

# Geometry simplification for interactive visualization of complex engineering data

J.L. Pajon,<sup>a</sup> P. Guilloteau<sup>b</sup>

<sup>a</sup> *Institut Français du Pétrole, France*

<sup>b</sup> *Centre d'Infographie, Université de Marne-la-Vallée, France*

## Abstract

Interactive walkthroughs in complex engineering data allow to check various spatial properties such as manufacturability, operability and maintainability. These data are so voluminous that various stratagems are necessary to achieve real time frame rates on current graphics workstations. A very common stratagem is to define for each object several levels of detail in order to choose the appropriate level at each frame. The choice of a specific level of detail can depend on various factors: performances must not be obtained at the cost of an important degradation of image quality, except perhaps when the viewer moves rapidly inside the scene. Some techniques for simplifying geometries are reviewed as well as some methods for selecting simplified objects. Some examples of applications are then given.

## 1 Introduction

Design of gas, oil and power plants generates great amounts of geometric objects. Concurrent use of databases for storing these objects and powerful 3D graphics tools allows to visualize selected parts of such plants. To achieve great interactivity on current graphics workstations, the number of displayed polygons and thus objects must be bounded.

A method for displaying more objects at interactive frame rates is to simplify temporarily the objects during 3D manipulations and to restore complete geometries when objects are in a fixed position. Several levels of detail, i.e. several polygonal models of various complexities, can thus be precomputed for each object and then the choice of the model to be displayed depends on the distance between the object and the viewer. Some virtual reality applications require such simplifications, in order to achieve natural manipulations in a continuously changing environment. In this article, we study algorithms described in the literature and practical implementations in various softwares.

Usually, different levels of detail are obtained by simplifying (i.e. removing polygons from) a given initial model. Many methods can be used for achieving such a process. Simplifications of parametric surfaces are easy to obtain. For more general unstructured surfaces, various algorithms can be designed. Some can be very brute-force and dramatically reduce the number of polygons of each model by subsampling points in a systematic way, without taking in account its topology and the curvature of specific regions. More elaborate algorithms have been designed to retain as much as possible the specific features of the initial object. Unfortunately, these algorithms cannot guarantee in all cases a significant reduction of the number of polygons.

After a review of these methods and a comparative description of their various benefits and disadvantages, we will show how these methods allow to perform in a natural way different types of operations while navigating inside virtual oil refineries with various input and output devices, like 6D trackers and head-mounted displays. In particular, it makes easier to verify the overall design of the project for both draughtsmen and engineers and also various technical characteristics such as the accessibility of some critical parts of the plant.

## 52 Visualization and Intelligent Design in Engineering

### 2 Review of some strategies for interactive walkthroughs

Apart from the display of various levels of detail, there are various other strategies for improving frame rates. We would like to stress three approaches which are based on the same idea. To alleviate the work of the graphics processors, they send only a fraction of the scene to them. This fraction is an estimated superset of the polygons that can be viewed. Unfortunately these methods are inefficient each time all objects are visualized.

At each frame, the viewing space is defined as a truncated pyramid in world coordinates. If the bounding box of an object is outside this space, this object can be omitted by the graphics hardware. Such simple culling strategies are often very efficient and their systematic use can only be very helpful. They are available in some graphics libraries like Performer of SGI.

In the design process, it is often easy to identify separate worlds or universes that are invisible from each other. When the user is in one of these worlds, he cannot see anyone of the objects that lie inside another world and that therefore do not need to be sent to the graphics processors. To access from one world to another one, specific elements or portals can be defined. Such strategies have become very popular in virtual reality applications, in particular for building walkthroughs.

Based on the previous idea, a more systematic approach consists in subdividing the space into numerous cells and to precompute for each cell the set of all cells that are visible from it. The algorithms described in the literature [1,5,15] are very specific to architectural models because they rely on the fact that most walls are aligned along two orthogonal axes. In this case, they are very efficient because they use precomputed structures to detect many invisible objects, and they avoid the storage in memory of the entire data, but their usefulness in more general cases can be discussed. Various other visibility precomputations can be imagined but they can be very costly because of their high combinatorial complexity.

### 3 Review of some algorithms for polygonal simplification

#### 3.1 Application Areas and Main Objectives

In various application areas, there is a need for simplifying polygonal models. We can make a distinction between two different ways of generating data, by sensor acquisition or by modelling.

