

Rendering on a Budget: A Framework for Time-Critical Rendering

James T. Klosowski* Cláudio T. Silva†

IBM T. J. Watson Research Center

Abstract

We present a technique for optimizing the rendering of high-depth complexity scenes. *Prioritized-Layered Projection* (PLP) does this by rendering an estimation of the visible set for each frame. The novelty in our work lies in the fact that we do not explicitly compute visible sets. Instead, our work is based on computing *on demand* a priority order for the polygons that maximizes the likelihood of rendering visible polygons before occluded ones for any given scene. Given a fixed budget, *e.g.* time or number of triangles, our rendering algorithm makes sure to render geometry respecting the computed priority.

There are two main steps to our technique: (1) an occupancy-based tessellation of space; and (2) a solidity-based traversal algorithm. PLP works by computing an occupancy-based tessellation of space, which tends to have smaller cells where there are more geometric primitives, *e.g.*, polygons. In this spatial tessellation, each cell is assigned a *solidity* value, which is directly proportional to its likelihood of occluding other cells. In its simplest form, a cell's solidity value is directly proportional to the number of polygons contained within it. During our traversal algorithm, cells are marked for projection, and the geometric primitives contained within them actually rendered. The traversal algorithm makes use of the cells' solidity, and other view-dependent information to determine the ordering in which to project cells. By tailoring the traversal algorithm to the occupancy-based tessellation, we can achieve very good frame rates with low preprocessing and rendering costs.

In this paper, we describe our technique and its implementation in detail. We also provide experimental evidence of its performance and briefly discuss extensions of our algorithm.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Visible line/surface algorithms

Keywords: Polygon rendering, visibility ordering, occlusion culling

1 Introduction

Recent advances in graphics hardware have not been able to keep up with the increase in scene complexity. In order to support a new set of demanding applications, a multitude of rendering algorithms have been developed to both augment

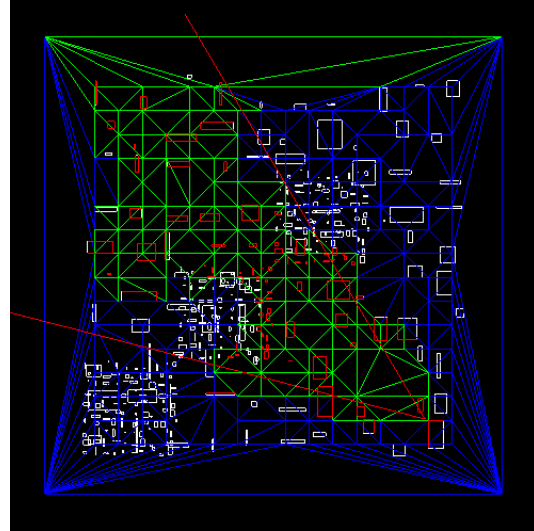


Figure 1: The Prioritized-Layered Projection Algorithm. PLP attempts to prioritize the rendering of geometry along layers of occlusion. The input geometry, line segments in two-dimension, is drawn in white and the spatial tessellation, a Delaunay Triangulation, is drawn in blue. Cells that have been projected by the PLP algorithm are highlighted in green and rendered geometry is drawn in red. The view frustum is also highlighted as red line segments. In this particular example, a budget of 500 line segments was used. The 2D prototype implementation does not enforce star-shaped constraints (or penalties) on the front.

and optimize the use of the hardware. An effective way to speed up rendering is to avoid rendering geometry that cannot be seen from the given viewpoint, such as geometry that is outside the view frustum, faces away from the viewer, or is obscured by geometry closer to the viewer. Quite possibly, the hardest part of the visibility-culling problem is to avoid rendering geometry that can not be seen due to its being obscured by closer geometry. In this paper, we propose a new algorithm for solving the visibility culling problem. Our technique is an effective way to cull geometry with a very simple and general algorithm.

Our technique optimizes for rendering by estimating the visible set for a given frame, and only rendering those polygons. It is based on computing on demand a priority order for the polygons that maximizes the likelihood of projecting visible polygons before occluded ones for any given scene. It does so in two steps: (1) as a preprocessing step, it computes an occupancy-based tessellation of space, which tends to have

*Visual Technologies, IBM T. J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598; jklosow@watson.ibm.com.

†Work performed while at IBM. The author's current address: AT&T Labs-Research, 180 Park Ave., PO Box 971, Florham Park, NJ 07932; csilva@research.att.com.

smaller spatial cells where there are more geometric primitives, *e.g.*, polygons; (2) in real-time, rendering is performed by traversing the cells in an order determined by their intrinsic solidity and some other view-dependent information. As cells are projected, their geometry is scheduled for rendering (see Fig. 1). Actual rendering is constrained by a user-defined budget, *e.g.* time or number of triangles.

Some highlights of our technique:

- **Budget-based rendering.** Our algorithm generates a projection ordering for the geometric primitives that mimics a depth-layered projection ordering, where primitives directly visible from the viewpoint are projected earlier in the rendering process. The ordering and rendering algorithms strictly adhere to a user-defined budget, making the PLP approach time-critical.
- **Low-complexity preprocessing.** Our algorithm requires inexpensive preprocessing, that basically amounts to computing an Octree and a Delaunay triangulation on a subset of the vertices of the original geometry.
- **No need to choose occluders beforehand.** Contrary to other techniques, we do not require that occluders be found before geometry is rendered.
- **Object-space occluder fusion.** All of the occluders are found automatically during a space traversal that is part of the normal rendering loop without resorting to image-space representation.
- **Simple and fast to implement.** Our technique amounts to a small modification of a well-known rendering loop used in volume rendering of unstructured grids. It only requires negligible overhead on top of view-frustum culling techniques.

Our paper is organized as follows. In Section 2, we give some preliminary definitions, and briefly discuss relevant related work. In Section 3, we propose our novel visibility-culling algorithm. In Section 4, we give some details on our prototype implementation. In Section 5, we provide experimental evidence of the effectiveness of our algorithm. In Section 6, we conclude the paper with some final remarks and future work.

