-the-art hardware. Many techniques have been proposed in last decades on reducing the polygon count while preserving the visual realism of complex scenes, including visibility culling, level-of-detail (LOD) modeling, and more recently, image-based rendering. Although image-based rendering scheme is capable of rendering complex scenes with photo-realistic images in the time complexity that is independent of the scene complex-

ity, it has been suffered from the limited viewing degree of freedom and some losses of image quality due to gaps and holes. Hybrid rendering that combines geometry- and image-based technique has become a viable alternative.

To take the advantages of both geometry- and image-based rendering techniques, we introduce a hybrid rendering scheme that aims to render a complex scene in a constant and high frame rate with only a little or an acceptable quality loss. The hybrid scheme consists of two stages: pre-processing stage and run-time stage. In the pre-processing stage, to exploit the spatial coherence, the x - y plane of a 3D scene is partitioned into equal-sized hexagonal navigation cells. To reduce hole problem due to self-occluding, each object outside a cell is represented either by a LOD mesh or by a *depth mesh* depending on its approximate *self-occluding error*. The object is represented by a LOD mesh of an appropriate resolution if its self-occluding error is over a user-specified tolerance; otherwise by an object-based *depth mesh*. The object-based *depth mesh* is derived from the object's original mesh based on the silhouette and depth variation on the rendered image viewed from the cell center. The resulting depth mesh is a view-dependent LOD model of the object's visible part that the resolution becomes coarser when the object is more distant while the silhouette is well preserved. In consequence, for each navigation cell, we have a set of LOD meshes and depth meshes (together with their cached images) for objects outside the cell. Since viewpoint is constrained in a particular cell during the navigation, we can further remove back-facing polygons of LOD meshes and occluded polygons for both type of meshes for any viewpoint inside the cell. This is accomplished by the proposed conservative back-facing culling and conservative occlusion culling. Based on our test, in average, conservative back-facing culling removes about 30 percent of polygons in LOD meshes.

At the run-time stage, LOD meshes and depth meshes associated with the navigation cell where current viewpoint lies are rendered. The depth meshes are texture mapped with the cached images by hardware accelerated *projective-alike texture mapping* which generates texture coordinates automatically. Run-time occlusion culling for the entire scene and back-facing culling for the objects inside the cell can be performed to further reduce the polygon count. To minimize the impact of the data loading while navigating across the cell boundary, a pre-fetching scheme is also developed to amortize the loading time to several previous frames.

2 RELATED WORK

View frustum culling that prevents the objects outside the view volume from being sent to the rendering pipeline is the basic visibility culling. Furthermore, back-face culling is also commonly used. In [15], a sub-linear algorithm has been proposed for culling back-facing polygons, but it suffers from the requirements of model partitioning. Zhang and et al. improved this by introducing *normal mask* which reduces the per polygon back-face test to only one logical AND operation [29].

Several run-time methods have been proposed that further cull out polygons that are occluded by others. There include occlusion culling using shadow frusta [14], hierarchical Z-buffer [10], and hierarchical occlusion map [30]. However, there are inevitable overheads doing occlusion culling at run-time. Cohen-Or and et al. [3] proposed a pre-processing algorithm for regional occlusion culling, but its performance depends heavily on a single strong occluder. Durand and et al. [7] proposed *extended projection* operations and Schaufler and et al. [20] proposed *blocker extension* to handle occluder fusion of multiple occluders. They both subdivide a scene into volumetric cell, and computes potential visibility set (PVS) for each volumetric cell.

With those visibility culling techniques, the remaining polygons might be still too many to achieve interactive rate. Level-of-detail (LOD) modeling has been very useful in further reducing the number of polygon that are visible and inside the view frustum. Distant objects get projected to small areas on the screen and hence can be represented with coarse meshes. On the other hand, nearby objects share larger screen areas and should be modeled by meshes of higher resolution. Many methods have been proposed to obtain LOD meshes; e.g., *edge collapsing* [13], *vertex clustering* [19], *vertex decimation* [22], *progressive mesh* [11, 28], and *view dependent LOD* [27, 12].

Geometry-based rendering based on visibility culling and LOD modeling alone usually still cannot meet interactive requirement for very complex scenes. Image-based rendering (IBR) has been a well known alternative. IBR takes parallax into account, and renders a scene by interpolating neighboring reference views [2, 1, 18, 23]. IBR has efficiency that is independent of the scene complexity, and can model natural scenes using photographs. It is, however, often constrained by the limited viewing degree of freedom. IBR in general has problems like folding, gap, and hole. Lumigraph [9] and light field rendering [16] have been proposed to reduce the *7D plenoptic function* $P(\theta, \phi, \lambda, V_x, V_y, V_z, t)$ to the *4D function* $P'(s, t, u, v)$ for static scenes. However, both require enormous storage for the extremely large number of images. Layered depth image (LDI) [25] is a good try to eliminate hole problems due to the visibility changes. LDI structure is more compact in the sense that redundant information has been reduced when several neighboring reference images are composed into a single LDI.

