THE JOURNAL OF VISUALIZATION AND COMPUTER ANIMATION
J. Visual. Comput. Animat. 2001; **12:** 277–286
Published online in Wiley InterScience (www.interscience.wiley.com). **(DOI:** 10.1002 / vis.269)

Exploiting coherence in hierarchical visibility algorithms

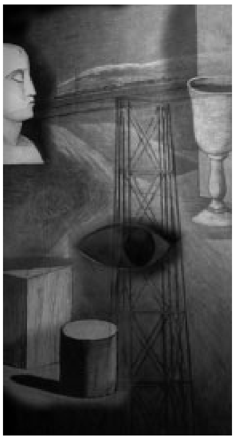

By Jiří Bittner and Vlastimil Havran*

We present a series of simple improvements that make use of temporal and spatial coherence in the scope of hierarchical visibility algorithms. The* hierarchy updating *avoids visibility tests of certain interior nodes of the hierarchy. The* visibility propagation *algorithm reuses information about visibility of neighbouring spatial regions. Finally, the* conservative hierarchy updating *avoids visibility tests of the hierarchy nodes that are expected to remain visible. We evaluate the presented methods in the context of hierarchical visibility culling using* occlusion trees. *Copyright © 2002 John Wiley & Sons, Ltd.

Received: 30 November 2001; Revised: 12 December 2001

KEY WORDS: visibility culling; coherence; 4D-tree

Introduction

Exploiting various types of coherence during image synthesis is one of the main goals of modern computer graphics. In the scope of visibility algorithms at least three types of coherence can be used: object space, image space, and temporal. Hierarchical visibility algorithms make some use of the spatial coherence inherently by utilizing a spatial hierarchy. For densely occluded scenes they may achieve a great benefit by quickly identifying groups of invisible objects that need not be considered for rendering.

A typical hierarchical visibility algorithm uses a *visibility test*, which classifies a node of the spatial hierarchy as completely visible, partially visible or invisible, depending on the visibility of the spatial region corresponding to that node. The visibility test is applied recursively starting at the root node. As soon as a node is found completely visible or invisible, the current branch of the traversal can be terminated, since visibility of all nodes in the current subtree is imposed by the visibility of the current node. In this paper we do not focus on the amount of image space or temporal coherence, which may be exploited by the visibility test itself. Instead we suggest a more general framework that is rather independent of the particular visibility algorithm.

*Correspondence to: J. Bittner, Center for Applied Cybernetics, Czech Technical University in Prague, Karlovo náměstí 13, 121 35 Praha 2, Czech Republic. E-mail: bittner@fel.cvut.cz

Traditional hierarchical visibility algorithms traverse the spatial hierarchy starting at the root node. Firstly, we propose a method that saves up to half of the visibility tests by skipping certain interior nodes of the hierarchy (assuming the spatial hierarchy corresponds to a binary tree). The skipping is guided by visibility classifications obtained during the previous invocation of the visibility algorithm. Secondly, we describe an algorithm that increases the amount of spatial coherence exploited. It reuses visibility classifications of hierarchy nodes already processed in the current pass of the algorithm. The nodes are processed in front-to-back order and the algorithm tries to determine visibility of the region corresponding to the current node by combining visibility states of neighbouring regions. If it fails, the usual visibility test is applied. Finally, we propose a conservative method that aims to avoid repeated visibility tests of nodes that probably remain visible.

Related Work

Some visibility algorithms exploit temporal coherence in a specialized way that reflects the principles of each such algorithm. Greene *et al.*[1] use the set of visible objects from one frame to initialize the *z-pyramid* in the next frame and so reduces 'overdraw' of the *hierarchical z-buffer*. Coorg and Teller[2] present a visibility algorithm that uses *relevant planes* which form a subset of visual events. They restrict the hierarchy traversal to nodes

Based on 'Exploiting Temporal and Spatial Coherence in Hierarchical Visibility Algorithms' by J. Bittner and V. Havran which appeared in *Proceedings of the Seventeenth Spring Conference on Computer Graphics*, April 2001, Budmerice, Slovakia, edited by Roman Durikovic and Silvester Czanner. © 2001 IEEE.

corresponding to planes that were crossed between successive viewpoint positions. Another method of Coorg and Teller[3] exploits temporal coherence by caching occlusion relationships.

Slater and Chrysanthou[4] have proposed a probabilistic scheme for view-frustum culling. They partition objects into groups, which are sampled according to their distance from the view-frustum. It is difficult to generalize this method for visibility algorithms, since the 'visible volumes' can be very complex, and usually they are not explicitly reconstructed. Moreover, this method is not conservative unless changes in viewing direction and position of the viewpoint are restricted.

The methods proposed in this paper can make use of temporal and spatial coherence in the scope of existing visibility algorithms that utilize a spatial hierarchy. Examples of these are algorithms based on hierarchical occlusion maps,[5] coverage masks,[6] shadow frusta[7] and occlusion trees.[8]

Overview

In order to describe modifications of the visibility algorithm we first restrict our discussion to one particular approach: *conservative hierarchical visibility culling*. Below, we give a short overview of the data structures and algorithms that are used in the scope of the proposed methods.

