

Real-Time Rendering of Complex Photorealistic Landscapes Using Hybrid Level-of-Detail Approaches

Carsten COLDITZ, Liviu COCONU, Oliver DEUSSEN and Christian HEGE

1 Introduction

It still remains a challenge to realistically render outdoor scenes with interactive frame-rates since these scenes consist of an immense number of highly complex geometric models. A square kilometre of a forest incorporates millions of plants, hundreds of thousands of trees, and numerous small shrubs. Even the explicit storage of the plant positions is extremely challenging and must therefore be avoided. The vast amount of geometry can only be handled using efficient level-of-detail algorithms that approximate larger parts of the scenes by very simple elements.

In this paper procedures are demonstrated by representing models and parts of the scene by dynamically changing sets of billboards. After a survey of existing literature in the following paragraphs, we first present algorithms to optimally approximate plant models using billboards. Sect. 3 introduces methods to realistically illuminate these billboards, which is important for a realistic rendition of the scene and for a smooth transition from billboard approximation to the detailed geometry. The geometry is still needed for close-ups of the plants since billboard clouds do not completely avoid parallax errors and cannot represent the details of a geometry-based model. To allow the approximation of larger parts of the scene, square tiles in the background consisting of many plant models are represented by combined groups of billboards that can be displayed with a single rendercall per tile and not with a rendercall per position. Following the presentation and discussion of the results, we will outline our future works.

1.1 Related Work

Related work can be divided into three sub-areas: level-of-detail algorithms, realistic rendering of complex objects, and billboard techniques.

1.1.1 Level-of-Detail

There has been a number of approaches to speed up the rendering of complex models. Automatic methods for reducing the geometric complexity of surfaces by triangle decimation have been developed by Hoppe [Hop]. Eck et al. [EDD*95] describe automatic polygonal reduction based on wavelet analysis. These methods can efficiently be applied to smooth objects. For trees, however, they are not well suited since the basic elements often consist of single triangles that cannot be simplified. An earlier work that explicitly dealt with this problem was presented by Rossignac and Borrel [RB93]. They proposed a method applicable to such data that determines groups of nearby vertices, and collapses them to a single vertex without being concerned with the topology. The method works, but the appearance of the models change since normals are altered. As early as 1985, Reeves and

Blau [RB85] represented trees using collections of disks in order to reduce the complexity of the foliage. Carefully selecting colors and shadowing effects resulted in very aesthetic images though the biological correctness of the models was very poor. Weber and Penn [WP] proposed a method of approximating trees by points and lines. The foliage was represented by a set of points that were placed inside agglomerated leaves of the foliage, which they referred to as “masses”. Lines were used for the tree skeleton. The idea of using point clouds for representing trees and other objects was further developed by several authors [SD01, DCSD02, DVS03]. A similar idea was presented by Wand et al. [WFP*]. For each triangle of a scene it was determined whether it was shown polygonally or point wise. Very efficient computing methods allowed displaying scenes of extreme complexity. Shade et al. [SGHS] used layered depth images (LDI) to render complex objects from pre-computed pixel-based representations with depth at places where full polygonal representations are too expensive and impostors such as billboards are too inaccurate. Chang et al. [CBL99] developed a hierarchical form of LDI that allowed them to generate different levels of detail. Both approaches are pixel-based, which results in low image quality for close-ups. Max [Max96] modeled and rendered trees hierarchically from pre-computed images with Z-buffers. A twig consists of some images of leaves, a branch of several images of twigs. Viewed from a distance, the whole tree is represented by one single image, if the viewer comes closer, the image is replaced by a combination of the first order branch images. A finer representation is achieved by replacing the images of first order branches with several twig images and so on. Though Max et al. [MDK99] accelerate the process by using texture hardware, the rendering times are far above of what is needed for interactive applications due to extensive texture transfer operations.