Hierarchical image caching proposed in [21, 24] is the first approach that combines geometry-based and image-based rendering aiming to achieve an interactive frame rate for complex static scenes. The cached texture possesses no depth information and, in turns, limits its life cycle. The image simplification schemes proposed in [5, 4, 26] represent background or distant scene using 2D cached depth meshes derived from the rendered images for some specific views. Such depth meshes are rendered by re-projection and texture mapping. In such approaches, folding prob-

lems and gaps resulting from the resolution changes can be eliminated; however, the hole problems due to visibility and self-occluding still remain. Moreover, disjointed objects might be rendered as connected objects, and depth meshes derived on the 2D cached images are in pixel-resolution, which might lead to geometric inaccuracy when re-projected into 3D space. In [6] proposed multi-layered impostors to restrict visibility artifacts to a given size, and as well as a dynamic update scheme to improve the resolution mismatch.

3 PROPOSED HYBRID SCHEME

A scene represented by traditional image-based representation would produce a lot of artifacts such as cracks and holes due to resolution mismatch and visibility change respectively. However, a scene represented by a single environment-based depth mesh (impostor) still produces rubber artifacts caused by incorrect connections between disjoint surfaces. Although, disjoint surfaces are identified, holes appear in a new view caused by the lack of information of occluded parts. Such that, we propose an *object-based depth meshes* scheme to greatly reduce the hole artifacts produced from the occlusion between objects, as well as a *self-occlusion error estimation* to restrict the hole artifacts produced from object's self-occlusion in a given size. Such estimation decides the representation of objects, those objects which will potentially result in holes smaller than a user-specified tolerance are represented by depth meshes; otherwise, by standard meshes of appropriate LOD resolutions. Hence, each object outside the cell will be checked to see if its *self-occluding error* is smaller than a predefined tolerance. The object will be associated with a depth mesh if such test is true, and a LOD mesh if it fails.

To reduce the redundancy of a regular-grid meshing, as well as the precision error caused by the projection from object space to image space, the *depth mesh* is simplified based on the original mesh and depth characteristic by an edge collapsing while preserving most of the important visual appearances.

Furthermore, to reduce the polygon count in scene navigation, for each navigation cell we do the conservative back-face culling for all LOD meshes, followed by a conservative occlusion culling for all depth meshes and LOD meshes, both aim to remove invisible polygons for any view inside the cell.

3.1 Pre-processing Stage

The processing steps in the pre-processing stage are:

1. Hexagonal spatial subdivision.
2. Selection of object's representation based on its self-occluding error estimation.
3. Depth mesh construction.
4. LOD mesh generation.
5. Regional conservative back-face culling.
6. Regional conservative occlusion culling.




Figure 1: The bird's eye view of the scene and the navigation path.

3.1.1 Hexagonal spatial subdivision

In order to utilize the spatial locality of a complex scene, we subdivide the x - y plane of the scene into $N \times M$ hexagonal navigation cells (Figure 1.) With the spatial subdivision, the scene data and viewpoints can be localized to cells, and, therefore, visibility culling, conservative back-facing and occlusion culling can be performed in pre-processing phase. A reason why the hexagonal subdivision, rather than rectangular subdivision, is used is that, in worse case, data of three adjacent cells need to be loaded, instead of four for rectangular subdivision, when navigating across the cell boundary. Table 1 depicts the maximum ratio of side faces could be seen from an inside point of the hexagonal and rectangular cell under different FOVs.

FOV($^\circ$)	120	90	60	45	30
Hexagonal	5/6	4/6	3/6	3/6	2/6
Rectangular	4/4	3/4	3/4	2/4	2/4

Table 1: Maximum ratio of side faces seen from a point inside the cell under different FOVs.

3.1.2 Self-occluding error estimation

The major problem of depth mesh representation is that it is only the visible part of the object viewed from the cell center hence has limited viewing degree of freedom. When a new view is far from the cell center, parts that are invisible originally might become visible and get rendered as holes. We propose a self-occluding error to estimate the maximum size of the hole that may appear when the object is represented by a depth mesh.

As shown in Figure 2, the maximum error occurs at the farthest view position V' from the cell center V . Let the cell size, i.e., the length of $\overline{VV'}$ be c , the distant between object and the cell center, i.e., the length of \overline{VO} be d , and the size in depth of the object itself, i.e., the length of \overline{OP} be l . The length of \overline{OC} is $l \cdot \tan \theta$, the angle θ between \overline{VP} and $\overline{V'P}$ is $\theta = \tan^{-1} \frac{c}{d+l}$, and s which is the projected size of \overline{OP} as well as \overline{OC} is:

$$s = \frac{\overline{AB}}{c} * \text{ImageRes.}$$




Figure 2: The maximum self-occluding error occurs at the position V' .

Since

$$\overline{AB} = \frac{\sqrt{3}c}{d} * \overline{OC} = \frac{\sqrt{3}c}{d} * \frac{cl}{d+l} = \frac{\sqrt{3}c^2l}{2d(d+l)},$$

therefore,

$$s = \frac{\sqrt{3}cl}{2d(d+l)} * \text{ImageRes.}$$

The test on self-occluding error is to check if the object's maximum projected error s is smaller than a user-specified tolerance in image precision. If it is, the object is represented by a depth mesh; otherwise by a LOD mesh.