2 Preliminaries and Related Work

The visibility problem is defined in [9] as follows. Let the scene, \mathcal{S} , be composed of modeling primitives (*e.g.*, triangles or spheres) $\mathcal{S} = \{P_0, P_1, \dots, P_n\}$, and a viewing frustum defining an eye position, a view direction, and a field of view. The visibility problem encompasses finding the points or fragments within the scene that are visible, that is, connected to the eye point by a line segment that meets the closure of no other primitive. For a scene with $n = O(|\mathcal{S}|)$ primitives, the complexity of the set of visible fragments might be as high as $O(n^2)$, but by exploiting the discrete nature of the screen, the Z-buffer algorithm [2] solves the visibility problem in time $O(n)$, since it only touches each primitive once. The Z-buffer algorithm solves the visibility problem by keeping a depth value for each pixel, and only updating the pixels when geometry closer to the eye point is rendered. In the case of high-depth complexity scenes, the Z-buffer might overdraw

each pixel a considerable number of times. Despite this potential inefficiency, the Z-buffer is a popular algorithm, widely implemented in hardware.

In light of the Z-buffer being widely available, and exact visibility computations being potentially too costly, one idea is to use the Z-buffer as filter, and design algorithms that lower the amount of overdraw by computing an approximation of the *visible set*. In more precise terms, define the visible set $V \subset \mathcal{S}$ to be the set of modeling primitives which contribute to at least one pixel of the screen.

In computer graphics, visibility-culling research mainly focussed on algorithms for computing conservative (hopefully tight) estimations of V , then using the Z-buffer to obtain correct images. The simplest example of visibility-culling algorithms are backface and view-frustum culling [11]. Backface-culling algorithms avoid rendering geometry that face away from the viewer, while viewing-frustum culling algorithms avoid rendering geometry that is outside of the viewing frustum. Even though both of these techniques are very effective at culling geometry, more complex techniques can lead to substantial improvements in rendering time. These techniques for tighter estimation of V do not come easily. In fact, most techniques proposed are quite involved and ingenious, and usually require the computation of complex object hierarchies in both 3- and 2-space.

Here again the discrete nature of the screen, and screen-space coverage tests, play a central role in literally all occlusion-culling algorithms, since it paves the way for the use of screen occupancy to cull 3D geometry that projects into already occupied areas. In general, algorithms exploit this fact by (1) projecting P_i in front-to-back order, and (2) keeping screen coverage information. Several efficiency issues are important for occlusion-culling algorithms:

- (a) They must operate under great time and space constraints, since large amounts of geometry must be rendered in fractions of a second for real-time display.
- (b) It is imperative that primitives that will not be rendered be discarded as early as possible, and (hopefully) not be touched at all. Global operations, such as computing a full front-to-back ordering of P_i , should be avoided.
- (c) The more geometry that gets projected, the less likely the Z-buffer gets changed. In order to effectively use this fact, it must be possible to merge the effect of multiple occluders. That is, it must be possible to account for the case that neither P_0 nor P_1 obscures P_2 by itself, but together they do cover P_2 . Algorithms that do not exploit *occluder-fusion* are likely to rely on the presence of large occluders in the scene.

A great amount of work has been done in visibility culling in both computer graphics and computational geometry. For those interested in the computational geometry literature, see [8, 9, 10]. For a survey of computer graphics work, see [20].

We very briefly survey some of the recent work more directly related to our technique. Hierarchical occlusion maps [21] solve the visibility problem by using two hierarchies, an object-space bounding volume hierarchy and another hierarchy of image-space occlusion maps. For each frame, objects from a pre-computed database are chosen to be occluders and used to cull geometry that cannot be seen. A closely related technique is the hierarchical Z-buffer [13]. In [1], an

extension of graphics hardware for occlusion-culling queries is proposed.

It is possible to perform object-space visibility culling. One such technique, described in [18], divides space into cells, which are then preprocessed for potential visibility. This technique works particularly well for architectural models. Additional object-space techniques are described in [6, 7]. These techniques mostly exploit the presence of large occluders, and keep track of spatial extents over time. In [4], a technique that precomputes visibility in densely occluded scenes is proposed. They show it is possible to achieve very high-occlusion rates in dense environments by pre-computing simple ray-shooting checks.

In [12], a constant-frame rendering system is described. This work uses the visibility-culling from [18]. It is related to our approach in the sense that it also uses a (polygon) budget for limiting the overall rendering time. Other notable references include [3], for its level-of-detail management ideas, and [16], where a scalable rendering architecture is proposed.

3 The PLP Algorithm

In this paper we propose the *Prioritized-Layered Projection* algorithm, a simple and effective technique for optimizing the rendering of geometric primitives. The guts of our algorithm consists of a space-traversal algorithm, which prioritizes the projection of the geometric primitives in such a way as to avoid (actually delay) projecting cells that have a small likelihood of being visible. Instead of conservatively overestimating V , our algorithm works on a budget. At each frame, the user can provide a maximum number of primitives to be rendered, *i.e.*, a polygon budget, and our algorithm, in its single-pass traversal over the data, will deliver what it considers to be the set of primitives which maximizes the image quality, using a solidity-based metric.

Our projection strategy is completely object-space based, and resembles* cell-projection algorithms used in volume rendering unstructured grids.

In a nutshell, our algorithm is composed of two parts:

Preprocessing. Here, we tessellate the space that contains the original input geometry with convex cells in the way specified in Section 3.1. During this one-time preprocessing, a collection of tetrahedra is generated in such a way as to roughly keep a uniform density of primitives per tetrahedron. Our sampling leads to large tetrahedra in unpopulated areas, and small tetrahedra in areas that contain a lot of geometry.