Many data are obtained from physical objects by acquisition techniques such as 3D scanners. Medical scanners produce volume data giving the value of some physical property in each cell of a regularly subdivided parallelepipedic domain while laser range scanners produce clouds of points distributed on the scanned surfaces. Let us summarize the problems encountered in these two cases.

Volume data contain very often millions of voxels (it is not uncommon to have more than 100 images of resolution  $512 \times 512$ ) that are often difficult to exploit with direct rendering methods, although many solutions exist now for some particular types of data, owing to a lot of algorithmical research, some progresses in graphics hardware like 3D-texturing on particular Silicon Graphics systems and some efficient commercial softwares. It is often appropriate to extract surfaces of interest, and in particular polygonal models, from these data. The models generated by most usual methods such as the Marching Cubes algorithm contain too many polygons and the cannot be efficiently rendered on many current machines. This is the case in particular for segmenting medical data or reverse-engineering. Here the main aim is really to replace an initial model by a more concise one that will be the only one to be analysed.

Range data are made of rectangular grids of distances from the sensor to an object. The grid spacings are independent from the shape of the object, such that regions of same extent with low or high curvature are digitized with the same number of points. In order to adapt the grid to the shape characteristics and to use as few points as possible to have a concise but nevertheless accurate representation, simplification methods are required [3].

Many other data are manually obtained by designers using modelling software. In this case, generally the initial data is always present in the visualization process and can be brought at any moment, allowing the user to concentrate on it and not on its simplified models. Such applications include interactive navigation through CAD data [11], visual simulators [10] and video-games.

Although our main interest is in visualization of CAD data, the understanding of simplification techniques and strategies is improved by considering the particular features of the various preceding application areas. Depending on these areas, various sets of rules and constraints can be used for modifying polygonal data.

For example in which way are the initial topology and geometry respected ?

Regarding topological constraints, we can distinguish two great classes of methods, according to the respect or not of the initial topology. Although it can seem very important to preserve topology in applications where, as said previously, focus is concentrated on the simplified data, in other applications where these data are visualized only temporarily during fast motions until the initial model is restored for more profound examination, the topology has not to be preserved absolutely. For example, in the algorithm described in [11], small distinct objects that are close can be merged together, and any object can be reduced to a segment if its shape is elongated or to a point if its shape is spherical. A further distinction [7] can

## Visualization and Intelligent Design in Engineering 53

be made between the methods that explicitly require input data to consist of only one mesh embedding a 2-manifold, or any type of polygonal data of various topologies.

Regarding geometrical constraints, there are also various strategies. For example, it can be crucial in some cases to preserve sharp features and therefore to keep edges and points that contribute to these features. Such edges can be identified as determining dihedral angles exceeding some threshold value.

To discuss the algorithms, we distinguish two great classes of methods:

1) **Clustering Algorithms:** an initial polygonal model is subsampled by regrouping primitives to form only one primitive in the final model. In this case, the two methods that will be described do not guarantee topology consistency.

2) **Iterative Algorithms:** an initial polygonal model is simplified or adjusted step by step until some criteria are fulfilled. A subset of the initial points is looked for, or a new surface can be fitted to the initial points.

Other distinctions can be made. In particular, some algorithms constraint the final model to be made of a subset of the initial points and some others do not deal with such constraints and in this case it is possible to have old or new points in the final model. We can also notice that some algorithms have been adapted to allow a sort of "continuous interpolation" between different levels of detail. The data structures required by these algorithms can also be considered.

We will also distinguish the way in which these methods deal with the two basic tasks they perform: modification of geometry (points positions) or of connectivity. Either these two tasks are made simultaneously at each iteration, or they are made successively in two separate steps.

### 3.2 Clustering Algorithms

In the two methods described in this section, the initial topology of the model is not constrained to be preserved in the coarser models. Old points can be replaced by a subset of them or by new points, that can be computed by averaging methods (which have often the disadvantage of reducing the size of the objects).

#### 3.2.1 Regular space partitioning [11]

The basic idea of this algorithm consists in regularly subdividing the bounding box of the scene and to merge into one point the vertices which lie inside each of the cells of this subdivision. The resulting polygonal models have less vertices. This algorithm is very effective in reducing the number of polygons very rapidly and it is easy to program as it requires only simple data structures and simple traversals of them. Continuous transition from one model to another one is allowed. The shape of simple surfaces is preserved, but the topology of more complex objects can change during simplification and in particular many holes can be introduced (aliasing effects). This sampling process eliminates all details that are smaller than the size of the grid cell, independently from their semantic importance, and it cannot suppress all redundancy: if a planar component is made of polygons that extend over several cells, it cannot replace all these polygons by only one. Furthermore, results are not invariant to translations or rotations of the model.

