Connexion and Tesselation: user modulated display of virtual industrial environments

M. Krus^{1,2}, P. Bourdot¹, S. Tonnoir², F. Guisnel², G. Thibault² and A. Osorio¹

 1
 LIMSI/CNRS, UPR 3251
 2
 EDF/DER, IMA/TIEM/CAO

 1
 BP 133
 1, av du Gle de Gaulle

 F-91403 Orsay Cedex
 F-92141 Clamart Cedex

 tel: +33 (0)1.69.85.81.21
 e-mail: krus@limsi.fr

 fax: +33 (0)1.69.85.80.88
 http://www.multimania.com/mkrus/

Abstract

In order to achieve interactive rates for the rendering of large industrial environments, it is necessary to reduce the complexity of the objects used such scenes. For objects representing piping, ventilation, etc. we use simple primitives which are tessellate at rendering time. We introduce an dynamic tessellation algorithm which adaptively controls the amount of polygons used to approximate the primitives based on the distance from the view point. The main contribution of our work is that we manage connection information so that the meshes of adjacent primitives remain continuous. An ergonomical study was conducted to determine the appropriate levels of details and find out how our algorithm could improve the performance of the user within the virtual environment.

Keywords: virtual environments, adaptive tessellation, user modulation.

Abstract

Afin d'obtenir les taux de rafraichissement nécessaires pour la navigation au sein d'environnements industriels, il faut réduire la compléxité des objets. Pour les objets qui représentent la tuyauterie ou la ventilation, nous utilisons des primitives simples qui sont converties en facettes au moment du rendu. Nous proposons une méthode de facettisation adaptative en fonction de la distance. La principale contribution de nos travaux réside dans l'utilisation d'une information de connexion afin de garantir que les maillages de primitives adjacentes soient continus. Une étude ergonomique a été réalisée afin de déterminer les niveaux de détails nécessaire et de montrer comment la facettisation adaptative pouvait améliorer les performances de l'utilisateur au sein de l'environnement virtuel.

Mots clés: environnements virtuels, facettisation adaptative, adaptation à l'utilisateur

Connexion and Tesselation: user modulated display of virtual industrial environments

M. Krus^{1,2}, P. Bourdot¹, S. Tonnoir², F. Guisnel², G. Thibault² and A. Osorio¹

1	LIMSI/CNRS, UPR 3251 BP 133 F-91403 Orsay Cedex	2	EDF/DER, IMA/TIEM/CAO 1, av du Gle de Gaulle F-92141 Clamart Cedex
	tel: +33 (0)1.69.85.81.21 fax: +33 (0)1.69.85.80.88		e–mail: krus@limsi.fr http://www.multimania.com/mkrus/

Abstract

In order to achieve interactive rates for the rendering of large industrial environments, it is necessary to reduce the complexity of the objects used such scenes. For objects representing piping, ventilation, etc. we use simple primitives which are tessellate at rendering time. We introduce an dynamic tessellation algorithm which adaptively controls the amount of polygons used to approximate the primitives based on the distance from the view point. The main contribution of our work is that we manage connection information so that the meshes of adjacent primitives remain continuous. An ergonomical study was conducted to determine the appropriate levels of details and find out how our algorithm could improve the performance of the user within the virtual environment.

Keywords: virtual environments, adaptive tessellation, user modulation.

1 Introduction

The scenes used for visual simulators of complex and hostile industrial environments, as found in training or project review applications, proved to be very expensive to render. They contain a large number of technical networks (piping, cable trays, ventilation) which are usually modeled by combinations of primitives such as cylinders, tori and cones. These primitives are converted to triangles at rendering time, leading to a huge triangle count. In order to maintain a sufficient frame rate, the resulting meshes must be simplified by building several levels of details which can be used at different distances from the observer. This technique often proves unsatisfactory for this type of objects mostly because the topological relationship between adjacent objects is ignored and discontinuities appear in rendering.

We have chosen a method of real time adaptive tessellation of the primitives based on the distance from the viewpoint. In order to prevent the creation of visual artefacts and ensure a continuity in the mesh, we describe the relations between primitives by the graph of their connections. By evaluating the density of points required at each connection, it is possible to use such a graph to manage the fitting between the tessellations of the connected primitives. While this provided significant performance gains, visual perception criteria needed to be evaluated in order to choose appropriate simplifications levels for the models representing technical components of industrial installations and determine the impact on certain tasks involving recognition of technical components or navigation in complex environments.



