GigaWalk: Interactive Walkthrough of Complex Environments

William V. Baxter III Avneesh Sud Naga K. Govindaraju Dinesh Manocha
University of North Carolina at Chapel Hill
{baxter,sud,naga,dm}@cs.unc.edu
http://gamma.cs.unc.edu/GigaWalk

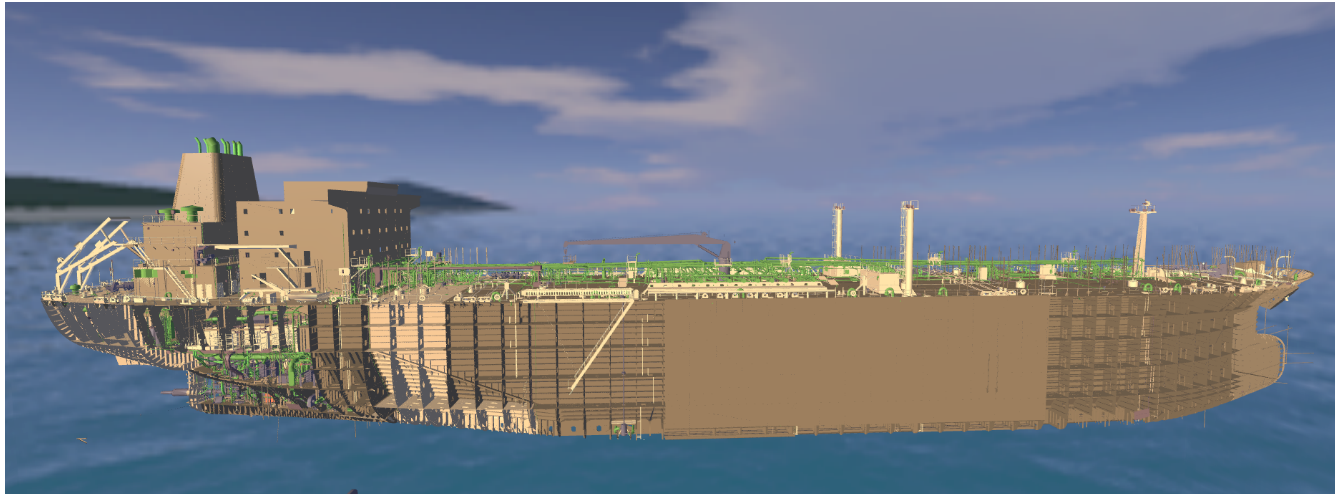

Figure 1: DoubleEagle Tanker: This 4 gigabyte environment consists of more than 82 million triangles and 127 thousand objects. Our algorithm can render it 11-50 frames per second on an SGI system with two IR2 graphics pipelines and three 300MHz R12000 CPUs.

Abstract

We present a new parallel algorithm for interactive walkthrough of complex, gigabyte-sized environments. Our approach combines occlusion culling and levels-of-detail and uses two graphics pipelines with one or more processors. We use a unified scene graph representation for multiple acceleration techniques, and we present novel algorithms for clustering geometry spatially, computing a scene graph hierarchy, performing conservative occlusion culling, and performing load-balancing between graphics pipelines and processors. The resulting system, GigaWalk, has been used to render CAD environments composed of tens of millions of polygons at interactive rates on an SGI Onyx system with two Infinite Reality rendering pipelines. Overall, our system's combination of levels-of-detail and occlusion culling techniques results in significant improvements in frame-rate over view-frustum culling or either single technique alone.

Keywords: Interactive display systems, parallel rendering, occlusion culling, levels-of-detail, Engineering Visualization.

1 Introduction

Users of computer-aided design and virtual reality applications often create and employ geometric models of large, complex 3D environments. It is not uncommon to generate gigabyte-sized datasets representing power plants, ships, airplanes, submarines and urban scenes. Simulation-based design and design review of such datasets benefits significantly from the ability to generate user-steered interactive displays or *walkthroughs* of these environments. Yet, rendering these environments at the required interactive rates and with high fidelity has been a major challenge.

Many acceleration techniques for interactive display of complex datasets have been developed. These include visibility culling, object simplification and the use of image-based or sampled representations. They have been successfully combined to render certain specific types of datasets at interactive rates, including architectural models [FKST96], terrain datasets [Hop98], scanned models [RL00] and urban environments [WWS01]. However, there has been less success in displaying more general complex datasets due to several challenges facing existing techniques:

Occlusion Culling: While possible for certain environments, performing exact visibility computations on large, general datasets is difficult to achieve in real time on current graphics systems [ESSS01]. Furthermore, occlusion culling alone will not sufficiently reduce the load on the graphics pipeline when many primitives are actually visible.

Object Simplification: Object simplification techniques alone have difficulty with high-depth-complexity scenes, as they do not address the problems of overdraw and fill load on the graphics pipeline.

Image-based Representations: There are some promising image-based algorithms, but generating complete samplings of large complex environments automatically and efficiently remains a difficult problem. The use of image-based methods can also lead to popping and aliasing artifacts.

1.1 Main Results

We present a parallel architecture that enables interactive rendering of complex environments comprised of many tens of millions of polygons. Initially, we pre-compute geometric levels-of-detail (LODs) and represent the dataset using a scene graph. Then at runtime we compute a *potentially visible set* (PVS) of geometry for each frame using a combination of view frustum culling and a two-pass hierarchical Z-buffer occlusion culling algorithm [GKM93] in conjunction with the pre-computed LODs. The system runs on two graphics rasterization pipelines and one or more CPU processors. Key features of our approach include:

1. A parallel rendering algorithm that is general and automatic, makes few assumptions about the model, and places no restrictions on user motion through the scene.

2. A clustering algorithm for generating a unified scene graph hierarchy that is used for both geometric simplification and occlusion culling.

3. A parallel, image-precision occlusion culling algorithm based on the hierarchical Z-buffer [GKM93, Gre01]. It uses *hierarchical occluders* and can perform conservative as well as approximate occlusion culling.

4. A parallel rendering algorithm that balances the computational load between two rendering pipelines and one or more processors.

5. An interactive system, GigaWalk, to render large, complex environments with good fidelity on two-pipeline graphics systems. The graphics pipelines themselves require only standard rasterization capabilities.

We demonstrate the performance of our system on two complex CAD environments: a coal-fired power plant (Fig. 2) composed of 13 million triangles, and a Double Eagle Tanker (Fig. 1) composed of over 82 million triangles. GigaWalk is able to render models such as these at 11-50 frames a second with little loss in image quality on an SGI Onyx workstation using two IR2 pipelines. The end-to-end latency of the system is typically 50-150 milliseconds.

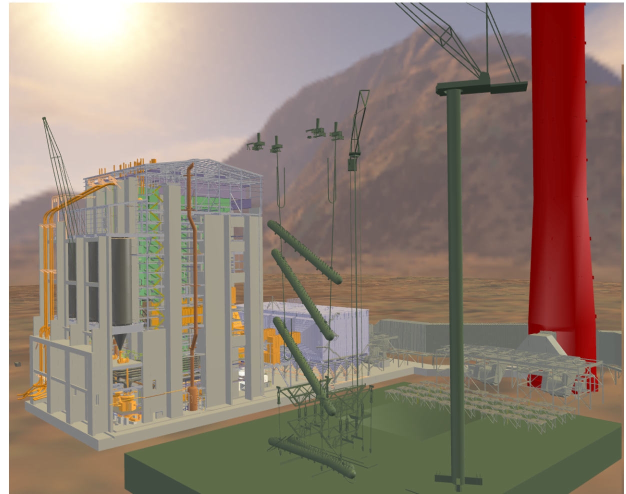

Figure 2: *Coal-Fired Power plant: This 1.7 gigabyte environment consists of over 13 million triangles and 1200 objects. GigaWalk can display it 12-37 frames per second on an SGI Onyx workstation using two IR2 graphics pipelines and three 300MHz R12000 CPUs.*

1.2 Organization

The rest of the paper is organized as follows. We give a brief survey of previous work in Section 2. Section 3 gives an overview of our approach. In Section 4 we describe the scene representation and preprocessing steps, including hierarchy computation. Section 5 presents the parallel algorithm for interactive display. We describe the system implementation and highlight its performance on complex models in Section 6.

2 Prior Work

In this section, we present a brief overview of previous research on interactive rendering of large datasets, including geometric simplification and occlusion culling algorithms, and other systems that have combined multiple rendering acceleration techniques.