Spatial Hierarchy

Hierarchical visibility culling utilizes a spatial hierarchy that is built over all objects of the scene. We have focused on kD-trees[9] because of their high flexibility and simplicity of building and traversal. A node of the kD-tree corresponds to an axis-aligned bounding box. Each leaf of the tree contains a list of references to objects that intersect the corresponding box.

The Node Visibility Test

The elementary step of the hierarchical visibility culling is the *node visibility test*, i.e., visibility classification of a single node of the hierarchy using a certain occlusion map. We assume that given a viewpoint and a viewing direction the visibility algorithm classifies visibility of the node as *completely visible*, *partially visible* or *invisible*. Although in this paper we do not focus on the visibility determination step itself, we give a brief description of one such algorithm (see Bittner *et al*.[8] for further details).

For each position of the viewpoint several large polygonal occluders are identified. These are used to build an *occlusion tree*, which results from merging 'shadow' frusta of each individual occluder. Briefly, the occlusion tree is a Binary Space Partitioning (BSP) tree,[10] which has its leaves classified as *in* or *out*, if they are occluded or unoccluded, respectively. The node visibility test is performed using constrained depth first search (DFS) on the occlusion tree. The final visibility classification is obtained by hierarchical combination of visibility states of nodes reached by the DFS.

Classical Approach

Classical hierarchical visibility culling proceeds as follows: starting from the root node of the hierarchy, view-frustum culling is applied on the current node.[11] If the node is outside the view-frustum it is classified as invisible. Otherwise, the node visibility test is performed. If the node is found to be visible all its descendants are visible. Similarly, if a node is invisible all its children are invisible. Descendants of nodes classified as partially visible are tested further to refine their visibility (see Figure 1). When the visibility of all leaves is known, objects from fully visible and partially visible leaves can be gathered and rendered using a low-level exact visibility solver (such as depth-buffer).

A simple improvement can be used to avoid visibility tests of hierarchy nodes that contain only few objects and so the estimated cost of rendering the objects is lower than the cost of the visibility determination. In such a case the node can be simply classified as visible.

Modifications Overview

In order to give an overview of the proposed modifications we first show how they are exploited in the scope of the hierarchical visibility algorithm (see Figure 2). The *hierarchy updating* test is applied first. This test eventually decides to skip all the remaining steps and to continue determining visibility of descendants of the current node. The *view-frustum* culling can report the node as invisible if it is outside the view-frustum. Otherwise, *visibility propagation* is applied, which can succeed in classifying the node as visible or invisible. The *conservative hierarchy updating*

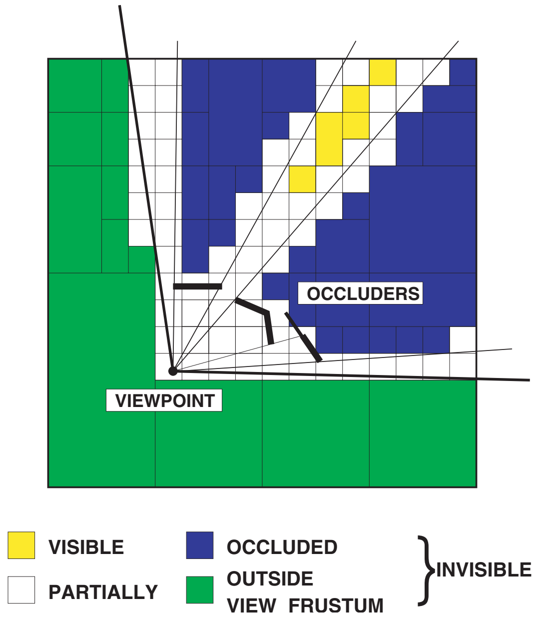

Figure 1. An example of hierarchical visibility culling. The node visibility test uses merged occlusion volumes of four occluders.

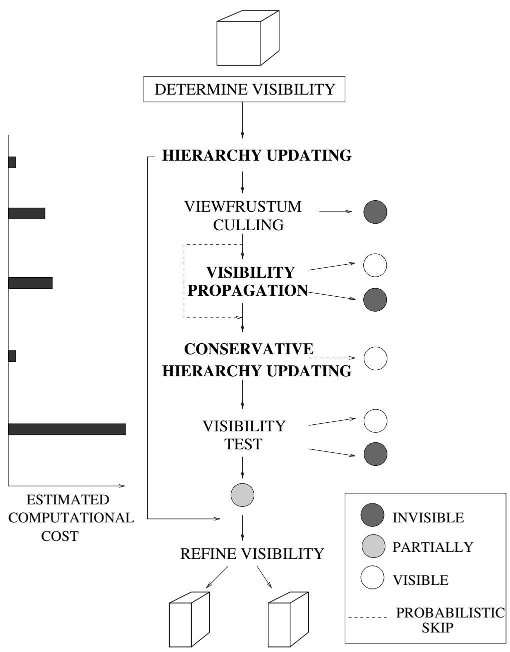

Figure 2. Series of steps determining visibility of a node of the hierarchy. The novel methods are highlighted.

classifies some nodes as visible with certain probability. If all previous steps failed in determining a node's visibility, the node visibility test is applied. Note that the steps are applied in order of increasing computational cost, reflecting the main idea of culling: use a more complicated test only when the simple one fails to find a solution.