(a) Two versions of the same object (b) Mesh discontinuities Figure 1: Sample tessellations of simple primitives

The rest of this paper is organized as follows: we briefly cover the state of the art in object simplification and visual perception of 3D objects. Then, we present the adaptive tessellation algorithm used in our visual simulator prototype. The following section presents the experimental protocol and the theoretical hypothesis we made. Finally, sect. 5 presents the results that were obtained for these experiments.

2 State of the art

2.1 Levels of details

2.1.1 Polygonal simplification

When using simple primitives such as cylinders, cones or tori, an offline tessellation algorithm is traditionally used to produce several meshes for each primitive at different levels of details (see fig. 1(a) showing the same object with 1906 and 368 triangles). Each simplification can then be used at a different distance, simpler models being used further away so that differences from the original model are less perceivable. Simplification of arbitrary meshes is usually achieved by polygonal simplification (see [1] for an overview). However, these methods have several drawbacks:

- Since topological information is not available in polygon-based models, there is no way for the simplification algorithms to know that, for example, a given set of polygons are forming a cylinder. They are thus usually incapable of producing simplification which preserve that topological constraint.
- Each simplified version should ideally be used at a given distance. However, it is not easy to know how many simplifications should be built and at what distance they should be used. It is a matter of finding a balance between the required quality of rendering and the size of the database. However, that knowledge is very application specific.
- Adjacent objects often appear disconnected when simplified, leading to discontinuities and cracks in the mesh (see fig.1(b)). This happens because the primitives are simplified independently of each other and that the knowledge of their connection is not used by (or not available for) the simplification algorithms. Some algorithms resolve this problem by leaving border vertices unchanged. However, in the case of simple primitives such as cylinders, no simplification could then be performed since all vertices are on a border.

2.1.2 Adaptive meshes

Some algorithms are capable of producing geomorphs which contain an entire range of simplifications. These are meshes for which the density of triangles can be interactively controlled. Hoppe uses this method so that parts of the object facing towards the observer will contain more triangles than other parts [2]. Similarly, SGI's library Performer uses the *Active Surface Definition* to combines several pre-computed levels of details and, at rendering time, construct a mesh by interpolating between two of the simplifications based on distance, so that the amount of detail is not uniform over the entire mesh. Both these methods though do not apply well to primitives such as those we have in the database. Furthemore, the connections between objects remain ignored.

2.2 Psychophisical considerations

2.2.1 Visual perception of levels of details

The use of levels of details is based on the assumption that the perception of details varies with the apparent size or speed of object. This was validated by Reddy [3] who used the characteristics of the human visual system to introduce contrast sensitivity functions in order to determine the perception of features in a scene and decide which version of an object should be chosen without the simplification being noticeable. He showed that level of details selection should combine speed and excentricity (distance from the center of field of view) because they individual contribution to reducing the amount of selected detail is negligible. In fact, he has proven that the optimization in size (or distance) are the most beneficial, contributing up to 95% to the performance gain in an application.

Additionnally, psychophysical literature on 3D object and scene recognition showed that the creation and identification of visual representations of 3D object follow different steps. The brain extracts the simple visual features (color, luminosity, contour, orientation, etc) followed by the more complex (orientations of traits in motion). At each step, the information is abstracted and compressed in a more symbolic representation.

By only considering the perceptual efficiency of the human visual system, Reddy did not take into account higher level visual processes such as identification, classification or object recognition. Since the simplifications permitted by the algorithm proposed by Reddy would not be perceivable by the human eye, they did not affect such processes. However, we considered this to be a limiting factor and postulate that further simplifications can be made which would be perceivable but would not affect the comprehension and performance in the simulator. We were interested in the interaction between levels of detail and the representation of the object at a cognitive level, by studying the effect of simplification and orientation on the recognition of an object. We have made the assumption that workers in the installations have a trade-related knowledge which influences their ability to recognize the objects based on the type of primitive shapes used to represent them, the level of tessellation of each shape and the orientation of the objects with respect to the observer.

2.2.2 Object recognition

Biederman [4] postulates that objects and scenes are represented by simple arrangements of volumetric primitives (such as boxes, cylinders, or cones) called "geons". A symbolic "grammar" would allow objects to be read. This seems to indicate that objects would

remain recognizable as long as the simplifications did not degrade these simple primitives beyond a certain level. Furthermore, it would mean that some parts of the object do not participate in the mental representation of the object and that they could be removed. However, selecting exactly which part can be ignored remains difficult in the general case.

Shepard and Metzler [5] showed the existence of prototypical orientations called "canonical views" in which mental representation of objects are stored in the brain (except for non realistic objects). When a user is presented with an object to identify, he performes a mental rotation until the orientation of the new objects corresponds to the canonical view stored in his memory. Shepard and Judo proved a few years latter that the recognition time is linearly proportionnal to the angle value of this rotation [6].