In another similarity to volume rendering, using the number of modeling primitives assigned to a given cell (*e.g.*, tetrahedron) we define its *solidity* value ρ , which is similar to the opacity used in volume rendering. In fact, we use a different name to avoid confusion since the accumulated solidity value used throughout our priority-driven traversal algorithm can be larger than one. Our traversal algorithm prioritizes cells based on their solidity value.

Generating such a space tessellation is not a very expensive step, *e.g.* taking only two minutes for a scene composed of

*Our cell-projection algorithm is different than the ones used in volume rendering in the following ways: (1) in volume rendering cells are usually projected in back-to-front order, while in our case, we project cells in *roughly* front-to-back order; (2) more importantly, we do not keep a strict depth-ordering of the cells during projection. This would be too restrictive, and expensive, for our purposes.




Figure 2: A snapshot of the PLP algorithm. Rendered geometry is shown in red. The cells of the spatial tessellation that are in the front are drawn in blue. The next cell to be projected is the (blue) one with the lowest solidity value.

one million triangles, and for several large datasets can even be performed as part of the data input process. Of course, for truly large datasets, we highly recommend generating this view-independent data structure beforehand, and saving it with the original data.

Rendering Loop. Our rendering algorithm traverses the cells in roughly front-to-back order. Starting from the seed cell, which in general contains the eye position, it keeps carving cells out of the tessellation. The basic idea of our algorithm is to carve the tessellation along *layers of polygons*. We define the layering number $\zeta \in \mathfrak{X}$ of a modeling primitive P in the following intuitive way. If we order each modeling primitive along each pixel by their positive[†] distance to the eye point, we define $\zeta(P)$ to be the smallest rank of P over all of the pixels to which it contributes. Clearly, $\zeta(P) = 1$, if, and only if, P is visible.

Finding the rank 1 primitives is equivalent to solving the visibility problem. Instead of solving this hard problem, the PLP algorithm uses simple heuristics. Our traversal algorithm attempts to project the modeling primitives by layers, that is, all primitives of rank 1, then 2 and so on. We do this by always projecting the cell in the front F (we call *the front*, the collection of cells that are immediate candidates for projection) which is least likely to be occluded according to its solidity values. Initially, the front is empty, and as cells are inserted, we estimate its accumulated solidity value to reflect its position during the traversal. (Cell solidity is defined below in Section 3.2). Every time a cell in the front is projected, all of the geometry assigned to it is rendered.

In Fig. 2, we can see a snapshot of our algorithm as it carves its way into space. The rendered geometry is shown in red and the cells in the front are shown in blue. Note how the algorithm has carved space around the body of the automobile.

[†]Without loss of generality, assume P is in the view frustum.




Figure 3: Occupancy-based spatial tessellation algorithm. The input geometry, a car with an engine composed of over 160K triangles, is shown in (a). Using the vertices of the input geometry, we build an error-bounded octree, shown in (b). The centers of the leaf-nodes of the octree, shown in yellow in (c), are used as the vertices of our Delaunay Triangulation.

There are several types of budgeting that can be applied to our technique, for example, a triangle-count budget can be used to make it time-critical. For a given budget of k modeling primitives, let T_k be the set of primitives our traversal algorithm projects. This set, together with S , the set of all primitives, and V , the set of visible primitives, can be used to define several statistics that measure the overall effectiveness of our technique. One relevant statistic is the *visible coverage ratio* for a budget of k primitives, ϵ_k . This is the number of primitives in the visible set that we actually render, that is, $\epsilon_k = \frac{|V \cap T_k|}{|V|}$. If $\epsilon_k < 1$, we missed rendering some visible primitives. PLP does not attempt to compute the visible set exactly. Instead, it combines a budget with its solidity-based polygon ordering. For a polygon budget of k , the best case scenario would be to have $\epsilon_k = 1$. Of course, this would mean that PLP finds all of the visible polygons.

In addition to the visible coverage ratio ϵ_k , another important statistic is the number of incorrect pixels in the image produced by the PLP technique. This provides a measure as to how closely the PLP image represents the exact image produced by rendering all of the primitives.

3.1 Occupancy-Based Spatial Tessellations

The underlying data structure used in our technique is a decomposition of the 3-space covered by the scene into disjoint cells. The characteristics we required in our spatial decomposition were:

- (a) **Simple traversal characteristics** - must be easy and computationally inexpensive to walk from cell to cell.
- (b) **Good projection properties** - depth-orderable from any viewpoint (with efficient, hopefully linear-time projection algorithms available); easy to estimate screen-space coverage.
- (c) **Efficient space filler** - given an arbitrary set of geometry, it should be possible to sample the geometry adaptively,

that is, with large cells in sparse areas, and smaller cells in dense areas.

(d) **Easy to build and efficient to store.**

We could have used any of a number of different spatial data structures, such as kd-trees or octrees, but we settled on a Delaunay triangulation, since it seems to most closely fill our needs. In making this decision, we were influenced by the work of Held *et al.* [14] on computing low-stabbing triangulations for collision detection; and Williams' MPVO [19], a linear-time algorithm for visibility ordering meshes mainly used for volume rendering.

In order to compute a spatial decomposition M , which adaptively samples the scene S , we use a very simple procedure that in effect just samples S with points, then constructs M as the Delaunay triangulation of the sample points, and finally assigns individual primitives in S to M . Fig. 3 shows our overall triangulation algorithm. Instead of accurately sampling the actual primitives (Fig. 3a), such as is done in [15], we simply construct an octree using only the original vertices (Fig. 3b); we limit the size of the octree leaves, which gives us a bound on the maximum complexity of our mesh; then we use the (randomly perturbed) center of the octree leaves as the vertices of our Delaunay triangulation (Fig. 3c). After M is built, we use a naive assignment of the primitives in S to M , by basically scan-converting the geometry into the mesh. Each cell $c_i \in M$, has a list of the primitives from S assigned to it. Each of these primitives is either completely contained in c_i , or it intersects one of its boundary faces. We use $|c_i|$, the number of primitives in cell c_i , in the algorithm that determines the solidity values of c_i 's neighboring cells. In a final pass over the data during preprocessing, we compute the maximum number of primitives in any cell, $\rho_{max} = \max_{i \in [1..|M|]} |c_i|$, to be used later as a scaling factor.