The rest of the paper is organized as follows: In the next section the hierarchy updating method is introduced. The sixth section describes a conservative modification of hierarchy updating. Visibility propagation is presented in the seventh section. Results and comparisons are presented in the eighth section. Finally, the ninth section discusses some topics for future work and concludes.

Hierarchy Updating

The hierarchical visibility algorithm can be seen as a traversal of the hierarchy, which is terminated either at leaves or nodes classified either as visible or invisible. Let us call such nodes the *termination nodes* and nodes that have been classified as partially visible the *opened nodes*. Denote sets of termination and opened nodes in the i-th frame $\mathcal{T}_i$ and $\mathcal{O}_i$, respectively. In the classical approach $\mathcal{T}_i \cup \mathcal{O}_i = \mathcal{V}_i$, where $\mathcal{V}_i$ is the set of all nodes visited in the i-th rendering frame.

Imagine the viewpoint is fixed. Visibility of all nodes of the hierarchy does not change and the sets $\mathcal{T}_i$, $\mathcal{O}_i$, and $\mathcal{V}_i$, are fixed as well. Nevertheless, the classical algorithm repeatedly tests visibility of all nodes $\mathcal{V}_i$.

Hierarchy updating is a modification that aims to eliminate the repeated visibility tests of the set of opened nodes from the previous frame. It skips all nodes of $\mathcal{O}_{i-1}$ and applies node visibility tests only on nodes of $\mathcal{T}_{i-1}$. In order to propagate eventual changes in visibility up into the hierarchy the visibility states determined at the termination nodes are pulled up according to the following rule: the visibility state of the node is updated as visible or invisible, if all its children have been classified as visible or invisible, respectively. Otherwise, it remains partially visible and thus opened. The pseudo-code of the hierarchical visibility algorithm with hierarchy updating is outlined in Figure 4. Note that the set of termination nodes is not maintained explicitly. Instead, each node contains its previous visibility classification. The *frame* variable is associated with each node that is used to identify nodes below the current termination nodes.

Visualization & Computer Animation

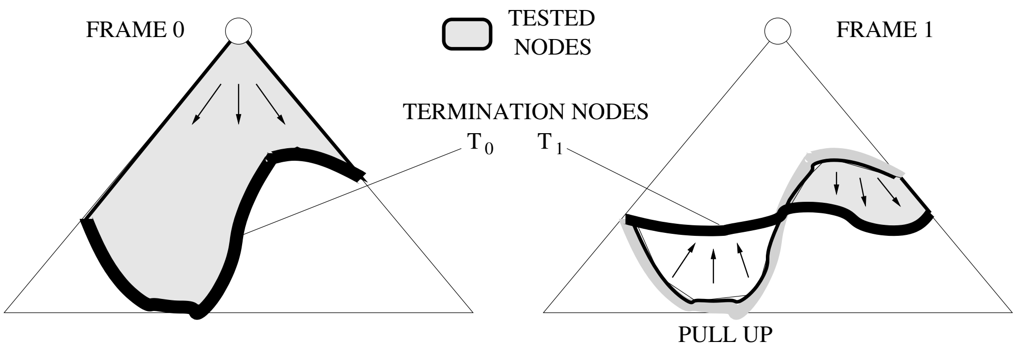

Figure 3. Illustration of hierarchy updating. Initially the algorithm proceeds starting at the root of the hierarchy (left). In the second frame the opened nodes $\mathcal{O}_0$ are skipped and the visibility tests are applied on the termination nodes $\mathcal{T}_0$ (and eventually 'below'). Visibility changes are propagated up to the hierarchy and the new set of termination nodes $\mathcal{T}_1$ is established.

Consequently, the modification does not change the final visibility classification, which is the same as that obtained using the classical approach. The behaviour of the modified hierarchical visibility algorithm is illustrated in Figure 3. Note that if the pull-up did not take place the algorithm could end up with the termination nodes being all leaves of the hierarchy. Hence, it would lose the advantages of the hierarchical algorithm.

```
Algorithm Hierarchical Visibility (NODE)
begin
  if NODE is leaf or NODE.visibility ≠ Partially
    (* termination nodes*)
    or NODE.frame < frame-1 then
  begin
    NODE.visibility ⟵ Test Visibility (NODE);
    NODE.frame ⟵ frame;
  end
  case NODE.visibility of
   VISIBLE : Render subtree of NODE;
   PARTIALLY :
     if NODE is leaf then Render NODE;
     else
       for all children C of NODE do
         Hierarchical Visibility (C);
     (* pull-up *)
     if visibility of all children equals v then
       begin
         NODE.visibility ⟵ v;
         NODE.frame ⟵ frame;
       end
   INVISIBLE : (* terminate the DFS *)
  end
end
```

Figure 4. Pseudo-code of the hierarchical visibility culling with hierarchy updating.