Blanz [7] showed that canonical views are the product of the complex interactions between experience, task and shape. The task asked from the observers influences the orientation in which the canonical view is instantiated. Thus, when subjects are asked to orient an objects as to provide the best impression to another observer, the orientations are not aligned with the axis of the object. At the opposite, when the subject indicates which direction the object is spontaneously pictured, the orientations correspond to the axis of the object.

Additionally, the task asked from the observers influences the orientation in which the canonical view is instantiated. This indicates that orientation will effect the recognition.

Most theories of perception focus on the processing of elemental attributes. But when we observe a scene, we see objects as a whole, independently of the individual elements which make up these shapes. The perceptual organization describes groups of several elements in a higher entity. The gestalt-theory (theory of shape) establishes the rules of perceptual organizations. In that sense, an object as a whole conveys a more significant meaning than the elements which compose it.

Similarly to Biederman's grammar for reading objects, there is a grammar for reading a complete environment. Brewer & Treyers [8] have shown our memory of places is driven by schemata that we construct with our familiarity with "types" of environments. A person structures and organizes his environment based on strategies acquired in his personal experience (this is known as "top-down" information flow, higher levels of abstraction influencing lower levels of recognition). This suggests that users with knowledge of the application area would tend to accept higher simplifications since their experience will help them recognize objects.

3 Adaptive tessellation

Since the current level of details methods fail to preserve the shape of objects in a way that is satisfactory from a psychological point of view, we use an dynamic tessellation method which adapts the complexity of an object to the current view point. To avoid discontinuities between adjacent objects, we introduce a new approach which constrains the adaptive tessellation of objects using connection information.

3.1 Basic geometric primitives

The main type of primitives we consider are those with a circular section such as cylinders, tori or cones. The number of triangles needed to approximate the curvature of such



Figure 2: Propagating the *up vector*

objects can be very important. If two primitives are placed in the scene in such a way that they share an extremity (the end points and the radius coincide) then a connection is constructed to store vertex, normal and texture coordinate information. A connection is either *bound* if it connects two different primitives, or *free* if it belongs to only one primitive. Tessellation is evaluated at each connection. Additionally, the torus is also tessellated adaptively on the number of segments used in the arc (see tap wheel of fig. 1(a)). Each segment in the torus is represented by implicit connections to which other primitives cannot connect. This way, we have a uniform way of building vertices, normals and texture coordinates for all the primitives.

3.2 Building a mesh

Constructing a mesh of several connected primitives can be viewed as a non uniform extrusion along a path linking the different connections, with constraints being propagated from one to the other.

When a scene is loaded, or when animations or user interactions causes an object to move, primitives are tested to determine if connections can be established and lists of connected primitives are build. Vertices are then constructed for each connection. Special care needs to be taken to make sure that points of the two connections of the same primitive are aligned, otherwise primitives would appear twisted. This is done by constructing, for the first primitive in a list, an axis system using the connexion's normal \vec{z} and defining the *up vector* \vec{y} in the direction where the first point of the circle will be. That up vector is then propagated from one connection to the other depending on the type of the primitive:

- for a cylinder or a cone, the axis system is simply copied from the first connection to the second. Changing it between the two extremities of a cylinder or a cone would introduce a twisting effect which would be noticeable.
- for a torus, the axis system is rotated around the axis of the torus by the same angle as the torus. This is done in steps for each of the implicite connections. Thus, the rotation of the up vector is made smoothly over all the connections (see fig. 2).

At each frame, the adaptive tessellation algorithm selects a number of points from the set stored in the connection based on two factors:

• the distance: for each object, the ergonomical study (see next section) has enabled us to define a range of distances between which tessellation should vary linearly

from the highest to the lowest possible value (which where also defined by the study).

• the radius of the primitives: this factor is in fact used to modify the value of the distance range discussed above. If a cylinder of one meter diameter reaches its lowest tessellation level at a distance of one hundred meters, then a cylinder of fifty centimeters diameter will reach it at only fifty meters.

Given a number of points for each connection, a triangle strip is build to represent the primitive. Since each connection might not be at the same distance from the user, the number of vertices may not be the same and the strip might contain degenerate triangles (reduce to a segment). While this is generally not desirable, it has no effect on the rendering (since points and normals coincide) and enable to retain the use of strips which are quicker to render.