2.1 Geometric Simplification

Simplification algorithms compute a reduced-polygon approximation of a model while attempting to retain the shape of the original. A survey of recent algorithms is presented in [Lue01].

Algorithms for simplifying large environments can be classified as either static (view-independent) or dynamic (view-dependent). Static approaches pre-compute a discrete series of levels-of-detail (LODs) in a view-independent manner [EM99, GH97, RB93, Sch97]. Erikson et al. [EMB01] presented an approach to large model rendering based on the hierarchical use of static LODs, or HLODs. We also use LODs and HLODs in our system.

At run-time, rendering algorithms for static LODs choose an appropriate LOD to represent each object based on the viewpoint. Selecting the LODs requires little run-time computation, and rendering static LODs

on contemporary graphics hardware is also efficient.

View-dependent, dynamic algorithms pre-compute a data structure that encodes a continuous range of detail. Examples include progressive meshes [Hop96, Hop97, XESV97] and hierarchies of decimation operations [LE97, ESV99]. Selection of the appropriate LOD is based on view-parameters such as illumination and viewing position. Overall, view-dependent LODs can provide better fidelity than static LODs and work well for large connected datasets such as terrain and spatially large objects. However, the run-time overhead is higher compared to static LODs, since all level-of-detail selection is done at the individual feature level (vertex, edge, polygon), rather than the object level.

2.2 Occlusion Culling

Occlusion culling methods attempt to quickly determine a PVS for a viewpoint by excluding geometry that is occluded. A recent survey of different algorithms is presented in [COCS01].

Several effective algorithms have been developed for specific environments. Examples include cells and portals for architectural models [ARB90, Tel92] and algorithms for urban datasets or scenes with large, convex occluders [CT97, HMC+97, SDDS00, WWS00, WWS01]. In this section, we restrict the discussion to occlusion culling algorithms for general environments.

Algorithms for occlusion culling can be broadly classified based on whether they are conservative or approximate, whether they use object space or image space hierarchies, and whether they compute visibility from a point or a region. Conservative algorithms compute a PVS that includes all the visible primitives, plus a small number of potentially occluded primitives [CT97, GKM93, HMC+97, KS01, ZMHH97]. The approximate algorithms identify most of the visible objects but may incorrectly cull some objects [BMH99, KS00, ZMHH97].

Object space algorithms can perform culling efficiently and accurately given a small set of large occluders, but it is difficult to perform the "occluder fusion" necessary to effectively cull in scenes composed of many small occluders. For these types of scenes, the image space algorithms typified by the hierarchical Z-buffer (HZB) [GKM93, Gre01] or hierarchical occlusion maps (HOM) [ZMHH97] are more effective.

Region-based algorithms pre-compute a PVS for each region of space to reduce the run-time overhead [DDTP00, SDDS00, WWS00]. This works well for scenes with large occluders, but the amount of geometry culled by a given occluder diminishes as the region sizes are increased. Thus there is a trade-off between the quality of the PVS estimation for each region and the memory overhead. These algorithms may be overly conservative and have difficulty obtaining significant culling in scenes including only small occluders.

2.3 Parallel Approaches

A number of parallel approaches based on multiple graphics pipelines have been proposed. These can provide scalable rendering on shared-memory systems or clusters of PCs. These approaches can by classified mainly as either object-parallel, screen-space-parallel, or frame-parallel [HEB+01, SFLS00]. Specific examples include distributing primitives to different pipelines by the screen region into which they fall (screen-space-parallel), or rendering only every Nth frame on each pipeline (frame-parallel).

Another parallel approach to large model rendering that shows promise is interactive ray tracing [AT99, WSB01]. The algorithm described in [WSB01] is able to render the Power Plant model at 4-5 frames a second with 640×480 pixel resolution on a cluster of seven dual processor PCs.

Garlick et al. [GBW90] presented a system for performing view-frustum culling on multiple CPUs in parallel with the rendering process. Their observation that culling can be performed in parallel to improve overall system performance is the fundamental concept behind our approach as well.

Wonka et al. [WWS01] presented a "visibility server" that performed occlusion culling to compute a PVS at run-time in parallel on a separate machine. Their system works well for urban environments; however, it relies on the *occluder shrinking* algorithm [WWS00] to compute the region-based visibility. This approach is effective only if the occluders are large and volumetric in nature.

2.4 Hybrid Approaches

The literature reports several systems that combine multiple techniques to accelerate the rendering of large models. For example, The BRUSH system [SBM+94] used LODs with hierarchical representation for large mechanical and architectural models. The UC Berkeley Architecture Walkthrough system [FKST96] combined hierarchical algorithms with object-space visibility computations [Tel92] and LODs for architectural models.

More recently, [ASVNB00] presented a framework that integrates occlusion culling and LODs. The crux of the approach is to estimate the degree of visibility of each object in the PVS and use that value both to select appropriate LODs and to cull. The method relies on decomposing scene objects into overlapping convex pieces (axis-aligned boxes) that then serve as individual "synthetic occluders". Thus the effective maximum occluder size depends on the largest axis-aligned box that will fit inside each object.

Another recent integrated approach uses a prioritized-layered projection visibility approximation algorithm with view-dependent rendering [ESSS01]. The resulting rendering algorithm seems a promising approach when approximate (non-conservative)

visibility is acceptable.

The UNC Massive Model Rendering (MMR) system [ACW+99] combined LODs with image-based impostors and occlusion culling to deliver interactive walkthroughs of complex models. A more detailed comparison with this system will be made later in Section 6.4.

Various proprietary systems exist as well, such as the one Boeing created in the 1990's to visualize models of large passenger jets. However, to the best of our knowledge, no detailed descriptions of this system are available, so it is difficult to make comparisons.

3 Overview

In this section, we give a brief overview of the main components of our approach. These components are simplification, occlusion culling, and a parallel architecture.

3.1 Model Simplification

Given a large environment, we generate a scene graph by clustering small objects, and partitioning large objects to create a spatially coherent, axis-aligned bounding box (AABB) hierarchy. Details of the hierarchy construction will be discussed in Section 4.

3.2 Parallel Occlusion Culling

At run-time, our algorithm performs occlusion culling, in addition to view frustum culling, based on the precomputed AABB scene graph hierarchy. We use a two-pass version of the hierarchical Z-buffer algorithm [GKM93] with a two-graphics-pipeline parallel architecture. In this architecture, occluders are rendered on one pipeline while the final interactive rendering of visible primitives takes place on the second pipeline. A separate software thread performs the actual culling using the Z-buffer that results from the occluder rendering. The architecture will be presented in detail in Section 5.

We chose to use the hierarchical Z-buffer (HZB) because of its good culling performance, minimal restrictions on the type of occluders, and for its ability to perform occluder fusion. Moreover, it can be made to work well without extra preprocessing or storage overhead by exploiting temporal coherence. The preprocessing and storage cost of GigaWalk is thus the same as that of an LOD-only system.

Occluder Selection: A key component of any occlusion culling algorithm is occluder selection, which can be accomplished in a number of ways. A typical approach uses solid angle to estimate a small set of good occluders [ZMHH97, KS00]. However, occluders selected according to such heuristics are not necessarily optimal in terms of the number of other objects they actually occlude. The likelihood of obtaining good occlusion can be increased by making the occluder set larger, but computational costs usually demand the set be as small as possible.

Our parallel approach, on the other hand, allows us to take advantage of the temporal coherence based occluder selection algorithm presented in [GKM93], which treats *all* the visible geometry from the previous frame as occluders for the current frame. This method makes use of frame-to-frame coherence and provides a good approximation to the foreground occluders for the current frame.

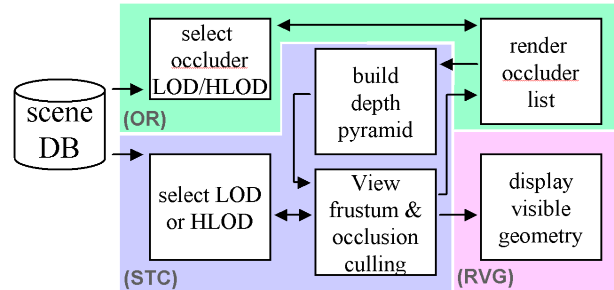

Figure 3: *System Architecture: Each color represents a separate process. The OR and RVG processes are associated with separate graphics pipelines, whereas the STC uses one or more processors.*

3.3 GigaWalk Architecture

Fig. 3 presents the overall architecture of our run-time system. It shows the three processes that run in parallel:

1. **Occluder Rendering (OR):** Using all the visible geometry from a previous frame as the occluder set, this process renders that set into a depth buffer. It runs on the first graphics pipeline.