For kD-trees $|\mathcal{O}_i| = |\mathcal{T}_i| - 1$. Thus hierarchy updating can save almost a half of the visibility tests that would be applied on the interior nodes of the hierarchy.

Conservative Hierarchy Updating

The hierarchy updating method ensures that on each path to a leaf node of the hierarchy at least one node is tested for visibility. We can further reduce the expected number of node visibility tests at the cost of the conservative behaviour of the modified algorithm. The conservative hierarchy updating produces a superset of visible nodes determined by hierarchy updating alone.

Due to the complexity of the occlusion volume it is difficult to predict changes in visibility unless a specialized visibility algorithm is involved.[2] To keep the conservative behaviour of the algorithm we cannot classify a node as invisible without really determining its visibility. Nevertheless, assuming visibility does not change significantly over successive frames, visibility states of visible and partially visible nodes do not have to be updated in each frame. We use a simple probabilistic sampling scheme. Visibility of a termination node that was classified as visible or partially visible in the last frame is updated with probability $1 - p_{skip}$.

With probability p_{skip} the node visibility test is skipped and the node is classified as visible. This method reduces the number of visibility tests applied on visible nodes of the hierarchy, but it does not immediatelly capture all changes in visibility. In such cases more nodes are classified as visible and consequently more objects are rendered comparedwith the non-conservative hierarchy updating. Nevertheless we show in the 'Results and Discussion' section that for the tested scenes and corresponding walkthrough paths we could determine such a p_{skip} that the total frame time was minimized.

Visibility Propagation

Hierarchical visibility culling already makes use of spatial coherence by utilizing a spatial hierarchy (kD-tree). However, we can further increase the amount of coherence exploited by reusing visibility information computed for neighbouring regions.

Suppose that the nodes of the spatial hierarchy are processed in front-to-back order with respect to the viewpoint. Using kD-tree this ordering is determined in a simple way.[10] First, we try to determine the visibility of the currently processed node by combining visibility classifications of its relevant neighbours. If the combination fails we revert to the node visibility test.

Let us denote the box corresponding to node N as B_N. The visibility of N can be determined combining the visibility of potentially visible faces $\mathcal{F}_{B_N}$ of B_N ($|\mathcal{F}_{B_N}| \leq 3$). Consequently, the visibility of a face $F \in \mathcal{F}_{B_N}$ can be determined combining visibility of appropriate *neighbour nodes*. If all faces of $\mathcal{F}_{B_N}$ are invisible the node N is invisible. Similarly, if all faces of $\mathcal{F}_{B_N}$ are visible and there is no occluder intersecting B_N, N can be classified as completely visible. Otherwise, the visibility propagation fails and the usual node visibility test must be applied. An example of a node that can be classified as invisible is depicted in Figure 5.

A neighbour node of N on a face F is a node U of the kD-tree with B_U lying in the opposite halfspace (induced by F) to B_N and having non-empty intersection with F. Naturally, we could keep a list of neighbour nodes for each face. Instead, we have used neighbour links (ropes) for kD-trees[12] that have low memory requirements and allow hierarchical visibility propagation.

Within each face F we associate a link to a neighbour node U that has a smallest box containing the face completely ($F \cap B_U = F$). When determining visibility of a face F there are three possible cases:

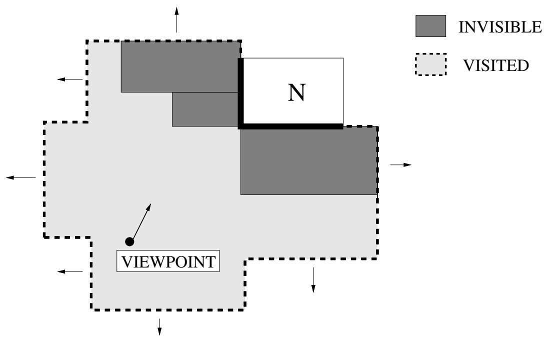

Figure 5. Node N can be classified as invisible since all its appropriate neighbours are invisible.

1. The link points to a node that is visible/invisible,
2. The link points to a node that is partially visible,
3. The link points to a node that has not been visited in the current frame.

The first case is trivial; the visibility of the face can be set immediately. In the second case we perform a constrained DFS and combine visibility of reached nodes. The search is constrained to nodes having non-empty intersection with the face F and terminates at the termination nodes $\mathcal{T}_i$. This process is illustrated in Figure 6. The visibility combination is performed using the same rule as in the pull-up pass of hierarchy updating (section on 'Hierarchy Updating'). Nevertheless, we can terminate the DFS whenever the combination results in partial visibility.

