MULTI-RESOLUTION GRAPHIC REPRESENTATION GENERATED BY WEIGHT-CONTROLLED VERTEX CLUSTERING FOR INTERACTIVE VISUALIZATION APPLICATIONS

INVENTORS

Low Kok Lim & Tan Tiow Seng

National University of Singapore

Department of Information Systems & Computer Science

Lower Kent Ridge Road

Singapore 119260

ABSTRACT

A method for execution by a data processor that prepares an object for display during interactive visualization. The method includes a step of preprocessing an original model, which is a polygonal boundary representation of the object, so as to produce a plurality of simplified models with each simplified model having fewer faces and vertices than the original model. The step of preprocessing includes the steps of grading, sorting vertices, triangulation, clustering, thick-edge computation, synthesis, reconstruction of normals, elimination, and storing the simplified model. During display of the object, the method further includes a step of selecting either the original model or one of the simplified models for display so as to achieve real-time interactive motion of the object.

FIELD OF THE INVENTION

The present invention relates generally to computer graphics and solid modeling, and more particularly, to methods and apparatus for representing and displaying an object at multiple levels of resolution.

BACKGROUND OF THE INVENTION

In traditional modeling systems, detailed models are created by applying versatile modeling operations (such as extrusion, constructive solid geometry, and freeform deformations) to a vast array of geometric primitives. For display, these models must usually be tessellated into polygonal approximations. Each of these polygonal approximations typically contain millions of polygons. Enormous computing power is required to render such complex polygonal surface models, in real-time and at interactive frame rates.

In interactive visualization applications, real-time interactive manipulation of viewing conditions is vital to the understanding of an object's geometry and its relative position in the virtual scene. The delays in real-time graphic feedback and the jerky object motion will lead to overshooting, reduced feeling of control, and thus make the user less productive.

In order to accommodate complex polygonal models while still maintaining real-time interactive frame rates, methods for approximating the polygonal models and using multi-resolution models have been proposed. Model simplification algorithms can be used to generate multiple polygonal models at varying levels of detail, and techniques are employed by the display system to select the appropriate level of detail and render it appropriately. These approximations must resemble the original models from all directions, and the transitions between one approximation and the next must be barely noticeable in order to effectively reveal details as objects approach the viewpoint.

One of these model simplification methods is disclosed in U.S. Pat. No. 5,448,686, "Multi-Resolution Graphic Representation Employing At Least One Simplified Model For Interactive Visualization Applications" issued to Borrel and Rossignac. Their method is also described in their paper titled "Multi-Resolution 3D Approximations for Rendering Complex Scenes", Modeling in Computer

Graphics, B. Falcidiena and T. L. Kunii, Eds. Springer-Verlag, pp. 455-465, 1993. This method is very fast in generating simplified models compared to other model simplification methods, and it is effective in achieving high data reduction rates. As such, it is very suitable for use in large Computer Aided Design (CAD) applications where designs are frequently edited and assemblies usually comprise hundreds of thousands of parts. However, the models it produces may be visually very different from the original models, and, moreover, the transitions between different levels of simplification during interactive viewing may not be satisfactory.

It is thus one object of this invention to provide a method for simplifying a polygonal surface model of an object such that (a) the simplified models have image fidelity close to the original model when displayed on the display screen, (b) the transitions between one level of simplification to the next are barely noticeable when the object is displayed using the simplified models, (c) the simplification process requires low computational cost, and (d) the method is effective in achieving high data reduction rates.

It is another object of this invention to provide methods and apparatus that employ one or more simplified object models for display to a viewer, the simplified models being stored and accessed as required. The aim is to provide real-time interactive frame rates during interactive visualization of the objects in the scene.

SUMMARY OF THE INVENTION

The invention described herein satisfies the above-identified needs and overcomes the limitations of the prior invention by Rossignac and Borrel. The present invention provides a method to construct a sequence of graphic models for an object, one for each desired level of simplification. Then, as the object moves in the object space during interactive viewing, the method successively switches to graphic models of the object having the most appropriate simplification resolutions. This results in a suppression of image detail, consequently with an improvement in real-time performance. The inventors have determined that many details in a displayed object can be eliminated, without a loss of visual information.

The method applies a basic principle both in the step of model simplification and the step of selecting the most appropriate model (either the original model or one of the simplified models) for displaying the object during interactive visualization.