Remarks: (1) The resolution of the octree we use is very modest. By default, once an octree node has a side shorter than 5% of the length of the bounding box of S , it is considered

```

Algorithm RenderingLoop()
1. while (empty( $F$ ) != true)
2.    $c = \min(F)$ 
3.   project( $c$ )
4.   if (reached_budget() == true)
5.     break;
6.   for each  $n; n = \text{cell\_adjacent\_to}(c)$ 
7.     if (projected( $n$ ) == true)
8.       continue;
9.      $\rho = \text{update\_solidity}(n, c)$ 
10.    enqueue( $n, \rho$ )

```

Figure 5: Skeleton of the *RenderingLoop* algorithm. Function *min*(F) returns the minimum element in the priority queue F . Function *project*(c) renders all the elements assigned to cell c ; it also counts the number of primitives actually rendered. Function *reached_budget*() returns true if we have already rendered k primitives. Function *cell_adjacent_to*(c) lists the cells adjacent to c . Function *projected*(n) returns true if cell n has already been projected. Function *update_solidity*(n, c) computes the updated solidity of cell n , based on the fact that c is one of its neighbors, and has just been projected. Function *enqueue*(n, ρ) places n in the queue with a solidity ρ . If n was already in the queue, this function will first remove it, and re-insert it with the updated solidity value. See text for more details on *update_solidity*().

a leaf node. This has shown to be quite satisfactory for all the experiments we have performed thus far. (2) Even though primitives might be assigned to multiple cells of M (we use pointers to the actual primitives), the memory overhead has been negligible. See Section 5.1.

3.2 Priority-Based Traversal Algorithm

Cell-projection algorithms [19, 17, 5] are implemented using queues or stacks, depending on the type of traversal (*e.g.*, depth-first versus breadth-first), and use some form of restrictive dependency among cells to ensure properties of the order of projection (*e.g.*, strict back-to-front).

Unfortunately such limited and strict projection strategies do not seem general enough to capture the notion of polygon layering, which we are using for visibility culling. In order for this to be feasible, we must be able to selectively stop (or at least delay) cell-projection around some areas, while continuing in others. In effect, we would like to project cells from M using a layering defined by the primitives in S . The intuitive notion we are trying to capture is as follows: if a cell c_i has been projected, and $|c_i| = \rho_{max}$, then the cells behind should wait until (at least) a corresponding *layer* of polygons in all other cells have been projected. Furthermore, in order to avoid any expensive image-based tests, we would prefer to achieve such a goal using only object-space tests.

In order to achieve this goal of capturing global solidity, we extend the cell-projection framework by replacing the fixed insertion/deletion strategy queue, with a metric-based queue (*i.e.*, a priority queue), so that we can control how elements get pushed and popped based on a metric we can define. We call this priority queue, F , the front. The complete traversal algorithm is shown in Fig. 5. In order to completely describe it, we need to provide details on solidity metrics and its update

```

float function update_solidity( $B, A$ )
/* refer to Fig. 7 */
1.  $\rho_B = \frac{|A|}{\rho_{max}} + (\vec{v} \cdot \vec{n}_B) * \rho_A$ 
2. if (star_shaped( $\vec{v}, B$ ) == false)
3.    $\rho_B = \text{apply\_penalty\_factor}(\rho_B)$ 
4. return  $\rho_B$ 

```

Figure 6: Function *update_solidity*(). This function works as if transferring accumulated solidity from cell A into cell B . ρ_B is the solidity value to be computed for cell B . $|A|$ is the number of primitives in cell A . ρ_{max} is the maximum number of primitives in any cell. \vec{n}_B is the normal of the face shared by cells A and B . ρ_A is the accumulated solidity value for cell A . The maximum transfer happens if the new cell is well-aligned with the view direction \vec{v} , and in star-shaped position. If this is not the case, penalties will be incurred to the transfer.

strategies.

Solidity. The notion of a cell’s solidity is at the heart of our rendering algorithm shown in Fig. 5. At any given moment, cells are removed from the front (*i.e.*, priority queue F) in *solidity order*, that is, the cells with the smallest solidity are projected before the ones with larger solidity. The solidity of a cell B used in the rendering algorithm is not an intrinsic property of the cell by itself. Instead we use a set of conditions to roughly estimate the visibility likelihood of the cell, and make sure that cells more likely to be visible get projected before cells that are less likely to be visible.

The notion of solidity is related to how difficult it is for the viewer to see a particular cell. The actual solidity value of a cell B is defined in terms of the solidity of the cells that intersect the closure of a segment from the cell B to the eye point. The heuristic we have chosen to define the solidity value of our cells is shown in Fig. 6.

We use several parameters in computing the solidity value.

- The normalized number of primitives inside cell A , the neighboring cell (of cell B) that was just projected. This number, which is necessarily between 0 and 1, is $\frac{|A|}{\rho_{max}}$. The rationale is that the more primitives cell A contains, the more likely it is to obscure the cells behind it.
- Its position with respect to the viewpoint. We transfer a cell’s solidity to a neighboring cell based on how orthogonal the face that is shared between cells is to the view direction \vec{v} (see Fig. 7).