3.1.3 Depth mesh construction

The cached image of an object is obtained by rendering the object using the cell center as the center of projection and using the cell's side face as the projection window. The *depth image* is the cached image augmented with the depth values. The simplest way to construct a *depth mesh* for an object is to use the regular-grid triangulation [17] performed on the depth image, which would, however, results in too many redundant triangles and produces rubber effects caused by incorrect connections. Furthermore, since it is performed in the image space, it always suffers from the precision error caused by the projection.

In order to reduce the number of the triangles on a depth mesh while preserving most of the visual appearances, several properties of the depth image could be adopted. The most important one is to use the depth coherence, by that we mean pixels of similar depth variation are likely to be on the same surface, and a pixel that has a sharp depth variation from adjacent pixels would have a high possibility to be on a contour edge. Moreover, the external contour edges of the rendered object on the image is the most important visual appearance, and hence must be included in the depth mesh. External contour edges can be easily derived by using the *contour extraction* in the field of image processing. If we can extract all the contour edges from the depth image, rubber effects caused by undetected gaps (C^0 discontinuity) between disjointed surfaces represented by a connected mesh and blur effects appearing at the sharp corners (C^1 discontinuity) represented by a flattened mesh could be greatly reduced.

In order to minimize the precision error caused by projection (from floating-point precision in world space into integer precision in image space), we simplify the depth mesh in both the image and the object space in three steps. Firstly, we categorize image pixels on the depth image based on the importance of its visual appearance and its characteristic into four categories:

- *external contour*: a pixel on the external contour extracted by *contour extraction*.
- *internal contour (gap) (C^0 discontinuity)*: a pixel whose Z value differs from neighboring pixels over a specified tolerance T_{C^0} (i.e., $|Z_i - Z_{i+1}| > T_{C^0}$, see Figure 3(a).)
- *sharp edge (C^1 discontinuity)*: a pixel whose Z variation differs from neighboring pixels over a specified tolerance T_{C^1} (i.e., $|(Z_{i-1} - Z_i) - (Z_i - Z_{i+1})| > T_{C^1}$, see Figure 3(b).)
- *interior*: other pixel whose Z value is different from the background Z value.

Secondly, vertices of object’s original mesh are projected again (but do not alter the value of depth image) with the same projection setup of the depth image to do the visibility test for each vertex. If Z value of a projected vertex equals to the Z value at the pixel on the depth image, the pixel on the depth image is representing the vertex of the original mesh; otherwise the vertex is behind other surfaces. For the former case, a weight relative to the category of the pixel is assigned to the vertex (The highest weight is assigned for *external contour* category, a lower weight is assigned for *internal contour* category, and so on.) On the other hand, for the later case (invisible vertex), the vertex gets the weight zero.

Lastly, weight-based edge collapsing is performed based on the weights of vertices to simplify the object’s original mesh. Moreover, to obtain a proper resolution of the simplified depth mesh and preserve the visual appearances, the edge with one of or both vertices’s weight smaller than the weight of the *sharp corner* category and whose projected size smaller than a user-specified length tolerance (in pixels) is simplified. As a result, a more optimized triangle aspect ratio is obtained and tiny triangles with respect to the view contain no important visual appearances are removed.

In addition, we use the JPEG algorithm to compress the cached images for reducing the storage requirements and reducing the loading time as well. Decompression needed at run-time can be supported by the state-of-the-art graphics hardware with texture decompression (eg. S3’s S3TC.) Bandwidth between host and graphics hardware can be greatly reduced as well.

3.1.4 Regional conservative back-face culling

As shown in Figure 4(b), for each polygon, we obtain the vector from one of six corner vertices of the navigation cell to the center of polygon, and do the dot product of the vector with polygon’s normal vector. If it is negative, the polygon is back-facing with respect to that corner vertex. If a polygon is back-facing for all six vertices of the cell, the polygon is back-facing with respect to any point inside the cell, and hence should be culled. In short, a polygon P

is back-facing for a navigation cell C , if

$$\begin{aligned} \text{dot_product}(P.\text{normal}, \text{vector}(C_i, P.\text{center})) < 0, \\ \text{for every } i = 0, \dots, 5, \end{aligned}$$

where C_i ’s are the corners of the navigation cell. A simple proof for the 2D case is as follows:

If a polygon P is back-facing with respect to both points A and B , P is back-facing with respect to any point on the line \overline{AB} (see Figure 4(a).) Any inner point I of a navigation cell is on a line $\overline{C_i E}$, where vertex E is on an edge $\overline{C_j C_{(j+1)\text{mod}6}}$ (see Figure 4(b).) Because, P is back-facing with respect to all corners C_k , $k = 0, \dots, 5$, P is back-facing with respect to E and therefore I .

Moreover, two corners are necessary and sufficient for this regional back-facing test for each polygon (see Figure 4(c).) For any point I' inside the navigation cell, $\overline{C_i I'}$ intersect $\overline{C_j P}$ at point D' , because D' is on $\overline{C_j P}$ and P is back-facing with respect to both C_i and C_j ; such that, P is back-facing with respect to D' , as well as I' . Note that, each polygon would required to test with different pair of corners.