4 Ergonomical experiments

The aim of these experiments is to evaluate the loss of information when simplifying virtual 3D objects. We tested the perceptual behavior of various types of users on two visual simulations (isolated 3D object first and later a complete virtual environment), in order to highlight the effect of tessellation, orientation and operability as explained in section 2.3. We wanted to answer three types of questions:

- how many triangles should be used for a primitive at a given distance?
- how does the type of the primitive, its adaptive tessellation and its orientation affect its identification?
- how does the field knowledge of the subjects affect their behavior in the simulator?

Four experiments were designed to answer these questions, under the supervision of A. Drouin and with the help of M.A. Amorim. Three groups participated to the psychophysical experiments:

- Group 1: workers from the piping and plumbing service of the CNPE (Energy Production Unit) of Tricastin.
- Group 2: computer scientists and engineers, expert in VR applications.
- Group 3: people non experts in either production installations or VR applications.

The comparison between the three groups will highlight the impact of schemata knowledge on the results (g1 vs. g3) and the amount of technical computer knowledge needed to complete the various tasks (g2 vs. g3).

4.1 Adjusting tessellation levels

The first two experiments (perception of virtual objects) consisted in presenting four industrial taps available in a CAD database. They were chosen because they were made of simple geometrical primitives such as cylinders or cones, and presented symmetry axis or plane. Additionally, they were presented from nine different angles: three azimuths (or heading, 0, 45 and 90°) and three elevations (or pitch, -45, 0 and 45°). Finally, similarities between facilitated confusions between taps depending on the distance and orientation (see fig. 3) and thus highlighted the importance of the individual shapes in the recognition.



Figure 3: Four taps for the experimentation

These experiments were designed to reconstruct the curve of perception levels for 3D objects based on the distance from the observer. In the first experiment, the taps were presented on screen at a fixed tessellation level (chosen among five), and a random orientation (chosen among nine), either close or far away from the observer. The subjects were asked to move the objects to a distance that they considered optimal (in terms of esthetics) for that configuration. We hypothesized that the distance should be proportional to the amount of degradation.

The second experiment was designed to optimize the maximum number of triangles needed when the objects are presented close up. We expected to find a perceptual barrier above which no extra details are necessary. The taps were presented close up, at various angles, at their best tessellation level. That level was degraded gradually over time until the subjects signaled that they had noticed to change in the quality of the rendering.

The answer should vary depending on the type of primitives which compose the taps and on their orientation. The angles which tend to expose the section of the primitives should result in the object being pushed further way. At the opposite, the particular case of 0° azimuth and 0° elevation, which tends to correspond to 2D schemas of the taps, should make identification easier and tessellation levels less perceivable.

4.2 The effect of degradation on recognition

Having tested the perception of the tessellation, we still needed to test if participants were capable of identifying randomly presented taps based on the mental representation they had acquired. Next to a "canonical" version of the tap shown at the highest tessellation level, an image of a tap was randomly presented various tessellation levels, distances and orientations. The subjects were asked if the two images displayed the same tap or not.

The errors made in this discrimination test should help us resolve the potential ambiguities that certain angles or degradations of essential primitives could introduce and highlight the influence of the simplification on the mental rotation needed to align the two objects.

4.3 Navigation and recognition within a virtual environment

The psychophysical experiments were designed to determine the perception of 3D objects isolated from their environment. This last experiment was intended to study the behavior of participants when they are required to process a complex environment of geometric and architectural configurations. We wanted to find out if an excessive degradation would impact on the efficiency of the simple viewpoint finding task.



Figure 4: Comparing real and virtual scenes

Subjects were shown a picture of an installation and asked to find the viewpoint from which the picture was taken in a virtual model of the same installation (a corridor of a nuclear power plant, see fig. 4) containing roughly 2600 primitives. The primitives in the environment had either fixed levels of tessellation (three levels determined by the previous experiments) or were tessellated adaptively, without the user being told which option was selected.

5 Results

5.1 Numerical results

The first two experiments showed that the optimal tessellation level for primitives in close view is 18 vertices on the section, or 36 triangles for a cylinder or a cone. The minimal number must not be less than 6 vertices (12 triangles) otherwise the section has the shape of a pentagon, which does seem to be identifiable as a circle. For a cylinder of diameter equal to one meter, the distance at which maximum tessellation must be obtained is 60cm. The lowest tessellation level most be reached at 30 meters. These numbers where used to control the rendering in the last experiment producing the following results:

- The fixed environment with 6 vertices (12 triangles) proved too degraded and inhibited the user's ability to recognize and discriminate objects.
- The fixed environment with 18 vertices resulted in the frame rate to decrease to lower rates which made navigation in the scene more cumbersome.
- The fixed environment with 12 vertices presented a satisfactory recognition performance and provided sufficient frame rate for a smooth navigation.
- The frame rates and triangle count for the adaptively tessellated environment were comparable to those obtained for a fixed tessellation of 12 vertices. However, triangles were better distributed since closer objects had more details than further ones. This lead to fewer identification errors and to a better recognition of objects in the environment.