The third case is solved as follows: if the link is pointing to a node that has not been visited in the current frame, there must be some termination node on

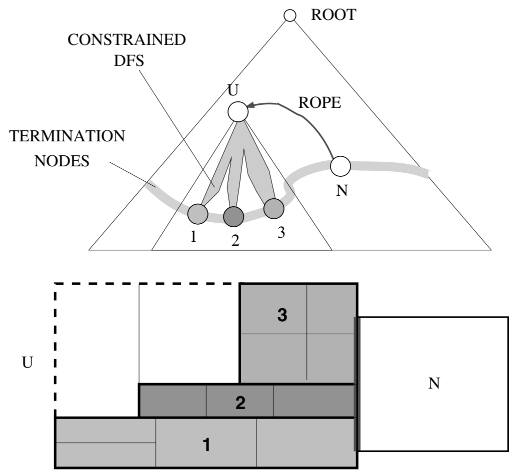

Figure 6. Hierarchical visibility propagation using ropes.

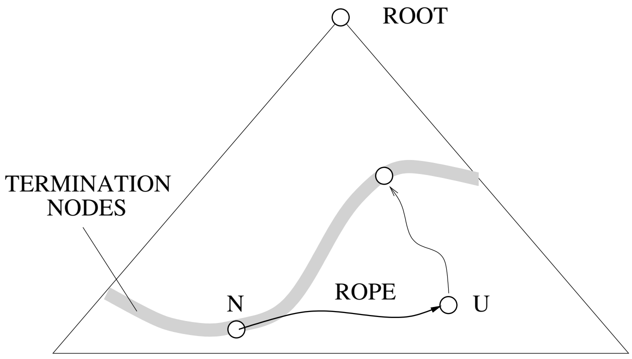

Figure 7. Lazy propagation of the visibility classification.

the path to the root. This path is followed until the termination node is reached (see Figure 7). Note that if the visibility states were propagated into subtrees of the termination nodes, the third case would never occur.

Temporal Coherence

The visibility propagation does not always succeed in determining the visibility of the processed node. In such a case it introduces an additional overhead into the visibility determination. However, we can use information obtained in the previous frame to guide the algorithm in the current frame.

Firstly, we avoid visibility propagation on nodes that we expect to remain partially visible and thus the visibility propagation would probably fail. To achieve this we apply the visibility propagation only on nodes that have not been classified as partially visible in the previous frame. Secondly, if for a given node the visibility propagation succeeded in the previous frame, it is applied in the current frame as well. Otherwise, it is applied with certain probability $p_{vp} < 1$.

Results and Discussion

We evaluated the proposed algorithms on two test scenes. The first scene (scene I) is a model of the fifth floor of the Soda Hall;* the second scene (scene II) is a building interior with precomputed lighting. The measurements were conducted using SGI O2 with 128 MB memory. In all measurements we used the visibility culling algorithm based on occlusion trees.[8]

The following methods were evaluated:

A—the classical approach;
B—hierarchy updating applied;
C—hierarchy updating + visibility propagation with probability $p_{vp} = 0.5$;
D—as **C** + conservative hierarchy updating with probability $p_{skip} = 0.5$.

For scene I the constructed kD-tree consisted of 1187 nodes, and of 1605 nodes for scene II. For each position of the viewpoint 16 occluders were identified and used to build the occlusion tree during walkthrough of scene I. For scene II we used 32 occluders, since the scene contained smaller patches resulting from the radiosity precomputed lighting. In both scenes a predetermined walkthrough path was followed for each measurement (see Figures 12 and 13 for scene snapshots). If not stated differently all presented values are averaged per one frame of the walkthrough.

The first six plots illustrate the dependence of the algorithms on the relative speed of the walk (Figures 8–10). A unit relative speed roughly corresponds to normal walking speed. We have measured the number of node visibility tests, the time spent by the hierarchical visibility determination and the total frame time.

All evaluated methods exhibit a very slow growth of the number of necessary node visibility tests. For a walk of relative speed 1.0 the following savings in average number of node visibility tests were achieved (compared to **A**):

- **scene I**—method **B** 47%, method **C** 50% and method **D** 67%;
- **scene II**—method **B** 49%, method **C** 51% and method **D** 72%.