1.1.2 Billboard-Representations

In contrast to these techniques, billboards have been used for the representation of tree models for many years, see, e.g., [RH94]. In its simplest form, a billboard represents a whole model and therefore it can only be used in the background since no parallax distortion can be handled. A better approximation can be achieved using volumetric representations. Jakulin uses sets of parallel billboards [Jak00] to represent trees, however, due to the parallel orientation artifacts occur. Neyret [Ney98] converts complex natural objects into volumetric textures, which are then raytraced. Mip-mapping is used to reduce texture data and to generate samples from distant trees. The method provides high quality images of very complex scenes at a moderate rendering time of several minutes, but it is not applicable to interactive purposes. Recently this technique was adapted to graphics hardware by the authors [DN04], however, an efficient level-of-detail mechanism on the billboards is needed to cope with the enormous amount of textures that are needed to represent large scenes. We use a similar technique to render larger parts of our scenes in the background and present a method for achieving level-of-detail on these volumetric textures. To represent tree models using billboards, the "billboard clouds" approach can be used [DDSD03]. It represents a geometry by a set of arbitrarily oriented billboards. In this work an automatic method for finding a good set is presented, although for tree models the proposed algorithm is not an optimal solution, since the used plane-space transform is not able to provide significant results due to the noncompact geometry of our tree models. Instead, we present a clustering algorithm on the basis of the model vertices that allows us to find good billboard approximations.

1.1.3 Realistic Lighting

Realistic lighting is also still a challenge, since the models, and therefore also the light interaction, is quite complex. Even a single tree consists of thousands of leaves which scatter the light. Franzke and Deussen provide raytracing algorithms to render trees in a botanically correct way [FD03]. A quite accurate method for offline rendering of landscapes covered with trees was proposed by Qin et al. [QNTN03]. The plants are represented using a multi-layered data structure, which makes it possible to represent direct and indirect lighting effects. However, the models are in part a two-dimensional approximation, which prevents using them in a close distance to the virtual camera. In [COL04] we introduce some methods to approximate complex lighting effects like translucency under realtime rendering conditions.

Our system avoids most of the restrictions of former approaches, and is described in the following by first presenting how to optimally represent a plant model using a set of billboards.

2 Billboard Approximation of Trees

As mentioned in the introduction, algorithms can be applied to smooth geometry, but cannot automatically be applied to plant models. The "billboard clouds" approach by Decoret et al. [DDSD03] allows to find nearly optimal sets of billboards for smooth objects by transforming all triangles in the scene into the dual plane-space, where they are represented as points. Subsequently, point clusters are determined and represented by billboards. Unfortunately, applying this method with tree foliage results in a nearly even distributed point set where only weak clustering appears. We modified the idea and applied a clustering algorithm directly to the vertices of the triangles that constitute the tree model. Each of the clusters is then represented by one or more billboards with dependency on the form of the respective point cloud. In Fig. 1(a)-(d) a tree geometry is shown together with three billboard approximations. Using the standard k-means clustering [Fuk90], a sub-optimal set of billboards is achieved. Especially for the trunk and the tree skeleton, visual artefacts can be seen. This is due to the clustering that does not take into account any model-specific characteristics. Besides k-means, the isodata algorithm is often used for clustering [BH65]. In contrast to the k-means approach, clusters are dynamically split or merged due to their size. We modified this approach in order to get mostly long and thin clusters instead of the compact forms. These clusters can better be represented using billboards (see below). The best result can be achieved if information about the model is available. In our approach, we receive the triangles of the geometric description in a hierarchical form, i.e., all triangles that belong to one branch, including its sub-branches, are stored subsequently in a file. The texture information on the triangles is used to distinguish between branches and leaves, this allows us to use geometric descriptions of plants delivered by a number of modeling tools. The clusters can be oriented along the branches of the tree which yields visually pleasing results. In Fig. 1(a)-(d) a poor tree is shown so that all billboards can be seen. With dependence on the desired total number of billboards a branching level is determined which is completely covered by a set of billboards including all sub-branches. In Fig. 1(d) the second level was chosen, so all

branches are represented by, in this case, crossed billboards. In Fig. 2, two other models are shown. Also in these cases the algorithm finds good solutions for a representation even though the structural description of the models is different.




