Billboard Clouds for Extreme Model Simplification

Xavier Décoret, Frédo Durand, François X. Sillion, Julie Dorsey

Introduction: We introduce *billboard clouds* – a new approach for extreme simplification in the context of real-time rendering. 3D models are simplified onto a set of planes with texture and transparency maps. We present an optimization approach to build a billboard cloud for a textured polygonal model and a given geometric error threshold. We demonstrate our technique on a large class of models, including smooth manifolds and composite objects.

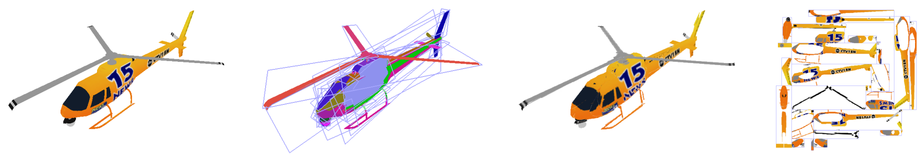

(a) (b) (c) (d)

Figure 1: Example of a billboard cloud: (a) Original model (5,138 polygons) (b) false-color rendering using one color per billboard to show the faces that were grouped on it (c) View of the (automatically generated) 32 textured billboards (d) the billboards side by side.

Approach: We iteratively replace large sets of faces by a textured plane. The texture is generated by projecting the associated set of faces on the plane. The faces and the plane are chosen so that the distance between vertices and their projection is bounded (error threshold), and that the projected area of faces is maximized. To approximate the whole model with a minimum number of planes, we uses a greedy approach that iteratively chooses planes that approximate great quantities of faces, which we call *dense* planes. Density is evaluated in plane space, using a discretization and adaptative refinement for efficiency. A good surface approximation is ensured by favoring planes that are "nearly tangent" to the model. This method does not require connectivity information, but still avoids cracks by projecting primitives onto multiple planes when needed. To avoid large textures with a lot of empty space, which can occur when distant parts of the model are simplified on the same plane, our algorithm takes into account the compacity of the generated textures when selecting set of faces to be simplified. For extreme simplification, our approach combines the strengths of mesh decimation and image-based impostors. Billboard clouds is a new representation that can be used not only for rendering but also quick shadow display or fast ray-object intersection as shown on Figure 2.

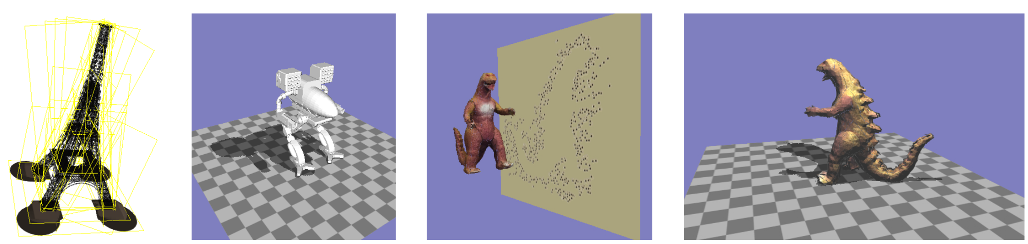

(a) (b) (c) (d)

Figure 2: Billboard clouds encode an object shape and appearance through a few quads and transparent textures. This representation can be used for fast rendering of complex objects (a), for shadow computations (b) or ray-object intersection (c). Storing additional information such as normal map in the textures also allow real-time relighting for improved realism (d).

Progress: The simplification of complex object with billboard clouds gives good results. The shape and details can be captured with a surprisingly low number of billboards. But this low number of billboards can not be straightforwardly compared to the number of faces in a mesh simplification. The rendering of a billboard cloud is penalized mainly by the fill rate : when 100 interleaved billboards are rendered, lots of pixels are drawn many times (the worst case being parallel billboards viewed from front). Therefore, a billboard clouds is useful when its projected screen size is not too large. When an object is far away, using a billboard cloud simplification gives higher quality results than a geometrically simplified version of the object[1] of same rendering cost.

Future: For billboard clouds to be interesting when viewed from closer range, we have to diminish the number of billboards. Unfortunately, if we want the billboard cloud to be viewable from any angle, this number of billboards cannot be lowered too much. The extrem case of a single billboard would obviously look "flat" from many viewpoints. Therefore we plan to investigate object based view-dependent billboards clouds. The space around an object is divided into regions as indicated on Figure 3. A billboard cloud is constructed for each of these regions or *viewcells*. By limiting the view-angle within a region, we expect to get a low billboard count for each cloud. Moreover, the distance range delimited by each viewcell also allows to optimize the resolution of billboards' textures. Therefore, the extra storage caused by multiple view-dependent clouds versus a single view-independent one should not be too high compared to the benefits. Preliminary results seems promising but there is still a lot to investigate, and notably the transitions.

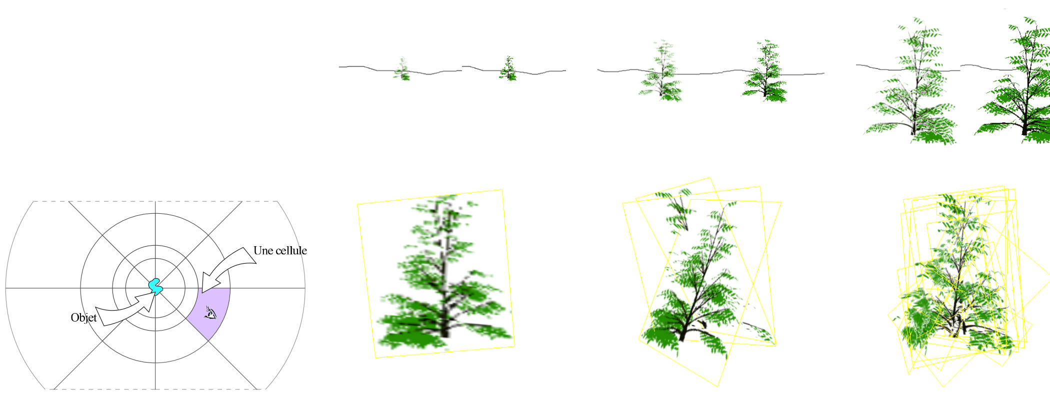

(object based viewcells) (close range) (middle range) (far range)

Figure 3: Examples of view-dependent billboards. A billboard cloud is build for each viewcell. Depending of the distance to the object, the algorithm detects that more and more billboards are required to represent the shape and appearance of a tree model (20 000 polygons). Note also that the further the billboard cloud is, the lower the resolution of the texture is.

Research Support: This research is supported by funding from the NTT Corporation.

References:

[1] M. Garland and P. Heckbert, "Surface simplification using quadric error metrics," *Proceedings of SIGGRAPH 97*, 1997.

[2] X. Décoret, *Preprocessing of large databases for interactive visualisation*, Ph.D. thesis, Université Joseph Fourier, Grenoble, october 2002.

[3] X. Décoret, F. Sillion, F. Durand, and J. Dorsey, "Billboard clouds," Tech. Rep. RR-4485, INRIA, Grenoble, June 2002.