

Model Simplification for Real-Time Rendering

MIT9904-20

Progress Report: July 1, 2002—December 31, 2002

Frédo Durand

Project Overview


The ever-increasing size of 3D models makes display acceleration a crucial issue. It can be addressed through two complementary approaches: fast culling of invisible portions of the scene, and model simplification. In this project, we are investigating new approaches to model simplification that address the limitations of current method

The simplification of a 3D model to a smaller number of primitives has for the most part relied on the greedy decimation of triangles, which can be seen as gradient descent in mesh space. Current solutions are very effective for outputs that still contain a significant number of primitives (several hundred triangles) and that have a continuous manifold structure. In contrast, *extreme simplification* aims at producing models with a number of primitives orders of magnitude smaller. In addition, we want to develop techniques that can simplify any class of models, including complex disconnected models such as vegetation.


We are working on a new representation called *billboard clouds* that bridges the gap between purely geometric (polygon-based) and image-based representations. Billboard clouds are very general primitives that consist in a set of rectangles with texture and alpha (transparency) masks. It is important to note that no only will billboard cloud offer a simplified version that permits faster display; they will also alleviate most flickering and aliasing issues by providing mip-mapping. Simplifying a model into a billboard cloud then reduces to the choice of a set of planes that best approximate the input model. This is an NP-hard geometric optimization problem that we address through heuristics.



Input model



Optimal set of planes




textures and
transparency for the
rectangles




Simplified model using
a cloud of 32 Billboards

Progress Through December 2002

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 use 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 adaptive 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 in the figure below.



The relighting of a Billboard cloud is enabled by the storage of normal maps.




Curve of the number of rectangles need for a given maximum error (in log-log).

Research Plan for the Next Six Months

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 extreme case of a single billboard would obviously look "flat" from many viewpoints. Therefore we plan to investigate object-based view-dependent billboard clouds. The space around an object will be divided into regions. A billboard cloud will be 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 will also permit the optimization of the texture resolution. Therefore, the

extra storage induced by multiple view-dependent clouds versus a single view-independent one should not be too high compared to the benefits. We also plan to work on the transition that occurs both between two billboard clouds representation with a different number of planes, and between billboard clouds and traditional geometry.



Objects simplified using billboard clouds. The rectangles are emphasized in yellow. Note that an object like the Eiffel tower would be extremely challenging to handle using traditional mesh simplification. In contrast, our representation using semi-transparent rectangles, which makes it very effective.