We also give preference to neighboring cells that are star-shaped [8] with respect to the viewpoint and the shared face. That is, we attempt to force the cells in the front to have their interior, *e.g.* their center point, visible from the viewpoint along a ray that passes through the face shared by the two cells. The reason for this is to avoid projecting cells (with low solidity values) that are occluded by cells in the front (with high solidity values) which have not been projected yet. This is likely to happen as the front expands away from an area in the scene where two densely occupied regions are nearby; we refer to such an area as a bottleneck. Examples of such areas can easily be seen in Figs. 1 and 4d. Actually, *forcing* the




Figure 7: Solidity Transfer. After projecting cell A , the traversal algorithm will add cells B and C to the front. Based upon the current viewing direction \vec{v} , cell B will accumulate more solidity from A than will cell C , however, C will likely incur the non-star-shaped penalty. \vec{n}_B and \vec{n}_C are the (respective) normals of the faces shared by the cell A 's neighboring cells. Refer to Fig. 6 for the transfer calculation.

front to be star-shaped at every step of the way is too limiting a rule. This would basically produce a visibility ordering for the cells (such as the one computed in [17, 5]). Instead, we simply *penalize* the cells in the front that do not maintain this star-shaped quality.[‡]

4 Implementation Details

We have implemented a system to experiment with the ideas presented in this paper. The code is a mix of Tcl/Tk and C++, and OpenGL for visualization. In all, we have about 6,000 lines of code. We briefly discuss the implementation of some of the main features of the code.

Data Structures. We need very simple data structures. Our current system only supports triangles for geometric primitives. Each triangle has pointers to its vertices, and a few flags, one of which is used to mark whether it has been rendered in the current scan. For the spatial tessellation, we represent each tetrahedron by pointers to its vertices, adjacency information is also required, as are a few flags for rendering purposes.

We keep the cells in the front in a priority queue. In our current implementation, we use an STL `set` to actually implement this data structure. Although simple and general, STL can add considerable overhead to an implementation. In our case, the number of cells in the front has been kept relatively small, and we have not noticed substantial slowdown due to STL.

Space Tessellation Code. We implemented our space tessellation code as three separate phases. First, vertices are inserted into a standard octree. We put a hard limit on the size of the octree leaf nodes based upon a user-defined error bound. Second, we compute a Delaunay triangulation of the centers of the leaves of the octree. For this, we used `qhull`, software written at the Geometry Center, University of Minnesota. Our highly constrained input is bound to have several

[‡]Even without this constraint, the algorithm seems to work fine, as can be seen in Figures 1 and 4d. Our 2D prototype does not make any use of star-shape constraints.

nasty degeneracies as all the points come from nodes of an octree, therefore we randomly perturbed these points and `qhull` had no problems handling them. Finally, we map triangles into the tetrahedra that contain them.

Rendering Loop Code. The rendering loop is basically a straightforward translation of the code in Fig. 5 into C++. Triangles are rendered very naively, one by one. We mark triangles as they are rendered, in order to avoid overdrawing triangles that get mapped to multiple cells. We also perform simple backface culling as well as view-frustum culling. We take no advantage of triangle-strips, vertex arrays, or other sophisticated OpenGL features.

Computing the Exact Visible Set. A number of benchmarking features are currently included in our implementation. One of the most useful is the computation of the actual *exact* visible set. We estimate V by using the well-know item buffer technique. In a nutshell, we color all the triangles with different colors, render them, and read the frame buffer back, recording which triangles contributed to the image rendered. After rendering, all the rank-1 triangles have their colors imprinted into the frame buffer.

Centroid-Ordered Rendering. In order to have a basis for comparison, we implemented a simple ordering scheme based on sorting the polygons with respect to their centroid, and rendering them in that order up to the specified budget. Our implementation of this feature is naive, and tends to be slow, since it needs to touch every single triangle in S .

5 Experimental Results

We performed a series of experiments in order to determine the effectiveness of PLP's visibility estimation. Our experiments typically consist of recording a flight path consisting of several frames for a given dataset, then playing back the final while varying the rendering algorithm used. We have three different strategies for rendering: (1) rendering every triangle in the scene at each frame, (2) centroid-based budgeting, or (3) PLP. During path playback, we also change the parameters when appropriate (*e.g.*, varying the polygon budget for PLP). Our primary benchmark machine is an IBM RS/6000 595 with a GXT800 graphics adapter. In all our experiments, rendering was performed using OpenGL with Z-buffer and lighting calculations turned on. In addition, all three algorithms perform view-frustum and backface culling to avoid rendering those triangles that clearly will not contribute to the final image. Thus, any benefits provided by PLP will be on top of the benefits provided by traditional culling techniques.

We report experimental results on two datasets:

City Model (CITY) The city model is composed of over 500K triangles (Fig. 8c). Each house has furniture inside, and while the number of triangles is large, the actual number of visible triangles per frame is quite small.

5 Car Body/Engine Model (5CBEM) This model has over 810K triangles (Fig. 9c). It is composed of five copies of a body and engine combination.

5.1 Preprocessing

Preprocessing involves computing an octree of the model, then computing a Delaunay triangulation of points defined by the octree (which is performed by calling `qhull`), and finally assigning the model geometric primitives to the spatial tessellation generated by `qhull`.

For the CITY model, preprocessing took 70 seconds, and generated 25K tetrahedra. Representing each tetrahedron requires less than 100 bytes (assuming the cost of representing the vertices is amortized among several tetrahedra), leading to a memory overhead for the spatial tessellation on the order of 2.5MB. Another source of overhead comes from the fact that some triangles might be multiply assigned to tetrahedra. The average number of times a triangle is referenced is 1.80, costing 3.6 MB of memory (used for triangle pointers). The total memory overhead (on top of the original triangle lists) is 6.1 MB, while storing all the triangles alone (the minimal amount of memory necessary to render them) already costs 50 MB. So, PLP costs an extra 12% in memory overhead.

For the 5CBEM model, preprocessing took 135 seconds (also including the `qhull` time), and generated 60K tetrahedra. The average number of tetrahedra that points to a triangle is 2.13, costing 14.7 MB of memory. The total memory overhead is 20 MB, and storing the triangles takes approximately 82 MB. So, PLP costs an extra 24% in memory overhead.

Since PLP’s preprocessing only takes a few minutes, the preprocessing is performed online, when the user requests a given dataset. We also support offline preprocessing, by simply writing the spatial tessellation and the triangle assignment to a file.