3.1.5 Regional conservative occlusion culling


It is very common that objects or part of objects are occluded by other objects in a complex scene, especially a densely occluded environment. In order to utilize the spatial coherence of occlusion, we perform a regional conservative occlusion culling in the pre-processing stage for both LOD meshes and depth meshes.

The extended projection [7] provide us a useful tool with a little modification for our hexagonal subdivision scheme. This extended projection can also handle the case of multiple occluders by using occluder fusion. The selection of occluders is based on the meshes’ projected size. Only those meshes whose projected sizes are larger than a user-specified threshold are selected to be occluders.

3.2 Run-time stage

In the run-time stage, we do the following steps:


1. At program start-up time, we setup a lowest priority thread for pre-fetching the geometry and image data of neighboring cells.
2. Ensure that the geometry and image data for the current navigation cell is loaded into memory.
3. Perform a run-time normal cluster based back-face culling for the objects inside the current navigation cell.
4. Perform a run-time occlusion culling for all meshes.
5. Remain the remain meshes. Depth meshes are the rendered using the *projective-alike texture mapping*.
6. Pre-fetch data of neighboring cells when CPU load is relative low.



(a) Gap, C^0 discontinuity.

(b) Sharp corner, C^1 discontinuity.

Figure 3: gap and sharp corner extraction.



(a) P is back-facing with respect to any point on \overline{AB} , if P is back-facing with respect to both A and B .

(b) P is back-facing with respect to any inner point I in a navigation cell, when P is back-facing with respect to all corners.

(c) P is back-facing with respect to the whole gray area (contains the navigation cell), if P is back-facing with respect to both C_i and C_j .

Figure 4: Regional back-face culling.

3.2.1 Rendering

Objects inside the navigation cell can be seen from any direction, it is impossible to determine the visibility during the pre-processing stage. Those polygons inside the navigation cell can be group into clusters according to its normal [15] in the pre-processing stage. During the run-time stage, we can quickly cull out the whole back-facing cluster of polygons according to the viewing direction and the FOV.

Although there are considerable overheads, it is a beneficial approach to reduce the polygons sent into graphic pipeline by applying a run-time occlusion culling for a densely occluded environment. To perform the culling, we generate an occlusion map similar to the idea proposed in [30]. Only LOD meshes, depth meshes, and original meshes inside the navigation cell whose projected area larger than a pre-specified threshold are selected to be occluders. Though this approach of occluders selection is not optimized, it is advantageous not to spend too much time on selecting occluders.

3.2.2 Projective-alike texture mapping

A *projective-alike-texture-mapping* method is developed to map the cached image onto depth mesh in such a way that the texture coordinate of each vertex can be generated automatically by the standard OpenGL (`glTexGen()`).

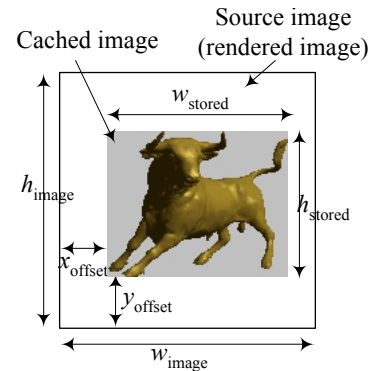


Figure 5: A cached image is a part of a source image.

In most cases, a object-based depth mesh located in part of a source image (e.g. see Figure 5.) To minimize the storage requirement of cached images, only necessary rectangular part of source image is stored as cached image.

Texture coordinates (s, t) maps from image coordinates (x, y) can be derived by $(s, t) = ((x - x_{\text{offset}})/w_{\text{stored}}, (y - y_{\text{offset}})/h_{\text{stored}})$, where x_{offset} and y_{offset} are the offset of the cached image relative to source image, and w_{stored} and h_{stored} are width and height of the cached image, respectively. The GPU of graphics hard-




Figure 6: Re-projection from source to destination image, T_1 and T_2 are the camera matrix of source image I_1 and destination image I_2 , respectively..

ware can also do the transformation of re-projection from the source image coordinates to the destination image coordinates by multiply the inverse of the camera matrix of the source image (Figure 6), allowing CPU leisurely does the visibility culling, level-of-detail selection, pre-fetching, and so on.

This method does not need to specify texture coordinates, and in consequence, reduces the bandwidth needed between CPU and graphics accelerator, and does not require additional memory for storing texture coordinate at each vertex. Moreover, vertices of a depth meshes are allowed to be stored in source image coordinate system, which requires 16-bit unsigned integer for x and y , and 32-bit floating point for z . As a result, 8 bytes is necessarily sufficient for each vertex, compared to 12 bytes per vertex if all x , y , and z are stored as world coordinates as well as 8 bytes for texture coordinates s and t .

3.2.3 Cell Transition

A major problem arises in spatial-subdivision approach is how to achieve smooth and unnoticeable transition between cells. When the view point moves across the cell boundary and makes a transition from cell A to cell B , we will switch the geometry set from G_A to G_B , depth mesh set from D_A to D_B , and cached image set from I_A to I_B . Accordingly, it may spend a lot of time in loading data from disk.