2. **Scene Traversal, Culling and LOD Selection (STC):** This process computes the HZB using the depth buffer computed by OR. It traverses the scene graph, computes the visible geometry and selects appropriate LODs based on the user-specified error tolerance. The visible geometry is used by RVG for the current frame and OR for the next frame. It runs on one or more processors.

3. **Rendering Visible Scene Geometry (RVG):** This process renders the visible scene geometry computed by STC. It uses the second graphics pipeline.

More details of the run-time system are given in Sections 5 and 6.

4 Scene Representation

In this section, we present an object clustering-based algorithm to automatically compute a scene graph representation of the geometric dataset.

CAD datasets often consist of a large number of objects which are organized according to a functional, rather than spatial, hierarchy. By "object" we mean

simply the lowest level of organization in a model or model data structure above the primitive level. The size of objects can vary dramatically in CAD datasets. For example, in the Power Plant model a large pipe structure, which spans the entire model and consists of more than 6 million polygons, is one object. Similarly, a relatively small bolt with 20 polygons is another object. Our rendering algorithm computes LODs, selects them, and performs occlusion culling at the object level; therefore, the criteria used for organizing primitives into objects has a serious impact on the performance of the system. Our first step, then, is to redefine objects in a dataset based on criteria that will improve performance.

4.1 Unified Scene Hierarchy

Our rendering algorithm performs occlusion culling in two rendering passes: pass one renders occluders to create a hierarchical Z-buffer to use for culling, pass two renders the objects that are deemed visible by the HZB culling test. Given this two-pass approach, we could consider using a separate representation for occluders in pass one than for displayed objects in pass two [HMC+97, ZMHH97]. Using different representations has the advantage of allowing different criteria for partitioning and clustering for each hierarchy. Moreover, it would give us the flexibility of using a different error metric to create simplified occluders, one optimized to preserve occlusion power rather than visual fidelity.

Despite these potential advantages, we used a single unified hierarchy for occlusion culling and levels-of-detail based rendering. A single hierarchy offers the following benefits:

• **Simplicity:** A single representation leads to a simpler algorithm.

• **Memory And Preprocessing Overhead:** A separate occluder representation would increase the storage overhead, and increase the overall preprocessing cost. This can be significant for gigabyte datasets, depending upon the type of occluder representation used.

• **Conservative Occlusion Culling:** Our rendering algorithm treats the visible geometry from the previous frame as the occluder set for the current frame. In order to guarantee conservative occlusion culling, it is sufficient to ensure that exactly the same set of nodes and LODs in the scene graph are used by each process. Ensuring conservative occlusion culling when different representations are used is more difficult.

4.1.1 Criteria for Hierarchy

A good hierarchical representation of the scene graph is crucial for the performance of occlusion culling and the overall rendering algorithm. We use the same hierarchy for view frustum culling, occluder selection, occlusion tests on potential occludees, hierarchical simplification, and LOD selection. Though there has been considerable work on spatial partitioning and bounding volume hierarchies, including top-down and bottom-up strategies and spatial clustering, none of them seem to have addressed all the characteristics desired by our rendering algorithm. These include good spatial localization, object size, balance of the hierarchy, and minimal overlap between the bounding boxes of sibling nodes in the tree.

Bottom-up hierarchies lead to better localization and higher fidelity LODs. However, it is harder to use bottom-up techniques to compute hierarchies that are both balanced and have minimal spatial overlap between nodes. On the other hand, top-down schemes are better at ensuring balanced hierarchies and bounding boxes with little or no overlap between sibling nodes. Given their respective benefits, we use a hybrid approach that combines both top-down partitioning and hierarchy construction with bottom-up clustering.

4.2 Hierarchy Generation

In order to generate uniformly-sized objects, our preprocessing algorithm first redefines the objects using a combination of partitioning and clustering algorithms (see Fig. 5). The partitioning algorithm takes large objects and splits them into multiple objects. The clustering step groups objects with low polygon counts based on their spatial proximity. The combination of these steps results in a redistribution of geometry with good localization and emulates some of the benefits of pure bottom-up hierarchy generation. The overall algorithm proceeds as follows:

1. Partition large objects into sub-objects in the initial database (top-down)
2. Organize disjoint objects in and sub-objects into clusters (bottom-up)
3. Partition again to eliminate any uneven spatial clusters (top-down)
4. Compute an AABB bounding volume hierarchy on the final redefined set of objects (top-down).

Next, we present the algorithms for clustering and partitioning in detail.

4.2.1 Partitioning Objects

We subdivide large objects based on their size, aspect ratios and polygon count into multiple sub-objects, since long, thin objects with large bounding boxes are less likely to be occluded. The algorithm proceeds as follows:

1. Check if an object meets the splitting criteria :
    - the number of triangles is above a threshold t_1, and,
    - the size (bounding box diagonal) is greater than a threshold s_1, or

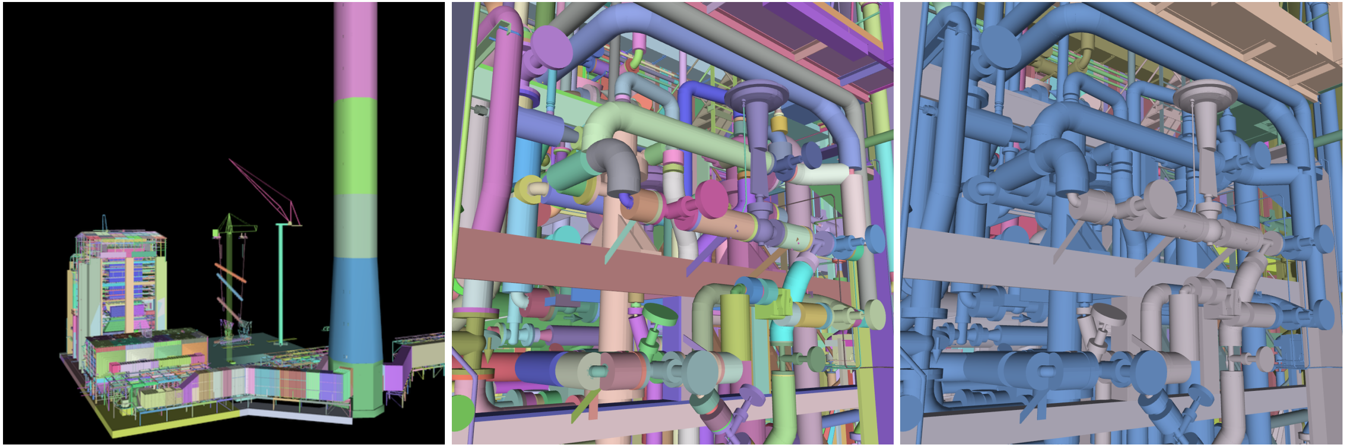

(a) Partitioning & Clustering on Power Plant (b) Original Objects in Double Eagle (c) Partitioning & Clustering on Double Eagle

Figure 4: *The image on the left shows the application of the partitioning and clustering algorithm to the Power Plant model. The middle image shows the original objects in the Double Eagle tanker model with different colors. The right image shows the application of the clustering algorithm on the same model. Each cluster is shown with a different color.*

- ratio of largest dimension of bounding box to smallest dimension is above a threshold r_1.

2. Partition the object along the longest axis of the bounding box. If that results in an unbalanced partition, choose the next longest axis.

3. Recursively split the child objects in the same manner.

See Section 6.3.1 for the parameter values used for this and the subsequent stages of our preprocess.

4.2.2 Clustering

Many approaches for clustering disjoint objects are based on spatial partitioning, such as adaptive grids or octrees. However, they may not work well for complex, irregular environments composed of a number of small and large objects. Rather we present an object-space clustering algorithm extending a computer vision algorithm for image segmentation [FH98]. The algorithm uses minimum spanning trees (MST) to represent clusters. It uses local spatial properties of the environment to incrementally generate clusters that represent global properties of the underlying geometry. The resulting clusters are neither too coarse nor too fine. The algorithm imposes this criteria by ensuring that it combines two clusters only if the internal variation in each cluster is greater than its external variation (by using Hausdorff metric). In each cluster, the maximum weight edge (which denotes the internal variation) of the MST representing the cluster denotes the maximum separation between any two "connected" objects in the cluster. A minimum weight edge connecting two different clusters (which denotes the external variation between the two clusters) estimates the separation between objects of one cluster with those of the other. In other words, it denotes the minimum radius of dilation necessary to connect at least one point of one cluster to that in another. The algorithm is similar to *Kruskal's* algorithm [Kru56] for generating a forest of minimum spanning trees. The overall algorithm proceeds as follows:

1. Construct a graph $G(V, E)$ on the environment with each object representing a vertex in the graph. Construct an edge between two vertices if the distance between the two vertices (objects) is less than a threshold D. The distance function is defined as the shortest distance between the bounding boxes of the two objects. If the two boxes are overlapping, then it is zero.