5.2 Rendering

We performed several rendering experiments. During these experiments, the flight path used for the 5CBEM is composed of 200 frames. The flight path for the CITY has 160 frames. For each frame of the flight path, we computed the following statistics:

- (1) the exact number of visible triangles in the frame, estimated using the item-buffer technique.
- (2) the number of visible triangles PLP was able to find for a given triangle budget. We varied the budget as follows: 1%, 2%, 5% and 10% of the number of triangles in the dataset.
- (3) the number of visible triangles the centroid-based budgeting was able to find under a 10% budget.
- (4) the number of wrong pixels generated by PLP.
- (5) time (all times are reported in seconds) to render the whole scene.
- (6) time PLP took to render a given frame.
- (7) time the centroid-based budgeting took to render a given frame.

Several of the results (in particular, (1), (2), (3), (5), and (6)) are shown in Table 1, and Figs. 8 and 9. The centroid rendering time (7) is mostly frame-independent, since the time is dominated by the sorting, which takes 6–7 seconds for the 5CBEM model, and 4–5 seconds for the CITY model. We

Dataset/Budget	1%	2%	5%	10%
City Model	51%	66%	80%	90%
5 Car Body/Engine Model	44%	55%	67%	76%

Table 1: Visible Coverage Ratio. The table summarizes ϵ_k for several budgets on two large models. The city model has 500K polygons, and the five car body/engine model has 810K polygons. For a budget of 1%, PLP is able to find over 40% of the visible polygons in either model.

collected the number of wrong pixels (4) on a frame-by-frame basis. We report worst-case numbers. For the CITY model, PLP gets as many as 4% of the pixels wrong; for the 5CBEM model, this number goes up, and PLP misses as many as 12% of the pixels, in any given frame.

PLP seems to do quite a good job at finding visible triangles. In fact, looking at Figs. 8a, and 9a, we see a remarkable resemblance between the shape of the curve plotting the exact visible set, and the PLP’s estimations. In fact, as the budget increases, the PLP curves seem to smoothly converge to the exact visible set curve. It is important to see that this is not a random phenomena. Notice how the centroid-based budgeting curve does not resemble the visible set curves. Clearly, there seems to be some relation between our heuristic visibility measure (captured by the solidity-based traversal), and actual visibility, which can not be captured by a technique that relies on distance alone.

Still, we would like PLP to do a better job at approximating the visible set. For this, it is interesting to see where it fails. In Figs 8d and 9d, we have 10%-budget images. Notice how PLP loses triangles in the back of the cars, (in Fig. 9d) since it estimates them to be occluded.

With respect to speed, PLP has very low overhead. For 5CBEM, at 1% we can render useful images at over 10 times the rate of the completely correct image, and for CITY, at 5% we can get 80% of the visible set, and still have four times faster rendering times.

Overall our experiments have shown that: (1) PLP can be applied to large data, without requiring large amounts of preprocessing; (2) PLP is able to find a large amount of visible geometry with a very low budget; (3) PLP is useful in practice, making it easier to inspect large objects, and in culling geometry that cannot be seen.

6 Conclusion and Future Work

In this paper, we proposed the Prioritized-Layered Projection algorithm. PLP renders geometry by carving out space along layers, while keeping track of the solidity of these layers as it goes along. PLP is very simple, requiring only a suitable tessellation of space where solidity can be computed (and is meaningful). The PLP rendering loop is a priority-based extension of the traversal used in depth-ordering cell projection algorithms developed originally for volume rendering.

We use PLP as our primary visibility-culling algorithm. Two things are most important to us. First, there is no offline preprocessing involved, that is, no need to simplify objects, pre-generate occluders, and so on. Second, its flexibility to adapt to multiple machines. In essence, in our application we were mostly interested in obtaining good image accu-

racy across a large number of machines with minimal time and space overheads. For several datasets, we can use PLP to render only 5% of a scene, and still be able to visualize over 80% of the visible polygons; if this is not enough, it is simple to adjust the budget for the desirable accuracy. A nice feature of PLP is that the visible set is stable, that is, the algorithm does not have major popping artifacts as it estimates the visible set from nearby viewpoints.

We see several other uses of PLP's rendering framework. A particularly intriguing one is to exploit PLP's ability to determine a large number of the visible polygons at low cost in terms of projected triangles (e.g., PLP can find over 40% of the visible polygons while only projecting 1% of the original geometry) to improve the performance of other occlusion-culling techniques. For instance, at each frame, HOM [21] projects geometry to create occlusion maps. Instead of relying on preprocessed simplified geometry, HOM could rely on PLP's output for its set of occluders.

Although PLP's budget-based framework is useful as is, there are several interesting avenues for new research. We would like to understand better the relation of the solidity measure to the actual set of rendered polygons. Changing our solidity value computation could possibly lead to even better performance. For example, accounting for front facing triangles in a given cell by considering their normals with respect to the view direction. The same is true for the mesh generation. Another class of open problems are related to further extensions in the front-update strategies. At this time, a single cell is placed in the front, after which the PLP traversal generates an ordering for all cells. We cut this tree by using a budget. It would be interesting to exploit the use of multiple initial seeds. Clearly, the best initial guess of what's visible, the easier it is to continue projecting visible polygons.

Another potential use of the PLP technique is in level-of-detail (LOD) selection. The PLP traversal algorithm can estimate the proportion of a model that is currently visible, which would allow us to couple visibility with the LOD selection process, as opposed to relying only on screen-space coverage tests.

Acknowledgements

We would like to thank Dirk Bartz and Michael Meissner for the city model; Bengt-Olaf Schneider for suggesting adding star-shape constraints to the front; Fausto Bernardini, Paul Borrel, William Horn, and Gabriel Taubin for suggestions and help throughout the project; and the Geometry Center of the University of Minnesota for ϵ hul1.