Here, we develop a pre-fetch mechanism which pre-loads the geometry and image data of neighboring cells when CPU load is relative low. It will amortize the loading time to several inside-cell frames and reduce the difference of rendering time between an inside-cell frame and a cross-cell-boundary frame. We can easily attain this by setting the priority of a pre-fetch thread to be lower than others. As a result, we will not be interrupted by the loading of newly navigated data during the cell transition and obtain a more smooth frame rates and an unnoticeable transition. Note that the smaller cell size is the more frequent the transition is.

4 EXPERIMENTS

Our test scene is a statuary park which is consists of 784 objects with 880, 840 polygons on an area of 1200×1000 ; see Figure 1. Figure 7 represents images of two different

views of our scene. Notice that the standing horse in Figure 7(b) is represented by a depth mesh, in a contrary, the front cattle in Figure 7(d) is represented by a LOD mesh.

The test platform is a PC with an AMD ThunderBird 1200Mhz CPU, 512MB main memory, and an nVidia GeForce3 with 64MB DDR RAM graphics accelerator.

4.1 Image Quality

To identify how much quality-loss of our proposed method, we use the signal-to-noise ratio SNR(dB), defined as follows:

$$\text{SNR} = 10 * \log \frac{\sum_{x=0}^{M-1} \sum_{y=0}^{N-1} \hat{f}(x, y)^2}{\sum_{x=0}^{M-1} \sum_{y=0}^{N-1} [\hat{f}(x, y) - f(x, y)]^2}$$

where $\hat{f}(x, y)$ is the pixel color of the approximated image at position (x, y) , M and N are the dimensions of the image.

Before applying SNR, the RGB color is mapped to a single luminance value Y since our human eyes are more sensitive to the changes in luminance than to the changes in chrominance. Such mapping [8] is

$$Y = 0.299 * R + 0.587 * G + 0.114 * B.$$

Note that, human eyes in general would not able to distinguish between two images that have an signal-to-noise ratio (SNR) greater than 25dB.

4.2 Cell Size and Self-Occluding Error Consideration

Four settings of the cell size and self-occluding error tolerance are used in our experimental test: two different cell sizes (50 and 100) with two different self-occluding error tolerances (1-pixel and 3-pixel.)

The distributions of LOD meshes and depth meshes for these four different settings are shown in Figure 8. A LOD mesh is marked as a red cross while a depth mesh as a blue point. As we can expect, the higher self-occluding error we can tolerate, the more objects are represented by depth meshes. In the case of cell size 50 and 1-pixel error tolerance, there are four LOD meshes, while there is only one LOD meshes for the setting of 3-pixel error tolerance. However, with no doubt, image quality for 3-pixel tolerance is lower than that of 1-pixel tolerance.

Figure 9 compares the image quality obtained by the proposed scheme and by rendering with pure geometry for the four different settings. Similarly, Figure 10 shows the frame time comparison but with no run-time occlusion culling for our scheme. Table 2 lists the polygon simplification ratio for the objects represented by LOD and depth meshes, the average frame rate, and average image quality for the four settings. All of them show that, more polygons are simplified for the higher error tolerance setting with the cost of higher quality-loss. That is, a trade-off between the performance and image quality. Note that, the quality differences between different cell sizes are almost the same, mainly due to the fact that the selection of object's representation is designed to ensure the bounded self-occluding error. Hence, the cell size has few impact on the image quality under our proposed scheme. However, as the cell size increases, more objects are represented by LOD meshes, more objects are put inside the navigation

cell, and the number of potential visible polygons increases also, so the number of polygons increases dramatically. As a result, worse performance will be found.

Table 3 depicts the culling rate of the conservative back-facing culling, in which about 30% polygons of LOD meshes are culled out for the case of cell size 50. In general, the larger size a cell has, the fewer back-facing rate is.

At the time of writing this paper, we still work on several minor bugs of this pre-processing occlusion culling program, therefore, we have no information about how many percentage of polygons will be culled, at this moment.

	Polygons number in LOD meshes	Back-face-culled polygon number	Culled rate
size 50, 3-pixel	9,511	2,891	30.40%
size 100, 3-pixel	8,887	2,018	22.71%

Table 3: Average culling rate of the regional back-face culling for all cells.

4.3 Performance

We use the setting of cell size 50 and 3-pixel self-occluding error tolerance for the further testing. Four rendering configurations are used to do the performance comparison:

- **A: (Pure geometry)** The original meshes of the scene are rendered using the traditional graphics pipeline.
- **B: (Pure geometry w/ view frustum culling)** Same as A, but with software view frustum culling.
- **C: (Proposed scheme w/ run-time occlusion culling)** The scene is rendered by the proposed scheme with view frustum culling and run-time occlusion culling.
- **D: (Proposed scheme w/o run-time occlusion culling)** The scene is rendered by proposed scheme with view frustum culling but *without* run-time occlusion culling.

Since we want to identify how much benefit comes from the depth mesh representation alone, the configuration C and D use original meshes instead of LOD meshes to represent those objects which the self-occluding error exceeding the tolerance.