The basic principle employed by the method is based on the following argument. In a synthetic scene, when an object is far away from the viewpoint, its image size is small. Due to the discreteness of the image space, many points on the object are mapped onto the same pixel, and this happens often when the object's model is complex and the image size is relatively small. For points mapped to the same pixel, only one point appears on the image at the pixel, and the others are eliminated by hidden-surface removal. This is a wastage in processing as many of such points are processed but never make their way to the final image. A potential solution to cutting down this wasteful processing is to identify those points that are going to fall onto the same pixel and use a new point to represent them. Only this new point is sent for rendering. Points that are most likely to be mapped to the same pixel, regardless of the direction of the viewpoint, are those points which are geometrically very close to one another in the object space.

The model simplification step of the method applies the above principle by using three-dimensional cells to identify vertices which are geometrically close to one another. These cells are placed in the object space to group vertices into clusters. Vertices in the same cluster are considered close to one another, and a new representative vertex is then created to replace these vertices in the cluster. Indirectly, determining the closeness of the vertices also helps to determine the closeness of the polygons. For example, two rectangles are close together if their corresponding vertices are close to each other. By using different clustering cell sizes, the definition of "closeness" will be different, and this allows the original model to be simplified into models of different levels of detail.

During display of the object, the size of the object space covered by a pixel on the display screen is determined, and the model selection step of the method selects the most appropriate model of the object for display. The model selection criterion is based on the determined size of the object space covered by a pixel, and the size of the clustering cells used in generating each of the simplified models.

More specifically, there is disclosed a method for execution by a data processor that prepares an object for display during interactive visualization. The method includes a step of preprocessing an original model, which is a polygonal boundary representation of the object, so as to produce a plurality of simplified models with each simplified model having fewer faces and vertices than the original model. The step of preprocessing includes the following steps:

(1) Grading—a weight is computed for each vertex in the original model to reflect the vertex's visual importance. The value of the vertex's weight is a function of the angles between all pairs of edges incident to the vertex, and the lengths of these incident edges.

(2) Sorting vertices—the vertices in the original model need to be sorted in non-increasing order of their weights.

(3) Triangulation—polygons in the original model are decomposed into non-overlapping triangles without introducing new vertices into the original model. The graded and triangulated original model needs to be stored for subsequent generation of other simplified models.

(4) Clustering—vertices are grouped into clusters based on geometric proximity. The positions of the clustering cells are based on the positions of the original vertices and the weights of the vertices.

(5) Thick-edge computation—for each group of triangles which are going to degenerate into the same edge (due to Step (6) — synthesis), compute an estimate for the thickness of the group of triangles.

(6) Synthesis—a representative vertex is computed to replace the vertices in each cluster and thus simplify some triangles into edges and points.

(7) Reconstruction of normals—the vertex-normals of each triangle are reconstructed.

(8) Elimination—duplicated triangles, edges and points in the simplified model are removed.

(9) Storing the simplified model for subsequent use.

Subsequent generation of other simplified models of the object requires execution of only Step (4) to Step (9) for each simplified model with a new clustering cell size for Step (4). A larger clustering cell size will generate a more simplified model (less detailed) while a smaller clustering cell size will generate a less simplified model (more detailed).

During display of the object, the method further includes a step of selecting either the original model or one of the simplified models for display. The step of selecting a model includes a step of determining the size of the object space covered by a pixel on the display screen, and the step of selecting a model selects either the original model or one of the simplified models as a function of at least the determined size of the object space covered by a pixel, and the size of the clustering cells used in generating each of the simplified models.

The method achieves the objects of the present invention by using:

(a) weight-controlled vertex clustering in which the positions of the clustering cells are based on the positions of the original vertices and the weights of the vertices. This helps to preserve important features of the object, improves the approximation quality, and thus leads to simplified models with higher image fidelity to the original model. Using weight-controlled vertex clustering also results in more consistent placements of important features from one simplified model to the next. The more consistent placements of important features lead to smoother transitions from one level of simplification to the next when object is displayed using the simplified models;

(b) thick-edges to further improve the image quality of the displayed simplified model. Each thick-edge is rendered as a multi-pixel-wide thick-line graphics primitive on the display screen, and a dynamic normal is computed for the thick-line to improve its shading and intensity;

(c) an appropriate clustering cell size to achieve the desired level of simplification and data reduction rate;

(d) a model selection step to select either the original model or one of the simplified models, for display of the object during interactive visualization, as a function of at least the determined size of the object space covered by a pixel, and the size of the clustering cells used in generating each of the simplified models.