Fig. 1: Tree approximation: a) original model; b) 2 crossed plane billboards c) billboard approximation using k-means algorithm; d) approximation using improved clustering and hierarchical model information



Fig. 2: Other approximations of plants using the different billboard techniques

2.1 Visual Representation of Clusters

To determine the form of a cluster that should be represented by billboards, a bounding box is generated. Finding a good box is surprisingly not easy. In our experiments we first tried to work with a regression plane with all points having a minimal square distance. The plane then served as a basis for the bounding box. Unfortunately, the visual results are often not optimal. A better solution was provided by finding a bounding box with minimal volume. Analytic algorithms for determining such boxes exist, however they work quite inefficiently in the case of complex point clouds [BHP99]. So we developed a simple optimization procedure that jitters a given initial bounding box in order to find a better one.

By rotating the box in one-degree steps, we found good solutions in an efficient way. After obtaining a bounding box for each point cloud, depending on their form a representation is chosen. A nearly flat box with two long and a small side can be represented by a single billboard. If the points form a box with more equal side lengths, three crossed billboards are used. The billboard textures are created by rendering all geometry that is represented by the point cloud. Rendering of the billboard textures can be performed in two variants: in the simplest form pre-lighting is used and the billboards are created for a single light-position including all shadows. In this case, per billboard one texture is stored. Using special pixel and vertex shaders, one is able to perform normal mapping. Here the normals of the geometry elements are mapped onto the billboards, and then encoded in a second texture as color values. During hardware rendering the texture is obtained and the normals are used to shade the pixels of the billboards. This method allows us to render the same realistic lighting behaviour for the billboards as for the geometry models.

3 Realistic Lighting

As described in [COL04] for a realistic lighting impression at the rendering process of plant models, it must be possible to flip the direction of the normals dynamically. This is done inside a fragment-shader-program. There is no problem for the geometric representation of the model, because every vertex has its own normal and as a result, every fragment has an interpolated normal that can be updated. But standard billboards only have the color information stored inside their textures, and do not contain any normal information. Hence, if an observer is changing the point of view looking at a geometry based model, the normals can be flipped, and the lighting behaviour appears realistic. In contrast, standard billboards just store one single illumination state, and if the viewer changes the observer position, the impression of the illumination cannot be changed. To eliminate this drawback we have to save the normal information from the model and store it into a separate normal texture. Now we are able to update the direction of the normals at the billboards in the same way we already handled the normals of the geometry model.

In Fig. 3, the comparison between a geometry model of a tree, a billboard model of the tree without normal textures, and a billboard model with normal textures, illustrates the qualitative differences.




Fig. 3: Upper row: sideview of the tree (geometry model with normals, standard billboard model without normal textures, billboard model with normal textures)
Lower row: topview of the tree (same classification as above)

4 Level of Detail Rendering

To crossfade smoothly between the geometry model and the billboard representation of the model we use a dynamic transition. At the close-up range, only the model with full geometry complexity will be displayed. In the distance only the billboards will be rendered. Between these two static states the representation will be changed dynamically. The mode of operation is simple and guarantees good visual results without the need of complex blending operations. With increasing distance, the geometry of the original model will be reduced and in the inverse way, the value of the alpha channel test for the billboard rendering process will be changed. Consequently the billboard becomes more visible with every step. In Fig. 4 (a)-(f), the blending process is described. Fig. 4 (a) and (f) shows the start and the end state, at first, the whole complexity of the geometry model and at last the pure billboard representation. At the start state, the value of the alpha channel test for the billboard rendering part is 1.0. Then, it is decreased in tenth part steps to 0.5 at the end condition.