Hierarchy updating (method **B**) saves almost half of

*http://graphics.lcs.mit.edu/ ~becca/research/SodaHall

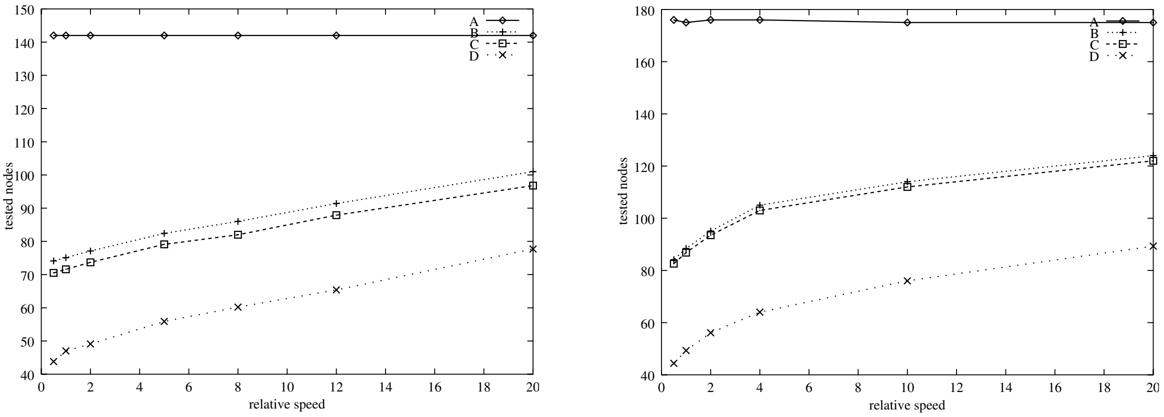

(a) (b)

Figure 8. Dependence of the number of node visibility tests on the relative speed of the walk for scene I (a) and scene II (b).

the node visibility tests, as expected. We have observed that visibility propagation (method **C**) succeeds in determining visibility of only a few nodes that usually correspond to rather large regions. The **D** method significantly decreases the number of node visibility tests. This is paid by a higher number of nodes classified as partially visible or visible (details follow further in the text).

Figure 9 shows that the time spent by hierarchical visibility culling was roughly proportional to the number of node visibility tests. Nevertheless, we can observe that the time spent by visibility propagation (method **C**) is not recovered by the savings in number of node visibility tests. In particular, this follows from the fact that the node visibility test using the occlusion tree is almost as fast as the visibility propagation.

In Figure 10 we can observe the conservative behaviour of method **D**. When the viewpoint moves slowly, this method achieves better frame times than the others. As the relative speed of the walk increases the visibility states of many nodes change quickly. Hence 'reusing' some previously visible nodes leads to a larger set of nodes to render and the frame time is increased.

Finally, we measured the behaviour of the conservative hierarchy updating algorithm in dependence on the probability p_{skip} (Figure 11). We can observe local minima in the average frame time at $p_{skip} = 0.5$ for scene I and $p_{skip} = 0.6$ for scene II. For probabilities greater than this minimumsavings in visibility classification do not recover the time necessary for rendering otherwise invisible objects.

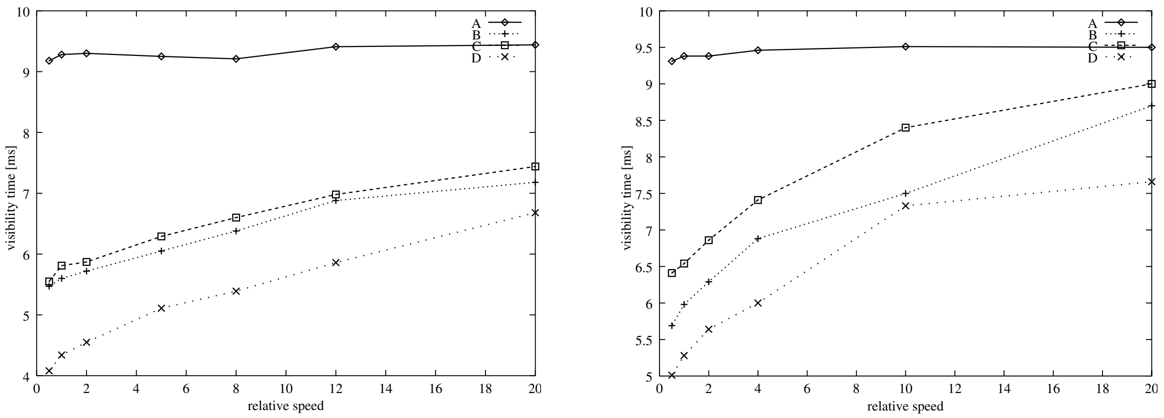

(a) (b)

Figure 9. Average time spent by the hierarchical visibility algorithm for scene I (a) and scene II (b).

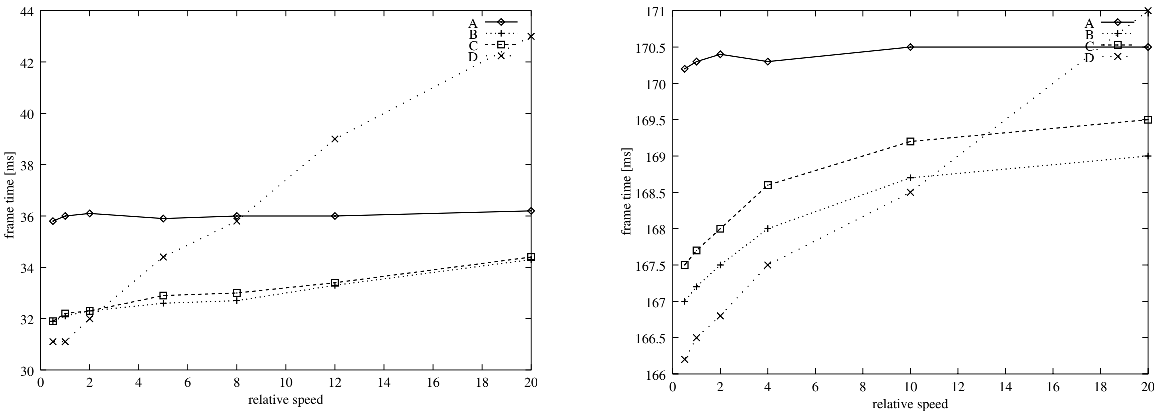

(a) (b)

Figure 10. Average frame time in dependence on the relative speed of the walk for scene I (a) and scene II (b).