2. Sort E into $\pi = (o_1, o_2, \ldots, o_k)$, $o_i \in E$, $i = 1, \ldots, k$ in a non-decreasing order based on edge weights, $w(o_i)$. Start with a forest F^0 where each vertex v_i represents a cluster.

3. Repeat Step 4 for the set of edges $o_q = o_1, \ldots, o_k, k = \|E\|$.

4. Construct forest F^q from F^{q-1} as follows. Let $o_q = (v_i, v_j)$, i.e., edge o_q connects vertices v_i and v_j. If there is no path from v_i to v_j in F^{q-1} and $w(o_q)$ is small compared to the internal variation of components containing v_i and v_j, and if the bounding volume of the resultant component is less than a maximum volume threshold, then add o_q to F^{q-1} to obtain F^q, otherwise add nothing. Mathematically, if $C_i^{q-1} \neq C_j^{q-1}$ and $w(o_q) \leq \mathcal{MI}(C_i^{q-1}, C_j^{q-1})$, then $F^q = F^{q-1} \cup\ o_q$ where C_i^{q-1} denotes the cluster containing vertex v_i in F^{q-1} and C_j^{q-1} is the cluster containing vertex v_j.

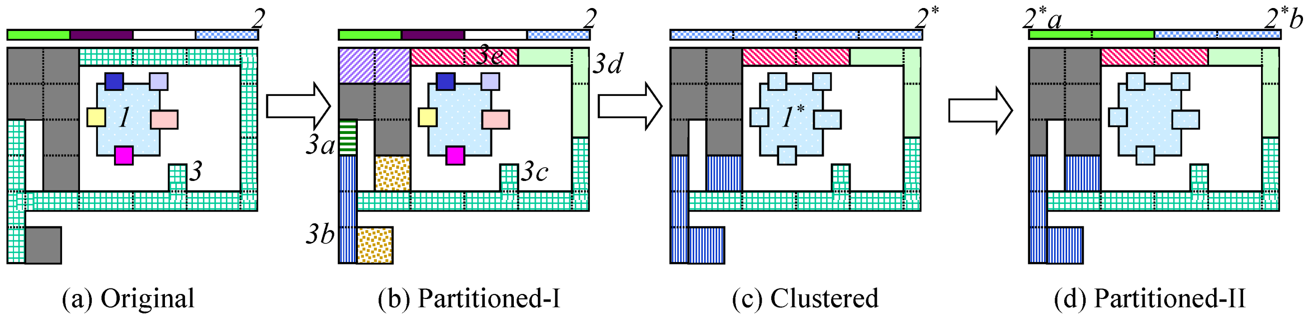

(a) Original (b) Partitioned-I (c) Clustered (d) Partitioned-II

Figure 5: *Our clustering and partitioning process applied to a 2D example. Each different color represents a different object at the end of a stage. (a) The model's original objects. This object distribution captures a number of features common in CAD models in which objects are defined by function rather than by location. (b) The initial partitioning stage serves to split objects with large bounding boxes. This prevents objects like* 3*, whose initial bounding box intersects most of the others, from causing clustering to generate just one large cluster. (c) After clustering, the group of small objects around* 1 *have all been merged to form* 1^**. The row of objects,* 2*, have also been merged into one cluster,* 2^**, as well, but one which has a poor aspect ratio. (d) The final partition splits* 2^* *into two separate objects.*

else $F^q = F^{q-1}$, where

$$\mathcal{MI}(C_1, C_2) = \min(\mathcal{I}(C_1) + K/\|C_1\|, \mathcal{I}(C_2) + K/\|C_2\|)$$

and

$$\mathcal{I}(C) = \max_{e \in MST(C,E)} w(e) \qquad (1)$$

The $K/\|C_i\|$ terms bias the results toward clusters of cardinality bounded by $O(K)$, where K is user-specified. We set a maximum volume threshold V to ensure final clusters are not too large in size. It is reasonably fast in practice and generates good spatial clusters. Fig. 4 shows its application to the Power Plant and Double Eagle models.

Overall, the clustering algorithm successfully merges small objects into larger groups with good localization. It improves the performance of the culling algorithm as well as the final image fidelity.

4.2.3 Partitioning Clusters

We finally repartition clusters with an uneven distribution of geometry, splitting objects about their center of mass. This is similar to the partitioning algorithm presented in Section 4.2.1. However, we use higher thresholds for the size and triangle count to avoid splitting clusters back into small objects, but tighter bounds for the aspect ratio to force splitting of uneven clusters.

4.2.4 Hierarchy Generation

We compute a standard AABB bounding volume hierarchy in a top-down manner on the set of redefined objects generated after clustering and partitioning. The bounding boxes are assigned to left and right branches of the hierarchy using their geometric centers to avoid overlap. The redefined objects become the leaf nodes in the AABB hierarchy.

4.3 HLOD Generation

Given the above scene graph, the algorithm computes a series of LODs for each node. The HLODs are computed in a bottom-up manner. The HLODs of the leaf nodes are standard static LODs, while the HLODs of intermediate nodes are computed by combining the LODs of the nodes with the HLODs of node's children. We use a topological simplification algorithm to merge disjoint objects [EM99].

The majority of the pre-computation time is spent in LOD and HLOD generation. The HLODs of an internal node depend only on the LODs of the children, so by keeping only the LODs of the current node and its children in main memory, HLOD generation is accomplished within a small memory footprint. Specifically, the memory usage is given by

$$\texttt{main_memory_footprint} \leq sizeof(AABBHierarchy) + \max_{N_i \in SG}(sizeof(N_i) + \sum_{C_j \in Child(N_i)} sizeof(C_j))$$

4.4 HLODs as Hierarchical Occluders

Our occlusion culling algorithm uses LODs and HLODs of nodes as occluders to compute the HZB. They are selected based on the maximum screen-space pixel deviation error on object silhouettes.

The HLODs our rendering algorithm uses for occluders can be thought of as "hierarchical occluders". A hierarchical occluder associated with a node i is an approximation of a group of occluders contained in the subtree rooted at i. The approximation provides a lower polygon count representation of a collection of object-space occluders. It can also be regarded as object-space

occluder fusion.

5 Interactive Display

In this section, we present our parallel rendering architecture for interactive display of complex environments. Here we describe in detail the operations performed by each of the two graphics pipelines and each of the three processes: occluder rendering (OR), scene traversal and culling (STC) and rendering visible geometry (RVG), which run synchronously in parallel (as shown in Fig. 6).

5.1 Run-time Architecture

The relationship between different processes and the tasks performed by them is shown in Fig. 3.

5.1.1 Occluder Rendering

The first stage for a given frame is to render the occluders. The occluders are simply the visible geometry from a previous frame. By using this temporal coherence strategy, the load on the two graphics pipelines is essentially balanced, since they render the exact same geometry, just shifted in time. The culling and LOD selection performed for displaying frame i naturally results in an occluder set for frame $i + 1$ that is of manageable size. A brief pseudo-code description is given in Algorithm 5.1.

```
Occluder_Render(δ,frame i)
  • get current instantaneous camera position
    (camera i)
  • while (more nodes on node queue from STC (i-1))
     * pop next node off the queue
     * select LOD/HLOD for the node according to
       error tolerance, δ, using camera i
     * render that LOD/HLOD into Z-buffer i
  • read back Z-buffer i from graphics hardware
  • push Z-buffer i onto queue for STC i
  • push camera i onto queue for STC i
```

ALGORITHM 5.1: Occluder_Render.

Since the list of visible geometry for rendering comes from the culling stage (STC), and STC gets its input from this process (OR), a startup procedure is required to initialize the pipeline and resolve this cyclic dependency. During startup, the OR stage is bypassed on the first frame, and STC generates its initial list of visible geometry without occlusion culling.