Table 4 lists the performance of each configuration. Configuration C and D spend additional time on loading neighboring cell data at the first frame, so, there are peaks at the first frame. Our proposed scheme w/o run-time occlusion culling has about 18.7 speedup factor compare to the pure geometric rendering while yields an average SNR about 22dB. As mention before, the overhead of run-time occlusion culling is unnegligible, and our test scene is not a densely occluded environment, so, configuration C is slower than D but still outperforms B.

Figure 11 depicts the frame time of every rendering frames on the same navigation path under these four rendering configurations. It shows that our proposed method provides a faster rendering with an acceptable quality-loss and has a more constant frame rate compared to the rendering configuration B.

	A	B	C	D
Frame Time(ms)	332.50	81.73	60.29	17.75
Frame Rate(fps)	3.01	12.24	16.59	56.35
Speed Up	1×(baseline)	4.07×	5.51×	18.7×

Table 4: Performance of the four testing configurations.

5 DISCUSSION & FUTURE WORK

In this paper, we have proposed a hexagonal spatial subdivision and a hybrid scheme for navigating complex scenes. The method can achieve a smooth, navigation with no apparent popping effects at an almost constant and interactive frame rate. By cooperating with LOD meshes and depth meshes, parallax error and popping effect will be minimized. Furthermore, with visibility preprocessing, polygons invisible from a region will never sent into graphics pipeline. A tradeoff between performance and quality requirement can be easily made by specifying the self-occluding error tolerance.

As the future works, we will firstly finish the pre-processing occlusion culling program to effectively cull out the polygons that are invisible from a cell region. We will also improve the depth mesh construction and simplification to have a better simplified depth mesh. Moreover, we will try to exploit the data coherence between neighboring cells to improve our pre-loading scheme.


References

- [1] S. E. Chen. Quicktime VR - an image-based approach to virtual environment navigation. In R. Cook, editor, *Proceedings of SIGGRAPH 95*, Annual Conference Series, pages 29–38. ACM SIGGRAPH, Addison Wesley, August 1995.
- [2] S. E. Chen and L. Williams. View Interpolation for Image Synthesis. In J. T. Kajiya, editor, *Computer Graphics (SIGGRAPH 93 Proceedings)*, pages 279–288, August 1993.
- [3] D. Cohen-Or, G. Fibich, D. Halperin, and E. Zadicario. Conservative Visibility and Strong Occlusion for Viewspace Partitioning of Densely Occluded Scenes. volume 17, pages 243–253, 1998.
- [4] L. Darsa, B. Costa, and A. Varshney. Walkthroughs of Complex Environments using Image-based Simplification. In *Computers & Graphics*, volume 22, pages 55–69, 1998.
- [5] L. Darsa, B. C. Silva, and A. Varshney. Navigating Static Environments Using Image-Space Simplification and Morphing. In *Symposium on Interactive 3D Graphics*, pages 25–34, 1997.
- [6] X. Decoret, G. Schauffer, F. X. Sillion, and J. Dorsey. Multi-Layered Impostors for Accelerated Rendering. In *Computer Graphics Forum 18(3), proceedings of EUROGRAPHICS'99*, Milano, Italy, September 1999.
- [7] F. Durand, G. Drettakis, J. Thollot, and C. Puech. Conservative Visibility Preprocessing using Extended Projections. In Kurt Akeley, editor, *Computer Graphics (SIGGRAPH 2000 Proceedings)*, Annual Conference Series, pages 239–248. ACM SIGGRAPH, Addison Wesley, July 2000.


	Polygon number of LOD mesh	Polygon number of Depth mesh	LOD mesh and Depth mesh	Simplified ratio of meshes outside the cell	Average frame rate(fps)	Image quality(SNR(dB))
size 50, 1-pixel	19,628	71,875	91,503	89.56%	48.86	22.66
size 50, 3-pixel	6,620	80,841	87,461	90.02%	56.35	21.86
size 100, 1-pixel	31,819	116,862	148,681	82.77%	38.15	21.65
size 100, 3-pixel	6,869	128,205	135,074	84.35%	46.75	21.60

Table 2: The average polygon simplification ratio for each cell, the average frame rate, and image quality for the four settings.