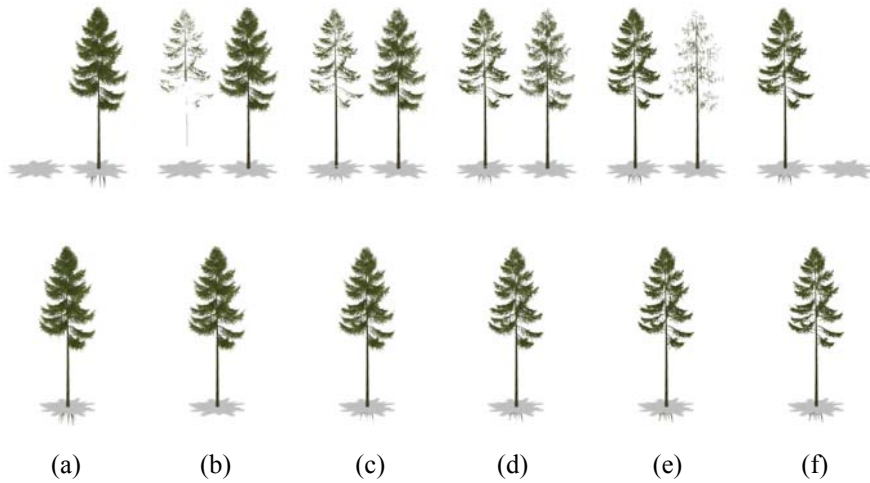


Fig. 4: Upper row: side by side view of the billboard (left) and the geometry model
Lower row: crossfaded view of the two representations

5 Results

Using our system we receive interactive frame rates even for complex scenes that consist of thousands of trees. In Table 1 some results are shown.

Scene	total number of plant positions	average number of displayed plant positions	average number of handled vertices	average framerate
Figure 5/6	365.000	20.000 – 25.000	500.000 – 15.000.000	4 - 7
Figure 7/8	390.000	5.000 – 10.000	15.000.000 – 150.000.000	2 - 5

Table 1: Frame rates of the system for varying scenes, resolution is 1024x768 pixels, with 4x antialiasing and 4x anisotropy filtering

Figures 5/6 , 7/8 and 9/10 show four additional views to some of our scenes and give hints about the visual quality of our approach.

More problems are caused by the blending between the billboard representation and the polygonal model. The full geometric complexity must be used in the direct vicinity of the virtual viewer. In a very dense outdoor environment the number of triangles in some square meters around the camera can reach hundreds of thousands and, thus, can avoid interactive

frame rates. This can only be circumvented by using billboards very early, which sometimes causes parallax problems. In the future we have to think about this last detail problem before reaching our goal of interactively rendering complex outdoor scenes with nearly no restriction to the scene size.




Fig. 5/6: Mixed scene, 300x300 meters, advanced lighting computation model




Fig. 7/8: Forest scene, 1.000x1.000 meters, advanced lighting computation model




Fig. 9/10: Small forest scene, 300x300 meters, advanced lighting computation model

6 References

- [BH65] BALL G., HALL D.: Isodata: A novel method of data analysis and pattern classification. Tech. rep., Stanford Research Institute Tech. Rep. NTIS-AD-699616, Stanford, 1965.
- [BHP99] BAREQUET G., HAR-PELED S.: Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. In SODA '99: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms (1999), Society for Industrial and Applied Mathematics, pp. 82–91.
- [CBL99] CHANG C.-F., BISHOP G., LASTRA A.: LDI tree: A hierarchical representation for imagebased rendering. Proceedings of SIGGRAPH 99 (1999), 291–298.
- [COL04] COLDITZ C.: Methoden zur Nachbildung natürlicher Beleuchtungseffekte bei der Echtzeitdarstellung von Pflanzenpopulationen. In Virtuelle und Erweiterte Realität, 1. Workshop der GI-Fachgruppe VR/AR (2004), SHAKER Verlag, pp. 127-136
- [DCSD02] DEUSSEN O., COLDITZ C., STAMMINGER M., DRETTAKIS G.: Interactive visualization of complex plant ecosystems. In IEEE Visualization 2002 (2002), IEEE, pp. 219–226.
- [DDSD03] DÉCORET X., DURAND F., SILLION F. X., DORSEY J.: Billboard clouds for extreme model simplification. ACM Transactions on Graphics 22, 3 (July 2003), 689–696.
- [DN04] DECAUDIN P., NEYRET F.: Rendering forest scenes in real-time. In Rendering Techniques '04 (Eurographics Symposium on Rendering) (june 2004), H. W. Jensen A. K., (Ed.), pp. 93– 102.
- [DVS03] DACHSBACHER C., VOGELGSANG C., STAMMINGER M.: Sequential point trees. ACM Transactions on Graphics 22, 3 (July 2003), 657–662.
- [EDD*95] ECK M., DEROSE T., DUCHAMP T., HOPPE H., LOUNSBERY M., STUETZLE W.: Multiresolution analysis of arbitrary meshes. SIGGRAPH 95 Conf. Proc. (1995), 173–182.
- [FD03] FRANZKE O., DEUSSEN O.: Accurate graphical representation of plant leaves. In Plant Modelling and Applications (2003), de Reffye P., (Ed.), Springer-Verlag.