References

- [1] D. Bartz, M. Meissner, and T. Huettner. Extending graphics hardware for occlusion queries in OpenGL. In *Proc. Workshop on Graphics Hardware '98*, pages 97–104, 1998.
- [2] E. Catmull. *A Subdivision Algorithm for Computer Display of Curved Surfaces*. PhD thesis, Department of Computer Science, University of Utah, 1974.
- [3] J. H. Clark. Hierarchical geometric models for visible surface algorithms. *Communications of the ACM*, 19(10):547–554, October 1976.
- [4] D. Cohen-Or, G. Fibich, D. Halperin, and E. Zadicario. Conservative visibility and strong occlusion for view-space partitioning of densely occluded scenes. *Computer Graphics Forum*, 17(3):243–253, 1998.
- [5] J. Comba, J. Klosowski, N. Max, J. Mitchell, C. Silva, and P. Williams. Fast polyhedral cell sorting for interactive rendering of unstructured grids. *Computer Graphics Forum (Proc. of Eurographics '99)*, To appear.
- [6] S. Coorg and S. Teller. Real-time occlusion culling for models with large occluders. In *Proc. 1997 Sympos. Interactive 3D Graphics*, pages 83–90, 1997.
- [7] S. Coorg and S. Teller. Temporally coherent conservative visibility. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 78–87, 1996.
- [8] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.
- [9] D. P. Dobkin and S. Teller. Computer graphics. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 42, pages 779–796. CRC Press LLC, Boca Raton, FL, 1997.
- [10] S. E. Dorward. A survey of object-space hidden surface removal. *Internat. J. Comput. Geom. Appl.*, 4:325–362, 1994.
- [11] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics, Principles and Practice, Second Edition*. Addison-Wesley, Reading, Massachusetts, 1990.
- [12] T. Funkhouser and C. Séquin. Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments. In *SIGGRAPH '93*, pages 247–254, 1993.
- [13] N. Greene, M. Kass, and G. Miller. Hierarchical Z-buffer visibility. In *SIGGRAPH '93*, pages 231–240, 1993.
- [14] M. Held, J.T. Klosowski, and J.S.B. Mitchell. Evaluation of Collision Detection Methods for Virtual Reality Fly-Throughs. In *Proc. 7th Canad. Conf. Comput. Geom.*, pages 205–210, Québec City, Canada, 1995.
- [15] P. M. Hubbard. Approximating polyhedra with spheres for time-critical collision detection. *ACM Transactions on Graphics*, 15(3):179–210, July 1996.
- [16] J. Rohlf and J. Helman. IRIS performer: A high performance multiprocessing toolkit for real-time 3D graphics. In *SIGGRAPH '94*, pages 381–395, 1994.
- [17] C. Silva, J. Mitchell, and P. Williams. An interactive time visibility ordering algorithm for polyhedral cell complexes. In *Proc. ACM/IEEE Volume Visualization Symposium '98*, pages 87–94, 1998.
- [18] S. Teller and C. Séquin. Visibility preprocessing for interactive walkthroughs. In *SIGGRAPH '91*, pages 61–69, 1991.
- [19] P. Williams. Visibility ordering meshes polyhedra. *ACM Transactions on Graphics*, 11(2):103–126, April 1992.
- [20] H. Zhang. *Effective Occlusion Culling for the Interactive Display of Arbitrary Models*. PhD thesis, Department of Computer Science, University of North Carolina, Chappel Hill, 1998.
- [21] H. Zhang, D. Manocha, T. Hudson, and K. Hoff III. Visibility culling using hierarchical occlusion maps. In *SIGGRAPH '97*, pages 77–88, 1997.