5.1.2 Scene Traversal, Culling and LOD Selection

The STC process first computes the HZB from the depth buffer output from OR. It then traverses the scene graph, performing view-frustum culling, occlusion culling and LOD error-based selection in a recursive manner. The LOD selection proceeds exactly as in [EMB01]: recursion terminates at nodes that are either culled, or which meet the user-specified pixel-error tolerance. A pseudo-code description is given in Algorithm 5.2. The occlusion culling is performed

```
Scene_Traversal_Cull(ε, frame i)
  • get Z-buffer i from OR i via Z-buffer queue
  • build HZB i
  • get camera i from OR i queue
  • push copy of camera i onto queue for RVG i
  • set NodeList[i] = Root(SceneGraph)
  • while (NotEmpty(NodeList[i]))
     * node = First(NodeList[i])
     * set NodeList[i] = Delete(NodeList[i],node)
     * if (View_Frustum_culled(node)) then next;
     * if (Occlusion_Culled(node)) then next;
     * if HLOD_Error_Acceptable(ε,node) then
        – push node onto queue for OR i + 1;
        – push node onto display queue for RVG i;
     * else
       set NodeList[i] = Add(NodeList[i],
                             Children(node));
```

ALGORITHM 5.2: Scene_Traversal_Cull.

by comparing the bounding box of the node with the HZB. It can be performed in software or can make use of hardware-based queries as more culling extensions become available.

5.1.3 Rendering Visible Scene Geometry

All the culling is performed by STC, so the final render loop has only to rasterize the nodes from STC as they are placed in the queue. See Algorithm 5.3.

```
Render_Visible_Scene_Geometry(frame i)
  • get camera i from STC i queue
  • while (more nodes on queue from STC i)
     • pop node off queue
     • render node
```

ALGORITHM 5.3: Render_Visible_Scene_Geometry.

5.2 Occluder Selection

Ideally, the algorithm uses the visible geometry from the previous frame ($i - 1$) as the occluders for the current frame to get the best approximation to the current foreground geometry. However, using the previous frame's geometry can lead to bubbles in the pipeline, because of the dependency between the OR and STC stages: STC must wait for OR to finish rendering the occluders before it can begin traversing the scene graph and culling. Fortunately, using the visible geometry from two frames previous can eliminate that dependency, and still provides a good approximation to the visible geometry for most interactive applications.

5.3 Trading Fidelity for Performance

The user can trade off fidelity for better performance in a number of ways. The primary control for achieving higher frame rates is the allowable LOD pixel error (see Figure 7).

Our system has been designed primarily to offer

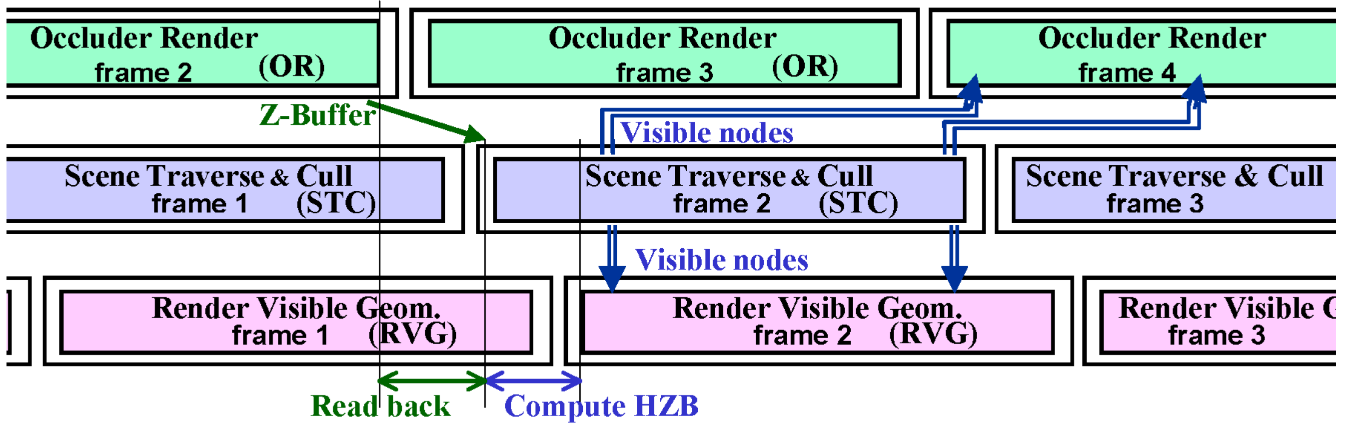

Figure 6: *Timing relationship between different processes. The arrows indicate data passed between processes during the computation of frame 2. Along with the other data indicated, viewpoints also travel through the pipeline according to the frame numbers. This diagram demonstrates the use of occluders from frame $i-2$ rather than $i-1$ (see Section 5.2).*

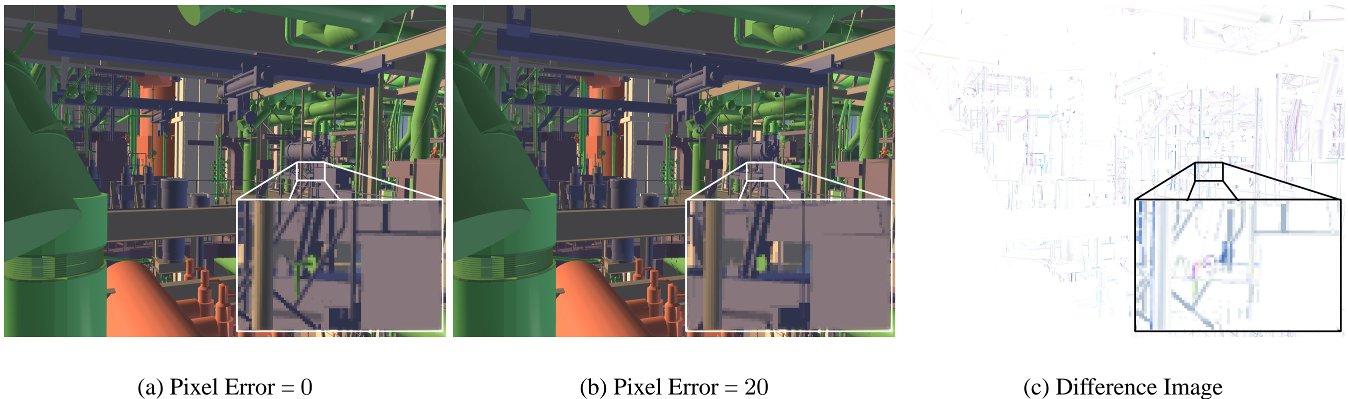

(a) Pixel Error = 0 (b) Pixel Error = 20 (c) Difference Image

Figure 7: *The Engine Room in the Double Eagle Tanker displayed without and with HLODs. The inset shows a magnification of one region. Original resolution 1280×960.*

conservative occlusion culling, and we report all of our results based on this mode of operation. Our system can guarantee conservative culling results for two reasons: 1) the underlying HZB algorithm used is itself conservative, and 2) for a given frame i we choose the exact same set of LODs for both OR and STC stages. By choosing the same LODs, we ensure that the Z-buffer used for culling is consistent with the geometry it is used to cull. Without this selection algorithm, conservativity is not guaranteed.

We have also modified our run-time pipeline in a number of ways to optionally increase frame rate or decrease latency, by allowing the user to relax the restriction that occlusion culling be performed conservatively:

• **Asynchronous rendering pipeline:** Rather than waiting for the next list of visible geometry from the culling stage (STC) to render frame $i+1$, the render stage (RVG) can instead proceed to render another frame, still using the geometry from frame i, but using the most recent camera position, corresponding to the user's most up-to-date position. This modification eliminates the extra frame of latency introduced by our method. The main drawback is that it may introduce occlusion errors that, while typically brief, are potentially unbounded when the user moves drastically.

• **Nth Farthest Z Buffer Values:** The occlusion culling can be modified to use not the farthest Z values in building the depth pyramid, but the Nth farthest [Gre01], thereby allowing for approximate "aggressive" culling.

• **Lower HZB resolution for occluder rendering:** The pixel resolution of the OR stage can be set smaller than that of the RVG stage. If readback from the depth buffer or HZB computation is relatively slow, this can improve the performance. However, using a lower resolution source for HZB allows for the possibility of depth buffer aliasing artifacts that can manifest themselves as small occlusion errors. In practice, however, we have not been able to visually detect any such errors when using OR depth buffers with as little as half the RVG resolution.