It is worth mentioning that our aim was not to evaluate the visibility algorithm itself, but rather to document the impact of the proposed methods. It is obvious that if the visibility algorithm was more demanding, the proposed methods would decrease the total frame time more significantly.

Conclusion and Future Work

In this paper we have introduced a series of modifications of the classical hierarchical visibility culling. The hierarchy updating proved to perform well in practice as it saves almost half of the visibility tests that would have to be applied using the classical approach. The savings would be less remarkable for hierarchies with higher branching factors, but our preliminary results indicate that kD-trees with arbitrarily positioned partitioning planes are much more effective for visibility culling than octrees or bounding volume hierarchies.

Suprisingly, we have observed that visibility propagation saves only a few visibility tests. This documents that the spatial coherence is already well exploited in the classical approach. Finally, we have shown that conservative hierarchy updating can improve the overall frame time for certain settings.

We have experimented with fixed probabilities used in both conservative hierarchy updating and the

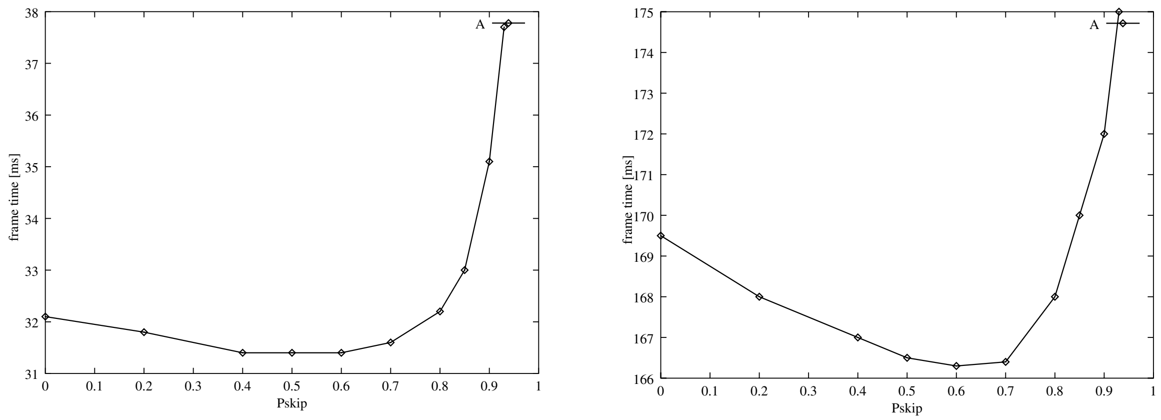

(a) (b)

Figure 11. Dependence of the average frame time on the probability p_{skip} using conservative hierarchy updating for scene I (a) and scene II (b).

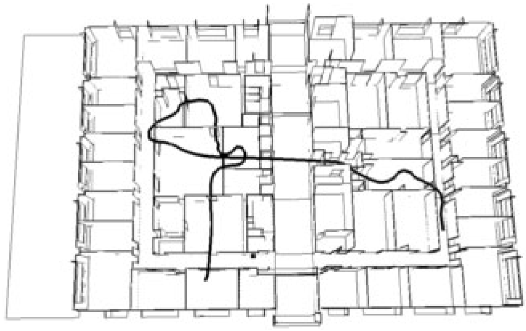

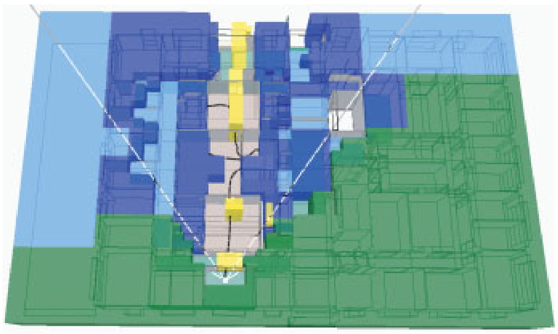

Figure 12. (Left) The path used for a walk through the model of scene I. For relative speed of walk equal to 1.0 the walk consists of 980 steps. (Right) An example of hierarchical visibility culling. The green regions are outside of the view-frustum. The few yellow regions in the viewing direction are completely visible. Invisible regions are shown in dark blue. The light-blue regions were found to be invisible by the visibility propagation algorithm. Partially visible regions are transparent.