- [Fuk90] FUKUNAGA K.: Introduction to Statistical Pattern Recognition. Academic Press, San Diego, CA, 1990.
- [Hop] HOPPE H.: Progressive meshes. In SIGGRAPH 96 Conf. Proc., pp. 99–108.
- [Jak00] JAKULIN A.: Interactive vegetation rendering with slicing and blending. In Proc. Eurographics 2000 (Short Presentations) (Aug. 2000), de Sousa A., Torres J., (Eds.), Eurographics.
- [Max96] MAX N.: Hierarchical rendering of trees from precomputed multi-layer Z-buffers. In Eurographics Rendering Workshop 96 (1996), Springer-Verlag (Rendering Techniques 1996), pp. 165–174.
- [MDK99] MAX N., DEUSSEN O., KEATING B.: Hierarchical image-based rendering using texture mapping hardware. Eurographics Rendering Workshop 99 (1999), 57–62.
- [Ney98] NEYRET F.: Modeling, animating, and rendering complex scenes using volumetric textures. IEEE Trans. on Visualization and Computer Graphics 4, 1 (1998), 55–70.
- [QNTN03] QIN X., NAKAMAE E., TADAMURA K., NAGAI Y.: Fast photo-realistic rendering of trees in daylight. In Proc. of Eurographics 2003 (2003).
- [RB85] REEVES W., BLAU R.: Approximate and probabilistic algorithms for shading and rendering structured particle systems. In Computer Graphics (SIGGRAPH '85 Conf. Proc.) (1985), vol. 19, pp. 313–322.
- [RB93] ROSSIGNAC J., BORREL P.: Multi-resolution 3D approximations for rendering complex scenes. In Geometric Modeling in Computer Graphics. Springer-Verlag, 1993, pp. 455–465.
- [RH94] ROHLF J., HELMAN J.: IRIS Performer: A high performance multiprocessing toolkit for real-time 3D graphics. In Proceedings of SIGGRAPH 94 (1994), Glassner A., (Ed.), pp. 381–394.
- [SD01] STAMMINGER M., DRETTAKIS G.: Interactive sampling and rendering for complex and procedural geometry. In Rendering Techniques 2001 (2001), Gortler S., Myszkowski C., (Eds.), Springer-Verlag, Vienna, pp. 151–162.
- [SGHS] SHADE J., GORTLER S., HE L., SZELISKI R.: Layered depth images. In SIGGRAPH 1998 Conf. Proc., pp. 231–242.
- [WFP*] WAND M., FISCHER M., PETER I., AUF DER HEIDE F. M., STRASSER W.: The randomized z-buffer algorithm: Interactive rendering of highly complex scenes. In SIGGRAPH 2001 Conf. Proc., pp. 361–370.

[WP] WEBER J., PENN J.: Creation and rendering of realistic trees. In SIGGRAPH 95 Conf. Proc., pp. 119–128.