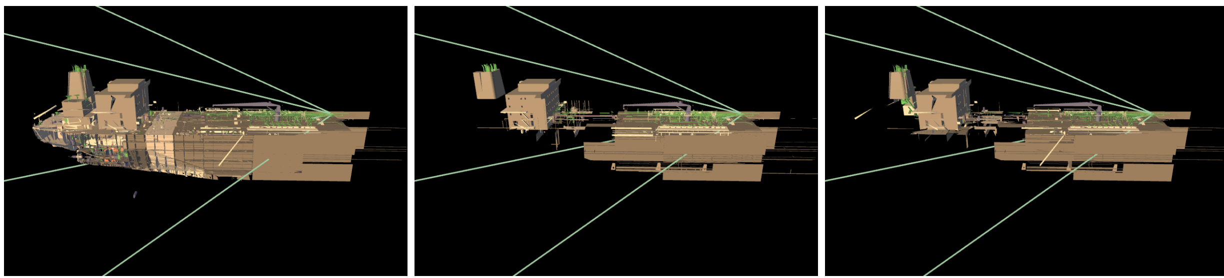

(a) Polygon Count = 202666 (b) Polygon Count = 3578485 (c) Polygon Count = 61771

Figure 8: *Comparison between different acceleration techniques from the same viewpoint. (a) Rendered with only HLODs. (b) Rendered with only HZB occlusion culling. (c) Rendered with GigaWalk using HLODs and HZB occlusion culling.*

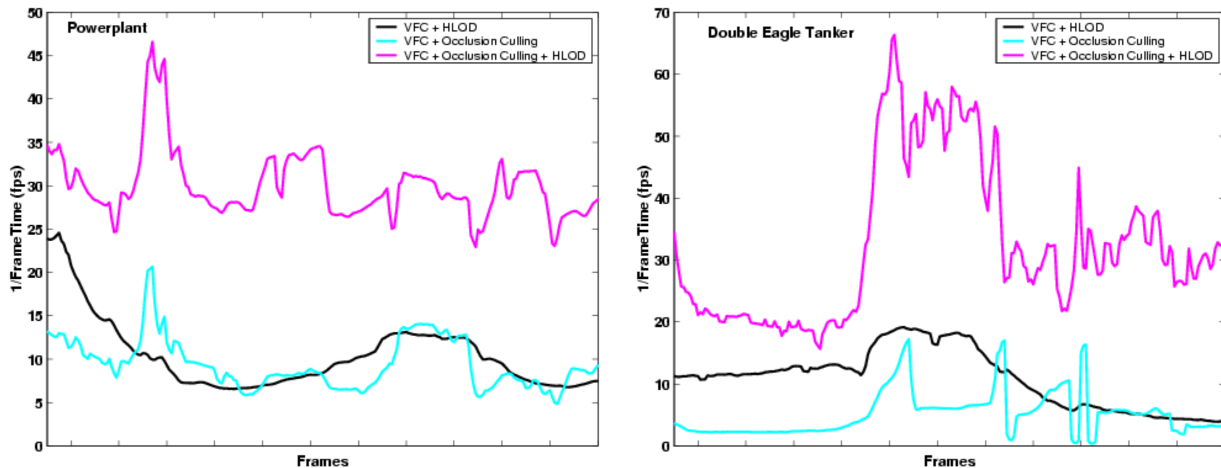

Figure 9: *Comparison of different acceleration algorithms on a path in the Power Plant model (left) and Double Eagle (right). The Y-axis shows the instantaneous frame rate. The combination of HLODs + occlusion culling results in 2-5 times improvement over using only one of them.*

6 Implementation and Performance

We have implemented our parallel rendering algorithm on a shared-memory multiprocessor machine with dual graphics rasterization pipelines: an SGI Onyx workstation with 300MHz R12000 MIPS processors, Infinite Reality (IR2e) graphics boards, and 16GB of main memory. Our algorithm uses three CPUs and two graphics pipelines.

All of the inter-process communication is implemented using a simple templated producer-consumer shared queue data structure of under 100 lines of C++ code. Each stage (OR,STC,RVG) is connected with the others using one or more instances of this queue data structure. Synchronization between the processes is accomplished by pushing sentinel nodes onto the shared queues to delimit the data at the end of a frame. The scene graph resides in shared memory, and is accessible to all processes.

We have tested the performance of GigaWalk on two complex environments, a coal-fired Power Plant (shown in Fig. 2) and a Double Eagle Tanker (shown in Fig. 1). The details about these environments are shown in Table 1. In addition to the model complexity, the table also lists the object counts after the clustering and partitioning steps.

6.1 Improvement in Frame Rate

GigaWalk is able to render our two example complex environments at interactive rates from most viewpoints. The frame rate varies from 11 to 50 FPS. It is more than 20 frames a second from most viewpoints in the scene. We have recorded and analyzed some example paths through these models (as shown in the video). In Fig. 9, we show the improvement in frame rate for each environment. The graphs compare the frame rate for each individual rendering acceleration technique alone and for the combination. Table 2 shows the average speedups obtained by each technique over the same path. The comparison between the techniques for a given viewpoint is shown in Fig. 8.

Env	Poly $\times 10^6$	Object Count: Init $\times 10^4$	Part[1] $\times 10^4$	Clust $\times 10^3$	Part[2] $\times 10^5$
PP	12.2	0.12	6.95	3.33	0.38
DE	82.4	12.7	2.21	2.31	1.2

Table 1: *A breakdown of the complexity of each environment.* **Poly** *is the polygon count.* **Init** *is the number of objects in the original dataset. The algorithm first partitions (***Part**^{1}*) objects into sub-objects, then generates clusters (***Clust***), and finally partitions large uneven spatial clusters* **Part**^{2}*. The table shows the object count after each step.*

Model	Average FPS: OCH	HLOD+VFC	OCC+VFC	VFC
PP	30.67	9.55	9.48	1.15
DE	29.43	9.76	3.27	0.02

Table 2: *Average frame rates obtained by different acceleration techniques over the sample path.* **FPS** = *Frames Per Second,* **HLOD** = *Hierarchical levels of detail,* **OCH** = *Occlusion culling with HLOD,* **OCC** = *Occlusion Culling,* **VFC** = *View Frustum Culling*

6.2 System Latency

Our algorithm introduces a frame of latency to rendering times. Latency can be a serious issue for many interactive applications like augmented reality. Our approach is best suited for latency-tolerant applications, namely walkthroughs of large synthetic environments on desktop or projection displays. The end-to-end latency in our system varies with the frame rate. It is typically in the range 50-150 ms. The high end of this range is achieved when the frame rate dips close to 10 frames a second. This latency is within the range that most users can easily adapt to (less than 300 ms) without changing their interaction mode with the model [EYAE99].

In many interactive applications, the dominant component of latency is the frame rendering time [EYAE99]. Through the use of our two-pass occlusion culling technique, our rendering algorithm improves the frame rate by a factor of 3-4. As a result, the overall system latency is decreased, in contrast to an algorithm that does not use occlusion culling.

6.3 Preprocessing

This section reports the values we used for the preprocessing parameters mentioned in Section 4 and gives details about the amount of time and memory used by our preprocess.

6.3.1 Preprocessing Parameters

For Partition-I (Sec. 4.2.1), empirically, we have obtained good results with t_1=1000, s_1=0.1$\times$`model_dimension`, r_1=2.5.

For Clustering (Sec. 4.2.2), we use $D = 0.1 \times$ `max_bounding_box_dimension`. We can also work on a complete graph but choosing a threshold in general works well and reduces the computational complexity. We choose $V = 10^{-3} \times$ `total_scene_bounding_box_volume` as the maximum volume threshold. For K we use 1500, which prevents too many closely-packed objects from all merging into a single cluster.

For Partition-II (Sec. 4.2.3) we use $s_2 = 2s_1$ and triangle count $t_2 = 10t_1$, but tighter bounds for the aspect ratio $r_2 = 0.5r_1$.

6.3.2 Time and Space Requirements

The preprocessing was done on a single-processor 2GHz Pentium 4 PC with 2GB RAM. The preprocessing times for the Double Eagle model were: 45 min for Partition-I, 90 min for Clustering, 30 min for Partition-II, 32.5 hours for out of core HLOD generation and 12 min for the AABB hierarchy generation. The size of the final HLOD scene graph representation is 7.6GB which is less than 2 times the original data size. The AABBTree hierarchy occupies 7MB space.

The main memory requirement for partitioning and clustering is bounded by the size of the largest object/cluster. For the Double Eagle it was less than 200MB for partitioning, 1GB for clustering and 300MB for out of core HLOD generation.