#### 3.2.2 Hierarchical object partitioning [12]

The idea of this algorithm is also to eliminate small details. Instead of removing all details contained inside cells of the same size, it replaces by one point sets of points that are close to each other. To perform this operation, it stores points in a specific tree where each node admits as children a cluster of points characterized by a dissimilarity measure: the strongest distance between any two of these points. All the points in a cluster can be replaced by a representative if the corresponding dissimilarity measure is smaller than the smallest detail size chosen by the user.

In contrast to the previous method, this method gives results invariant to rigid motion, as it relies on the intrinsic spatial properties of the model and not on an arbitrary subdivision of space. The authors state that this method can only deal with small models comprising less than some hundreds of points because the building of the hierarchical clustering tree can take some time. As the preceding one, it doesn't suppress all redundancy. An initial step performed with one of the following algorithms could be used for this task.

### 3.3 Iterative Algorithms for Selecting a Subpart of a Set of Initial Points

#### 3.3.1 Decimation of triangle meshes [14]

The basic approach starts from a given polygonal model and then tries to improve this model by progressively removing points from it and by filling with less triangles the holes thus created. The criterion for removing a point is its distance to a plane approximating the surface where its neighboring points are lying. If this distance is smaller than a specified value, the point is a candidate for removal. But it will be removed only if it is possible to find two vertices of the border of the hole such that the "split plane" which contains these points and which is orthogonal to the average plane is not intersected by the border edges of the hole.



Arbitrary topologies are accepted and preserved by this algorithm. For example, an edge that is shared by more than two triangles is acceptable but it will not be removed. Each iteration modifies both geometry and connectivity.

### 3.3.2 Adaptive Subdivision, Polygon Growth [3]

These methods have been described in the case of quadrilateral meshes commonly obtained in 3D digitizing techniques. Particular algorithms have been devised in this case. Rectangular areas of points are approximated by polygons. An adaptive algorithm allows to recursively subdivide such polygons if they do not satisfy some fitting criterion. Conversely, polygon growth algorithms allow to combine several polygons together until this same criterion is not satisfied.

## 3.4 Iterative Algorithms for Fitting Surfaces to a Set of Initial Points

The following algorithms are restricted to surface meshes.

### 3.4.1 Mesh Optimization [8]

As in preceding iterative approaches, the algorithm starts from an initial polygonal mesh, that is progressively modified to contain less vertices while preserving the initial shape. These two criteria are expressed in a mathematical way in order to allow numerical algorithms to be explicitly performed. The topology and geometry of the mesh are considered separately. For a fixed mesh, an inner, continuous minimization is made by varying locally the geometry, i.e. by modifying the location of some points. For a fixed geometry, an outer, discrete minimization is made by using three basic transformations (edge collapse, edge split and edge swap) for varying the mesh structure. Local heuristics are used to simplify the inner minimization algorithm. Nevertheless, this algorithm seems to involve complex computations.

### 3.4.2 Re-Tiling [16]

The retiling algorithm is a multi-step process that generates from an initial mesh a more concise representation. In an initial step, an arbitrary number of points (that will be the points of the surface that is looked for) is randomly distributed on the initial surface. These points are then submitted to repulsive forces that move them farther from each other in order to distribute them more evenly. This process can be adaptive to concentrate more points in regions of high curvature. These points are then added to the initial ones to form a *mutual tessellation*. The old points are then removed as in the decimation algorithm. Continuous interpolation from the initial model to the simplified one is possible. As in the previous algorithm, but in a very different way (in two successive steps and not in two imbricated loops) the two essential steps of this algorithm (location of new points on the surface by point repulsion, generation of a new mesh by mutual tessellation) allow to deal separately, in two successive steps (in contrast to the previous algorithms), with the geometry and then the connectivity of the mesh.

### 3.4.3 Deformable Models [4]

The basic idea is to approximate a set of points by a deformable mesh, which has initially a simple shape (a sphere or a cylinder for example) and which is progressively deformed to fit the data. This mesh is attracted towards the points of interest by forces that decrease at each iteration in order to minimize an energy functional. If these points form a surface, the final mesh gives a concise representation of this surface, as its number of points can be fixed arbitrarily low. Interactive tools allow also to add more points, to concentrate points on some regions, to cut the mesh and change its topology in order to obtain a more accurate fitting. To allow real time computations, a special class of mesh is used, a 2-simplex mesh: each vertex has a three-connectivity (it is linked to three vertices), which leads to hexagonal faces. This algorithm proceeds in exactly the reverse order of the preceding one: it deals first with the connectivity and second with the geometry of the mesh.