- [8] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*, chapter 4, page 228. Addison-Wesley, September 1993.
- [9] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The Lumigraph. In H. Rushmeier, editor, *Proceedings of SIGGRAPH 96*, Annual Conference Series, pages 43–54. ACM SIGGRAPH, Addison Wesley, August 1996.
- [10] N. Greene, M. Kass, and G. Miller. Hierarchical Z-Buffer Visibility. In J. T. Kajiya, editor, *Computer Graphics (SIGGRAPH 93 Proceedings)*, pages 231–238, August 1993.
- [11] H. Hoppe. Progressive Meshes. In H. Rushmeier, editor, *Proceedings of SIGGRAPH 96*, Annual Conference Series, pages 99–108. ACM SIGGRAPH, Addison Wesley, August 1996.
- [12] H. Hoppe. View-Dependent Refinement of Progressive Meshes. In *Proceedings of SIGGRAPH 97*, Annual Conference Series, pages 189–198. ACM SIGGRAPH, Addison Wesley, 1997.
- [13] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh Optimization. In J. T. Kajiya, editor, *Computer Graphics (SIGGRAPH 93 Proceedings)*, pages 19–26, August 1993.
- [14] T. Hudson, D. Manocha, J. Cohen, M. Lin, K. Hoff, and H. Zhang. Accelerated Occlusion Culling using Shadow Frusta. In *Proceedings of 13th Symposium on Computational Geometry*, pages 1–10, 1997.
- [15] S. Kumar, D. Manocha, B. Garrett, and M. Lin. Hierarchical Back-face Culling. In *7th Eurographics Workshop on Rendering*, pages 231–240, 1996.
- [16] M. Levoy and P. Hanrahan. Light Field Rendering. In H. Rushmeier, editor, *Proceedings of SIGGRAPH 96*, Annual Conference Series, pages 31–42. ACM SIGGRAPH, Addison Wesley, August 1996.
- [17] W. R. Mark, L. McMillan, and G. Bishop. Post-Rendering 3D Warping. In *Symposium on Interactive 3D Graphics*, pages 7–16, 1997.
- [18] L. McMillan and G. Bishop. Plenoptic Modeling: An Image-Based Rendering System. In R. Cook, editor, *Proceedings of SIGGRAPH 95*, Annual Conference Series, pages 39–46. ACM SIGGRAPH, Addison Wesley, August 1995.
- [19] J. Rossignac and P. Borrel. Multi-resolution 3-D Approximations for Rendering Complex Scenes. In B. Falcidieno and T. L. Kunii, editors, *Modeling in Computer Graphics*, pages 455–465, Berlin, 1993. Springer-Verlag.
- [20] G. Schaufler, J. Dorsey, X. Decoret, and F. X. Sillion. Conservative Volumetric Visibility with Occluder Fusion. In Kurt Akeley, editor, *Computer Graphics (SIGGRAPH 2000 Proceedings)*, pages 229–238. ACM SIGGRAPH, Addison Wesley, July 2000.
- [21] G. Schaufler and W. Stürzlinger. A Three-Dimensional Image Cache for Virtual Reality. In *Proceedings of Eurographics '96*, pages 227–236, August 1996.
- [22] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of Triangle Meshes. *Computer Graphics*, 26(2):65–69, 1992.
- [23] S. M. Seitz and C. R. Dyer. View Morphing. In H. Rushmeier, editor, *Proceedings of SIGGRAPH 96*, Annual Conference Series, pages 21–30. ACM SIGGRAPH, Addison Wesley, August 1996.
- [24] J. Shade, D. Lischinski, D. H. Salesin, T. DeRose, and J. Snyder. Hierarchical Image Caching for Accelerated Walkthroughs of Complex Environments. In H. Rushmeier, editor, *Proceedings of SIGGRAPH 96*, Annual Conference Series, pages 75–82. ACM SIGGRAPH, Addison Wesley, August 1996.
- [25] J. W. Shade, S. J. Gortler, L.-W. He, and R. Szeliski. Layered Depth Images. In *Proceedings of SIGGRAPH 98*, Annual Conference Series, pages 231–242. ACM SIGGRAPH, Addison Wesley, July 1998.
- [26] F. Sillion, G. Drettakis, and B. Bodelet. Efficient Impostor Manipulation for Real-Time Visualization of Urban Scenery. In *Proceedings of Eurographics '97*, pages 207–218, Budapest, Hungary, September 1997.
- [27] J. Xia and A. Varshney. Dynamic View-Dependent Simplification for Polygonal Models. In *IEEE Visualization*, pages 327–334, 1996.
- [28] S. K. Yang. Material-Preserving Progressive Mesh Using Geometry and Topology Simplification. Master’s thesis, Department of Computer Science and Information Engineering, National Chiao Tung University, 1999.
- [29] H. Zhang and K. E. Hoff III. Fast Backface Culling Using Normal Masks. In *Symposium on Interactive 3D Graphics*, pages 103–106, 1997.
- [30] H. Zhang, D. Manocha, T. Hudson, and K. Hoff. Visibility Culling Using Hierarchical Occlusion Maps. In *Computer Graphics*, volume 31 of *Annual Conference Series*, pages 77–88, 1997.




(a) View 1, rendered by pure geometries.



(b) View 1, rendered by the proposed scheme.

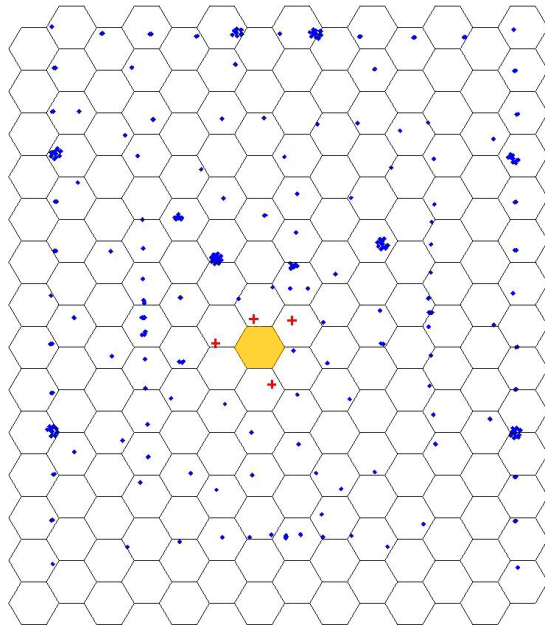


(c) View 2, rendered by pure geometries.