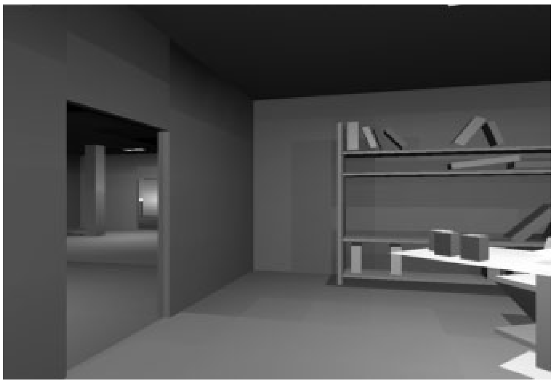

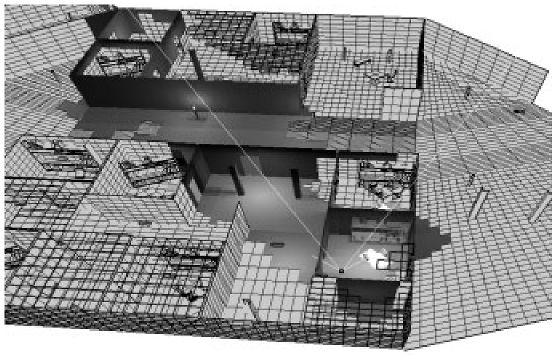

Figure 13. (Left) A camera view on the test path through scene II. (Right) Top view showing the part of the scene classified as invisible (dark gray).

probabilistic modification of the visibility propagation algorithm. More elaborate methods could be used that automatically adjust these probabilities according to the history of visibility changes.

The algorithm as described determines whether a node is visible or invisible. It could be extended to estimate whether a node that was partially visible in the previous frame remains partially visible. This modification would be of benefit in sparsely occluded environments where many small regions are classified as partially visible.

ACKNOWLEDGEMENTS

This research was supported by the Czech Ministry of Education under Project LN00B096 and the Aktion Kontakt OE/CZ, grant number 1999/17.

References

1. Greene N, Kass M, Miller G. Hierarchical Z-buffer visibility. In *Computer Graphics (Proceedings of SIGGRAPH '93)*, 1993; pp 231–238.
2. Coorg S, Teller S. Temporally coherent conservative visibility. In *Proceedings of the Twelfth Annual ACM Symposium on Computational Geometry*, Philadelphia, PA, May 1996.
3. Coorg S, Teller S. Real-time occlusion culling for models with large occluders. In *Proceedings of the Symposium on Interactive 3D Graphics*, New York, 27–30 April 1997. ACM Press; 83–90.
4. Slater M, Chrysanthou Y. View volume culling using a probabilistic caching scheme. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST-97)*, New York, 15–17 September 1997. ACM Press; pp 71–78.
5. Zhang H, Manocha D, Hudson T, Hoff III KE. Visibility culling using hierarchical occlusion maps. In *Proceedings of SIGGRAPH 97*, August 1997; pp 77–88.

6. Greene N. Hierarchical polygon tiling with coverage masks. In *Proceedings of SIGGRAPH '96*, August 1996; pp 65–74.
7. Hudson T, Manocha D, Cohen J, Lin M, Hoff K, Zhang H. Accelerated occlusion culling using shadow frusta. In *Proceedings of the Thirteenth ACM Symposium on Computational Geometry*, Nice, France, June 1997.
8. Bittner J, Havran V, Slavík P. Hierarchical visibility culling with occlusion trees. In *Proceedings of Computer Graphics International '98 (CGI'98)*, IEEE, 1998; pp 207–219.
9. Samet HJ. *Design and Analysis of Spatial Data Structures: Quadtrees, Octrees, and Other Hierarchical Methods*. Addison-Wesley, Reading, MA, 1989.
10. Fuchs H, Kedem ZM, Naylor BF. On visible surface generation by a priori tree structures. In *Proceedings of SIGGRAPH '80*, July 1980; pp 124–133.
11. Rohlf J, Helman J. IRIS performer: a high performance multiprocessing toolkit for real-time 3D graphics. In *Proceedings of SIGGRAPH '94*, July 1994; pp 381–395.
12. Havran V, Bittner J, Žára J. Ray tracing with rope trees. In *Proceedings of 13th Spring Conference on Computer Graphics*, Budmerice, 1998; pp 130–139.

Authors' biographies

Jiří Bittner currently works as a researcher at the Center for Applied Cybernetics in Prague, Czech Republic. He is a PhD candidate at the Department of Computer Science of the Czech Technical University, where he received his MS in 1997. His research interests include efficient rendering techniques, global illumination and computational geometry.

Vlastimil Havran received an MS in 1996 for computer science and engineering and PhD in 2001 for computer science, both at the Czech Technical University in Prague, Czech Republic. He was working partly in industry in 1999–2001 at IGP Company, Czech Republic. He currently holds a post-doctoral research position at Max-Planck-Institut, für Informatik, saarbrucken, Germany. Dr Havran's current research interests reside in efficient algorithms for computer graphics, particulary algorithms in visibility computation and global illumination.