## 4 Management of levels of detail in real-time visualization

Many approaches can be used to select levels of detail and to make transitions between them. The software designer must take care of the various strategies that will be at the disposal of the user as well as of the degree of control that will be left to him in order to adapt these strategies to his personnel needs.

We can list here some alternatives that can be thought of and some proposed solutions.

- Is there an optimal number of levels of detail for each application ? Can it be computed ? Can the user choose this number ? Usually there are approximately four or six levels of detail, which can include the initial model, its convex hull, its bounding-box, its center point and no point at all. A great number of levels of detail can require huge storage capabilities.

- Are all objects displayed with the same level of detail as is done in some softwares (such as Review of CadCentre) or is each object displayed with a particular level of detail depending on its characteristics and its location ? This second solution seems preferable. Let us consider now this option.
- Which rule can be used to change the level of detail of a particular object ? The most basic approach is to make it depend only on the distance of this object to the viewer as is done in visual simulators or on the size in pixels of each object (the two cases are the same if all objects have identical size). More complex functions than this distance can be evaluated at each visualization step in order to determine which level of detail to use. A general framework has been established in [6] to describe the situation and design efficient strategies. For each object  $O$ , let us suppose that several levels of detail  $L$  and several rendering algorithms  $R$  are available and that with each "tuple"  $(O, L, R)$  we can associate a benefit  $B(O, L, R)$  and a cost  $C(O, L, R)$ . To have the maximum of detail while maintaining some selected frame rate, the set  $S$  of "tuples" to be chosen must maximize the benefit  $\sum_S B(O, L, R)$  with the constraint that  $\sum_S C(O, L, R) < TargetFrameRate$ . Another constraint is that no object can be rendered several times.  $B(O, L, R)$  mainly depends on the perceived size of object  $O$  but it can also depend on the accuracy on its representation, its particular meaning in the application, its location relative to the center of the window, etc.  $C(O, L, R)$  depends on the number of polygons and points that are displayed as well as on the number of pixels covered by  $O$ . This constrained optimization problem guarantees a constant frame rate without degrading too much the image quality, but it is NP-complete as this is a continuous multiple choice knapsack problem. A greedy algorithm to approximate the solution it is to calculate the function  $B(O, L, R)/C(O, L, R)$  for each tuple and to select the tuples that have the highest values until the maximum cost is reached. Frame to frame coherence can be used to update this function incrementally at each step.

Is it possible to use fuzzy logic and genetic algorithms to ease the implementation of the preceding methods, as is suggested by some authors [13] ?

- Let us suppose that the only used criterion for switching models is distance. Is it possible to compute automatically at which distances switchings can be made ? Can simple rules be applied to obtain these distances ? Can they change during animation ? Furthermore, which rule can we apply to select the fraction of polygons that is retained from one distance to the following ? In [10], the author suggests to divide space into several parts of equal sizes (except perhaps the two regions nearest to the user, that can be two times smaller) and to reduce polygon density in proportion to the square of the distance.

As said previously, a more constraining rule would be for example to bound the maximal number  $n$  of polygons that can be displayed. A proposed solution would be the following. If there are five levels of detail, the last one consisting of void models, we can define for each of them an arbitrary maximum number of polygons such that these numbers comes to  $n$ . Let us suppose also that the objects are ordered according to their distances to the viewer. It is then easy to associate the first level of detail with the first objects until the number of cumulated polygons is greater than the specified bounding number, and so on for the following objects and the other levels of detail. This algorithm, simpler than the one described in the previous session, would give perhaps unnecessarily less details, as the maximum number of polygons for each level of detail is fixed.

- How to prevent transitions from one level of detail to another from being too distracting to the user ? In [10], the author suggests to progressively replace a more and more transparent object by a more and more opaque object. In order to avoid continual switching between two models when for example the viewer turns around the same object while maintaining approximately the same distance, the author suggests also to have two switching distances for each transition, the first one for allowing more detail and the second one for removing it.

## 5 Virtual Reality for Complex engineering Data

Virtual reality will be taken here in a very broad meaning as the simulation of tasks that can be performed in real world. In engineering industry, it is a common practice even today to make physical mock-ups of reduced size of complex CAD data before actually building them. The reduced size of these mock-ups is necessary because of the very great size of the future constructions but it has also the great advantage of allowing to detect more rapidly the flaws that can be contained in these data.