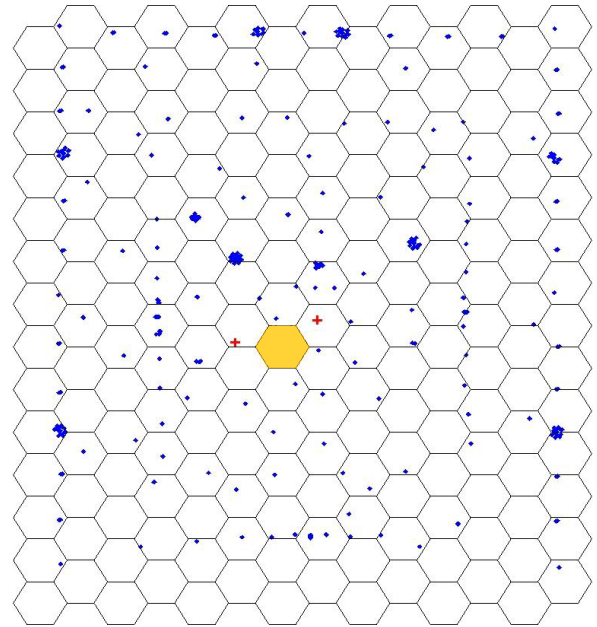


(d) View 2, rendered by the proposed scheme.


Figure 7: Two different views of our test scene.




(a) size 50, 1-pixel



(b) size 50, 3-pixel



(c) size 100, 1-pixel



(d) size 100, 3-pixel

Figure 8: Distribution of LOD meshes and depth meshes under four different settings.




Figure 9: The quality of the four different settings on the same navigation path.

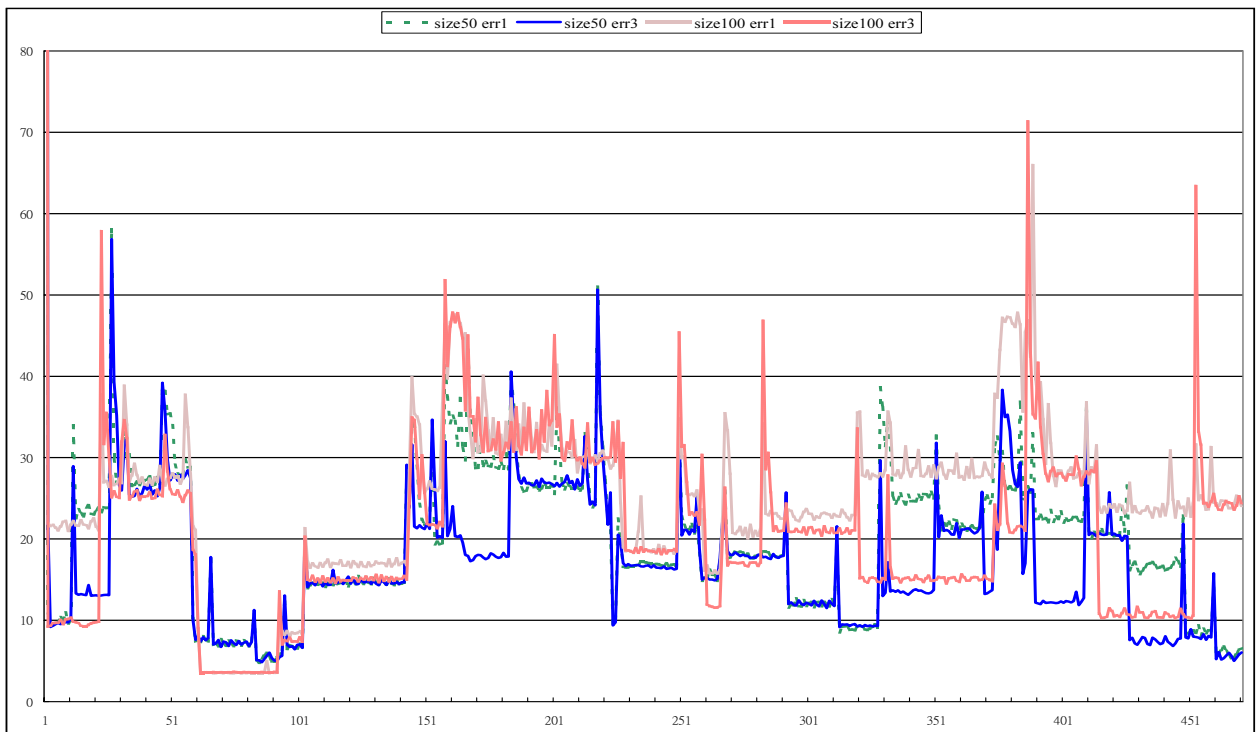


Figure 10: The frame time of the four different settings on the same navigation path.




Figure 11: The frame time of the four different configurations on the same navigation path.