5.2 Types of Primitives

The second experiment showed that for all types of primitives the best tessellation level was 18 vertices. However, the first experiment showed that the adjustment of the distance

varies for each type of primitive. Open angle cones (as in tap 2) are placed further away than closed angle cones or cylinders, mostly because the section (where the tessellation is most noticeable) is more likely to be visible in the first case.

Both experiments show the importance of the viewing angle. It appears that the 0° azimuth and 0° elevation corresponds to typical 2D views as found in industrial schematics of the taps and favors identification. Other things being equal, this angle enables degrade objects to be presented closer than for the worst case angle (90°/-45°). The third experiment confirms these results: confusions between taps increase as the angle moves away from the 0°/0° position.

These results show that not all primitives affect recognition in the same way. Most noticeably, those who expose their section perturb the perceptive judgment of "good shape" for the object (thus confirming the gestalt theory).

5.3 Field knowledge effect

Recognition levels change with respect to the population, the third group (non-expert subjects) saving as reference for the others. Results prove that dynamic visual perception is influenced by the operative knowledge as well as by the task that is assigned to the subject.

In the first experiment, the adjustment levels did not vary from one group to the other since they only depend on the visual ability to process the information. However, we discovered that the industry workers did not focalize its attention on the supporting piping (the bottom most horizontal cylinder in each tap) but rather concentrated on the other parts of the tap. This suggests that the environment is not processed in the same manner as the tap itself. This proves that things that are not directly related to the assigned task can be further simplified than others. It also appeared, as confirmed by the second experimentation, that the group of VR specialists, being accustomed to seeing computer generated images, required more details than average.

Concerning the third experiment, the group of industrial workers used its schemata knowledge of the objects to develop specific strategies. Their judgments are based on a detailed mental representation of the object to be identified. When this object is not representative enough, they try to find a similarities between the displayed object and their operative knowledge which sometimes causes them to confuse one tap with another.

The last experiment confirms, for all groups, that attention of the subject does not focalize on the supporting piping which could thus be further simplified. The operative knowledge of the members of the first group also interferes with their task of finding the appropriate viewpoint. Navigation strategies are influenced by the research of similarities with the professional experience, which sometimes leads them to consider only foreground objects and ignore pertinent information in the background.

6 Conclusion

In order to reduce the amount of triangles used to represent simple primitives, we have presented a dynamic tessellation algorithm which adapts the number of triangles to the observation distance using connection information so that alignment constraints can be propagated from one primitive to the next and that the meshes of adjacent primitives remain continuous. A psychophysical study was conducted to determine the appropriate levels of tessellation for a given distance. The study has also shown that the increase in the rendering performance helped the users navigate within the scene and that adaptive tessellation facilitated the readability of the scenes by preserving the appearance of visible objects.

This work is currently being extended to include other type of objects within our connection oriented framework. New primitives will add extra complexity on the shape of the connection (using parametrical curves for example), the points for these connections are only computed when the objet is modified. However, managing more than two connections per primitive would result in extra complexity for building the mesh at each frame.

References

- [1] Mike Krus, Patrick Bourdot, Françoise Guisnel, and Guillaume Thibault. Levels of detail and polygonal simplification. *ACM's Crossroads*, 3.4, 1997.
- [2] H. Hoppe. View-dependent refinement of progressive meshes. In *Computer Graphics* (*SIGGRAPH'97 Proceedings*), 1997.
- [3] Martin Reddy. *Perceptually Modulated Level of Detail for Virtual Environments*. PhD thesis, University of Edinburgh, 1997.
- [4] I. Biederman. Recognition by components: a theory of human image understanding. *Psychological Review*, (94):115–147, 1987.
- [5] R. Shepard and J. Metzler. Mental rotation of three-dimensional objects. *Science*, (171):701–703, 1971.
- [6] R. Shepard and S. Judo. Perceptual illusion of rotation of three-dimensional objects. *Science*, (191):952–954, 1976.
- [7] V. Blanz, M. J. Tarr, H. H. Bulthoff, and T. Vetter. What object attributes determine canonical views? Technical Report 42, Max-Planck-Institut f
 ür biologische Kybernetik Arbeitsgruppe, 1996.
- [8] W. F. Brewer and J. C. Treyens. Role of schemata in memory for places. *Cognitive Psychology*, (13):207–230, 1981.