6.4 Comparison with Earlier Approaches

A number of algorithms and systems have been proposed for interactive display of complex environments. These include specialized approaches for architectural, terrain and urban environments, as highlighted in Section 2. Given low depth complexity scenes, or scenes composed of large or convex occluders (e.g, architectural or urban models), our general approach is not likely to perform any better than special-purpose algorithms designed specifically to exploit such features. In the rest of this section, we compare GigaWalk with earlier systems that can handle general environments.

Some of the commonly used algorithms for general environments are based on a scene graph representation and using a combination of LODs or HLODs at each node [Clar76]. The SGI's Performer toolkit [RH94] and IBM's BRUSH system [SBM$^+$94] provided this capability. The BRUSH system used the LODs generated by a vertex clustering algorithm [RB93] that did not provide good error bounds on the quality of resulting approximations. Cohen et al. [Cohe96] used LODs generated using the "simplification envelopes" algorithm and integrated them into the Performer framework. However, the resulting algorithm was unable to compute drastic simplifications of large environments. Erikson et al. [EMB01] used a combination of LODs and HLODs computed using GAPS [EM99] for different nodes in the scene graph. However, none of these

approaches performed occlusion culling. As a result, they were unable to render high depth complexity models at interactive rates with good fidelity. GigaWalk also uses GAPS to compute LODs and HLODs of objects. However, our approach for computing a unified hierarchy that is used for LOD and HLOD selection at runtime, and occlusion culling is quite different from that presented in [EMB01].

The MMR system combined LODs and occlusion culling for near-field geometry with image-based textured meshes to approximate the far-field, in a cell-based framework [ACW+99]. While the combination of techniques proved capable of achieving interactive frame rates, the system had some drawbacks. No good algorithms are known for automatic computation of cells for a large and general environment. The MMR system used some model-dependent features (e.g. walkways in the Powerplant) for cell computation. On the other hand, the preprocessing and scene graph computation in GigaWalk is fully automatic. The image-based far-field representations used in the MMR system can result in dramatic popping and distortion when switching between different cells. Moreover, the memory overhead and preprocessing cost of creating six meshes and textures per-cell was quite quite high. The MMR system used a single graphics pipeline for occlusion culling as well as rendering the visible geometry and textured meshes. It used some preprocessing techniques to estimate the depth complexity of near-geometry in each cell and only used occlusion culling when the viewer is in a cell with high depth complexity. However, the overall benefit of occlusion culling was rather limited because of these heuristics and a single graphics pipeline based implementation. The cell based spatial decomposition in the MMR system resulted in a simple out-of-core rendering and prefetching algorithm. The runtime memory footprint was much smaller than the total size of the model and associated data structures (e.g. LODs and image-based representations). On the other hand, GigaWalk assumes that the entire scene graph and the LODs and HLODs are loaded in the main memory.

7 Conclusions and Lessons Learned

We have presented an approach to rendering interactive walkthroughs of complex 3D environments. The algorithm features a novel integration of conservative occlusion culling and levels-of-detail using a parallel algorithm. We have demonstrated a new parallel rendering architecture that integrates these acceleration techniques on two graphics pipelines and one or more processors. Finally, we have presented novel algorithms for partitioning and clustering, and generating an integrated hierarchy suitable for both HLODs and occlusion culling. We have demonstrated the performance of our system, GigaWalk, on two complex CAD environments. To the best of our knowledge, GigaWalk is the first system that can render such complex environments at interactive rates with good fidelity, i.e. no popping or aliasing artifacts.

There are many complex issues with respect to the design and performance of systems for interactive display of complex environments. These include load balancing, extent of parallelism and scalability of the resulting approach, the effectiveness of occlusion culling and issues related to loading and managing large datasets.

7.1 Load Balancing

There is a trade-off between the depth of scene graph, which is controlled by the choice of minimum cluster size, and the culling efficiency. Smaller bounding boxes lead to better culling since more boxes can be rejected, so less geometry is sent to RVG. On the other hand, more boxes increases the cost of scene graph traversal and culling in STC. In our system, scene traversal and object culling (STC) operate in parallel with rendering (OR and RVG). If our performance bottleneck is the rendering processes (RVG and OR), we can shift the load back to the culling (STC) process by creating a finer partitioning. Conversely, we can use a coarser partitioning to move the load back to RVG and OR. Thus, the system can achieve load balancing between different processes running on the CPUs or graphics pipelines by changing the granularity of partitioning.

7.2 Parallelism

Parallel graphics hardware is increasingly being used to improve the rendering performance of a walkthrough system. Generally, though, the speed-up obtained from using N pipelines is no more than a factor of N. Using a second pipeline for occlusion culling (i.e. $N = 2$), however, enables GigaWalk to achieve more than two times speed-up for scenes with high depth complexity. For low depth complexity scenes there is little or no speed-up, but there is no loss in frame rate as the occlusion culling is performed using a separate pipeline. However, our parallel algorithm introduces a frame of latency.

Note also that other parallel approaches [HEB+01, SFLS00] are fundamentally orthogonal to our approach, and could potentially be used in conjunction with our architecture as black-box replacements for the OR and RVG rendering pipelines.

7.3 Load Times

One of the considerations in developing a walkthrough system to render gigabyte datasets is the time taken to load gigabytes of data from secondary storage, which can be many hours. To speed up the system we have implemented an on-demand loading system. Initially the system takes a few seconds to load the skeletal representation of the scene graph with just bounding boxes. Once loaded, the user commences the walk-

through while a fourth, asynchronous background process automatically loads the geometry for the nodes in the scene graph that are visible. We have found that adding such a feature is very useful in terms of system development and testing its performance on new complex environments.

8 Limitations and Future Work

Our current implementation of GigaWalk has many limitations. The current system works only for static environments, and it would be desirable to extend it to dynamic environments as well, perhaps with a strategy similar to the one proposed in [EMB01].

The memory overhead of GigaWalk can be high. In the current implementation, viewing an entire model requires loading the scene graph and HLODs. It would be useful to develop an out-of-core rendering system that uses a finite memory footprint and uses prefetching techniques to load the visible nodes.

The preprocessing time for our largest dataset, the Double Eagle, was also higher than desired. Since most nodes in the scene graph are non-overlapping, the LODs can be generated independently. Thus the algorithm could compute the LODs and HLODs in parallel, using multiple threads. This could improve the preprocessing performance considerably, reducing the 32.5 hours spent on the Double Eagle to overnight.

One interesting possibility for further increasing the efficiency of GigaWalk would be to use speculative rendering in the RVG process to eliminate any time lost to pipeline stalls when RVG needs to wait for visible nodes from the STC process. Since we know that the geometry rendered during the last frame is likely to still be visible in the current frame, the RVG process could proceed to render nodes from the old list whenever the STC is not ready with new data. It would also be useful to use a progressive occlusion culling algorithm [GKM93], as opposed to a two-pass HZB algorithm. It may need extra support from the graphics hardware.

The algorithm described in this paper guarantees image quality in terms of a bound on screen-space LOD error, but there are no guarantees on the frame rates. The current system would be improved by the addition of a target-frame-rate rendering mode. Furthermore, the current system's use of static LODs and HLODs leads to some popping when switching between detail levels. We would like to explore view-dependent or hybrid view-dependent/static LOD-based simplification approaches that can improve the fidelity of our geometric approximations without increasing the polygon count, while decreasing popping artifacts.

Finally we are interested in adapting GigaWalk to a distributed system architecture using two graphics-capable PCs, connected using a standard network. We expect that the network bandwidth requirement will be less than 10-megabits per second, enabling the system to run over a standard ethernet LAN. Our current implementation runs on a shared-memory SGI workstation. We believe that taking advantage of the performance of current commodity PC graphics hardware will both improve the rendering performance and enable wider use of the resulting system.

Acknowledgments

Our work was supported in part by ARO Contract DAAD19-99-1-0162, NSF award ACI 9876914, ONR Young Investigator Award (N00014-97-1-0631), a DOE ASCI grant, and by Intel Corporation.

The Double Eagle model is courtesy of Rob Lisle, Bryan Marz, and Jack Kanakaris at NNS. The Power Plant environment is courtesy of an anonymous donor. We would like to thank Carl Erikson, Brian Salomon and other members of UNC Walkthrough group for their useful discussions and support. Special thanks to Dorian Miller for help measuring end-to-end latencies with his latency meter [MB02].