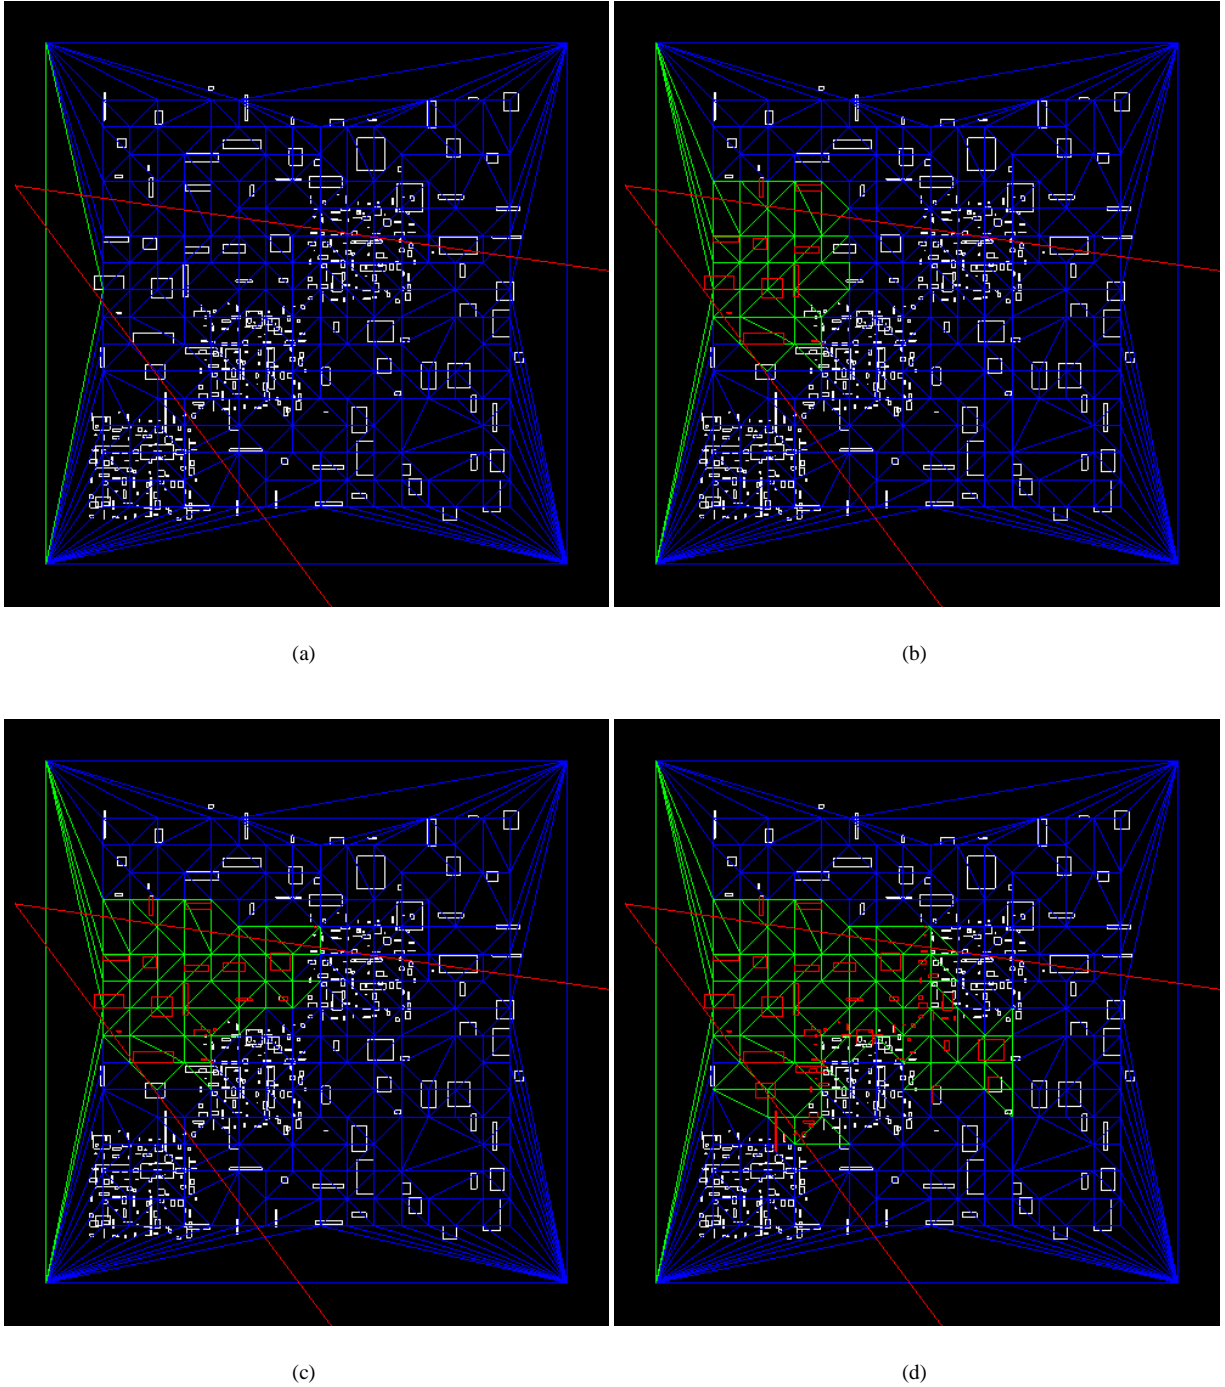

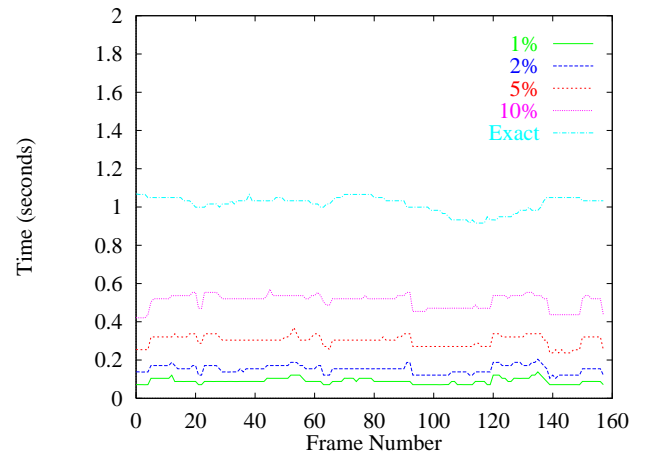



Figure 4: Priority-based traversal algorithm. In (a), the first cell, shown in green, gets projected. The algorithm continues to project cells based upon the solidity values. Note that the traversal, in going from (b) to (c), has delayed projecting those cells with a higher solidity value (*i.e.* those cells less likely to be visible) in the lower-left region of the view frustum. In (d), as the traversal continues, a higher priority is given to cells likely to have visible geometry, instead of projecting the ones inside of high-depth complexity regions. Note that the star-shaped criterion was not included in our 2D implementation.




(a)



(b)




(c)




(d)


Figure 8: CITY results. (a) The top curve, labelled Exact, is the number of visible triangles for each given frame. The next four curves are the number of the visible triangles PLP finds with a given budget. From top to bottom, budgets of 10%, 5%, 2%, and 1% are reported. The bottom curve is the number of visible triangles that the centroid sorting algorithm finds. (b) Rendering times in seconds for each curve shown in (a), with the exception of the centroid sorting algorithm, which required 4-5 seconds per frame. (c) Image of all the visible triangles. (d) Image of the 10% PLP visible set.




(a)



(b)



(c)



(d)

Figure 9: 5CBEM results. (a) The top curve, labelled Exact, is the number of visible triangles for each given frame. The next four curves are the number of the visible triangles PLP finds with a given budget. From top to bottom, budgets of 10%, 5%, 2%, and 1% are reported. The bottom curve is the number of visible triangles that the centroid sorting algorithm finds. (b) Rendering times in seconds for each curve shown in (a), with the exception of the centroid sorting algorithm, which required 6-7 seconds per frame. (c) Image of all the visible triangles. (d) Image of the 10% PLP visible set.