References

[ACW+99] D. Aliaga, J. Cohen, A. Wilson, H. Zhang, C. Erikson, K. Hoff, T. Hudson, W. Stuerzlinger, E. Baker, R. Bastos, M. Whitton, F. Brooks, and D. Manocha. Mmr: An integrated massive model rendering system using geometric and image-based acceleration. In *Proc. of ACM Symposium on Interactive 3D Graphics*, 1999.

[ARB90] J. Airey, J. Rohlf, and F. Brooks. Towards image realism with interactive update rates in complex virtual building environments. In *Symposium on Interactive 3D Graphics*, pages 41–50, 1990.

[ASVNB00] C. Andujar, C. Saona-Vazquez, I. Navazo, and P. Brunet. Integrating occlusion culling and levels of detail through hardly-visibly sets. In *Proceedings of Eurographics*, 2000.

[AT99] J. Alex and S. Teller. Immediate-mode ray-casting. Technical report, MIT LCS Technical Report 784, 1999.

[BMH99] D. Bartz, M. Meibner, and T. Huttner. Opengl assisted occlusion culling for large polygonal models. *Computer and Graphics*, 23(3):667–679, 1999.

[Clar76] Clark, J. Hierarchical Geometric Models for Visible Surface Algorithms. Communications of the ACM, 547-554, 1976.

[Cohe96] Cohen, J., Varshney, A., Manocha, D., Turk, G., Weber, H., Agarwal, P., Brooks, F., and Wright, W. Simplification Envelopes. In *Proc. of ACM SIGGRAPH*, 119-128, 1996.

[COCS01] D. Cohen-Or, Y. Chrysanthou, and C. Silva. A survey of visibility for walkthrough applications. *SIGGRAPH Course Notes # 30*, 2001.

[CT97] S. Coorg and S. Teller. Real-time occlusion culling for models with large occluders. In *Proc. of ACM Symposium on Interactive 3D Graphics*, 1997.

[DDTP00] F. Durand, G. Drettakis, J. Thollot, and C. Puech. Conservative visibility preprocessing using extended projections. *Proc. of ACM SIGGRAPH*, pages 239–248, 2000.

[EM99] C. Erikson and D. Manocha. Gaps: General and automatic polygon simplification. In *Proc. of ACM Symposium on Interactive 3D Graphics*, 1999.

[EMB01] C. Erikson, D. Manocha, and W. Baxter. HLODS for fast display of large static and dynmaic environments. *Proc. of ACM Symposium on Interactive 3D Graphics*, 2001.

[ESSS01] J. El-Sana, N. Sokolovsky, and C. Silva. Integrating occlusion culling with view-dependent rendering. *Proc. of IEEE Visualization*, 2001.

[ESV99] J. El-Sana and A. Varshney. Generalized view-dependent simplification. *Computer Graphics Forum*, pages C83–C94, 1999.

[EYAE99] S. Ellis, M. Young, B. Adelstein, and S. Ehrlich. Discrimination of changes of latency during voluntary hand movement of virtual objects. In *Proc. of the Human Factors and Ergonomics Society*, 1999.

[FH98] P. Felzenszwalb and D. Huttenlocher. Efficiently computing a good segmentation. In *Proceedings of IEEE CVPR*, pages 98–104, 1998.

[FKST96] T.A. Funkhouser, D. Khorramabadi, C.H. Sequin, and S. Teller. The ucb system for interactive visualization of large architectural models. *Presence*, 5(1):13–44, 1996.

[GBW90] B. Garlick, D. Baum, and J. Winget. Interactive Viewing of Large Geometric Databases Using Multiprocessor Graphics Workstations. In *SIGGRAPH '90 Course Notes (Parallel Algorithms and Architectures for 3D Image Generation), volume 28*, 1990.

[GH97] M. Garland and P. Heckbert. Surface simplification using quadric error bounds. *Proc. of ACM SIGGRAPH*, pages 209–216, 1997.

[GKM93] N. Greene, M. Kass, and G. Miller. Hierarchical z-buffer visibility. In *Proc. of ACM SIGGRAPH*, pages 231–238, 1993.

[Gre01] N. Greene. Occlusion culling with optimized hierarchical z-buffering. In *ACM SIGGRAPH COURSE NOTES ON VISIBILITY, # 30*, 2001.

[HEB+01] G. Humphreys, M. Eldridge, I. Buck, G. Stoll, M. Everett, and P. Hanrahan. WireGL: A scalable graphics system for clusters. *Proc. of ACM SIGGRAPH*, 2001.

[HMC+97] T. Hudson, D. Manocha, J. Cohen, M. Lin, K. Hoff, and H. Zhang. Accelerated occlusion culling using shadow frusta. In *Proc. of ACM Symposium on Computational Geometry*, pages 1–10, 1997.

[Hop96] H. Hoppe. Progressive meshes. In *Proc. of ACM SIGGRAPH*, pages 99–108, 1996.

[Hop97] H. Hoppe. View dependent refinement of progressive meshes. In *ACM SIGGRAPH Conference Proceedings*, pages 189–198. ACM SIGGRAPH, 1997.

[Hop98] H. Hoppe. Smooth view-dependent level-of-detail control and its application to terrain rendering. In *IEEE Visualization Conference Proceedings*, pages 35–42, 1998.

[Kru56] J.B. Kruskal On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of American Mathematical Society*, 7:48–50, 1956.

[KS00] J. Klowoski and C. Silva. The prioritized-layered projection algorithm for visible set estimation. *IEEE Trans. on Visualization and Computer Graphics*, 6(2):108–123, 2000.

[KS01] J. Klowoski and C. Silva. Efficient conservative visiblity culling using the prioritized-layered projection algorithm. *IEEE Trans. on Visualization and Computer Graphics*, 7(4):365–379, 2001.

[LE97] D. Luebke and C. Erikson. View-dependent simplification of arbitrary polygon environments. In *Proc. of ACM SIGGRAPH*, 1997.

[Lue01] D. Luebke. A developer's survey of polygon simplification algorithms. *IEEE CG & A*, pages 24–35, May 2001.

[MB02] D. Miller and G. Bishop. Latency meter: a device for easily monitoring VE delay. In *Proceedings of SPIE*, Vol. #4660 Stereoscopic Displays and Virtual Reality Systems IX, San Jose, CA, January 2002.

[RB93] J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering. In *Modeling in Computer Graphics*, pages 455–465. Springer-Verlag, June–July 1993.

[RH94] J. Rohlf and J. Helman. Iris Performer: A high performance multiprocessor toolkit for realtime 3D Graphics In *Proc. of ACM SIGGRAPH*, 381–394, 1994.

[RL00] S. Rusinkiewicz and M. Levoy. QQplat: A multiresolution point rendering system for large meshes. *Proc. of ACM SIGGRAPH*, 2000.

[SBM+94] B. Schneider, P. Borrel, J. Menon, J. Mittleman, and J. Rossignac. BRUSH as a walkthrough system for architectural models. In *Fifth Eurographics Workshop on Rendering*, pages 389–399, July 1994.

[Sch97] W. Schroeder. A topology modifying progressive decimation algorithm. In *Proceedings of Visualization'97*, pages 205–212, 1997.

[SDDS00] G. Schaufler, J. Dorsey, X. Decoret, and F. Sillion. Conservative visibility preprocessing using extended projections. *Proc. of ACM SIGGRAPH*, pages 229–238, 2000.

[SFLS00] R. Samanta, T. Funkhouser, K. Li, and J. P. Singh. Hybrid sort-first and sort-last parallel rendering with a cluster of pcs. *Eurographics/SIGGRAPH workshop on Graphics Hardware*, pages 99–108, 2000.

[Tel92] S. J. Teller. *Visibility Computations in Densely Occluded Polyheral Environments*. PhD thesis, CS Division, UC Berkeley, 1992.

[WSB01] I. Wald, P. Slusallek, and C. Benthin. Interactive distributed ray-tracing of highly complex models. In *Rendering Techniques*, pages 274–285, 2001.

[WWS00] P. Wonka, M. Wimmer, and D. Schmalstieg. Visibility preprocessing with occluder fusion for urban walkthroughs. In *Rendering Techniques*, pages 71–82, 2000.

[WWS01] P. Wonka, M. Wimmer, and F. Sillion. Instant visibility. In *Proc. of Eurographics*, 2001.

[XESV97] J. Xia, J. El-Sana, and A. Varshney. Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):171–183, June 1997.

[ZMHH97] H. Zhang, D. Manocha, T. Hudson, and K. Hoff. Visibility culling using hierarchical occlusion maps. *Proc. of ACM SIGGRAPH'97*, 1997.