In this way, we can distinguish two levels of virtual reality because of the two sorts of real tasks that can be imitated: evaluation of real mock-ups of reduced size, use of real equipment in an industrial plant.

In the first type of virtual experiment, how can we restore the user's sensations ? Let us describe some solutions:

With immersive displays (HMDs, Caves [2], etc) the user can be faced with a virtual mock-up that has exactly the same size as traditional mock-ups. In this way, he could perform almost exactly the same operations as he is used to. It would be interesting also to enable him to interact with some other people

## 56 Visualization and Intelligent Design in Engineering

who could share this evaluation. A Cave could be used for this purpose but in order to restore more exactly the original experience, a see-through HMD would be more appropriate. In contrast to the Cave which surrounds the user with a synthetic environment, a see-through HMD could allow the user to perceive the synthetic imagery only in a specific region of space. If he faces this region, he could see the virtual mock-up exactly in the same location that would have been occupied by a real one. Furthermore, these images could be superimposed not only on the real surroundings but also on the people he is discussing with. For example, a person who faces him at some distance would appear as being behind the virtual mock-up as his legs can be hidden by the corresponding images.

Because of the disadvantages of current immersive displays (cost, bulkiness, etc), we think it would be a good compromise to use a virtual workbench [9] on which stereoscopic images of the future plant could be projected. Furthermore, immersion may be not required in this application as even in real life virtual mock-ups are not built to be immersed into but to be seen from a certain distance. Let us add a final remark: the size of the workbench cannot be the same as that of a real mock-up. In this way, when the user wants to zoom in on some details by moving himself nearer to the workbench he cannot see all the scene. Another question is how much displacement is required from the user in order for him to have views at very different scales. Does this displacement have to be the same as in real world? In fact, the problem is that with a virtual mock-up, details on a more extended area cannot be seen just by moving the eye and that consequently physical motions are necessary.

A computer screen can also be used. The user has a much narrower window than before to view the model, so he must constantly change the position of this window to see more details. This can become a difficult task in complex environments as the user can easily be completely lost in the data. A way to simplify the task is to visualize at the same time a floormap of the model and to locate the position of this window on it, as in 3DIX of IBM.

In the second type of virtual experiment, the goal is to simulate the future use and maintenance of equipment in an industrial plant, the operability, the accessibility all parts of the plant for an operator or for a vehicle, etc. Here also the three preceding alternatives are possible. Immersive visualization seems more appropriate because it enables to approximate more closely natural human interactions.

## 6 Conclusion

We have reviewed some of the main concerns that we can have in view when complex engineering data are to be rendered in real time in order to perform "virtual reality" experiments. In particular, we have been interested in geometry simplification techniques, which have been described in some recent papers and which seem to be promising for interactive navigation in voluminous industrial data.

There are two main classes of simplifying algorithms. Either they remove redundant polygons while respecting the initial topology of given surfaces but in this way there is some limit to polygon reduction, or they can be more brute-force by removing features of any size and even by reducing entire objects to almost nothing. These two types of methods have their own advantages. There is clearly a need for software tools which allow a full range of simplification results, and a lot of work remains to be done to choose the best algorithms from many existing ones in each class, to improve them, and to combine them in integrated software packages, which must answer the particular needs of various application areas (virtual reality, visual simulation, games, but also medicine, computational fluid dynamics where finite elements computations can be made in grids obtained from real data, etc). Difficult compromises have to be made to conciliate high quality rendering with high frame rates, especially in virtual reality applications where small user's motions are constantly recorded. The situation is simpler in desktop systems, where detailed models can be displayed when there is no user's action.

Nowadays, many modellers offer some tools for simplifying data (MultiGen of **Software Systems**, 3D-Studio of **Autodesk**, Explore of **Wavefront**), which can then be used by visual simulation or virtual reality toolkits (Performer of **SGI**, Vega of **Paradigm Simulation**, Clovis of **Medialab**). Some virtual reality toolkits (WorldToolkit of **Sense8**, dVS of **DIVISION**) include their own utilities for achieving such tasks.

Some visualisation tools for engineering data (Review of **CadCentre**, 3DIX of **IBM**) integrate automatic simplification and levelling tools, which eases the final user's task. But these tools have not the flexibility of virtual reality toolkits which can be used to build a large set of interactive experiments and to perform various human factors studies.

Methods for generating levels of detail and for using them to accelerate frame rates are very various from one software solution to another one, and many research results do not seem to have been fully exploited yet. In particular, continuous interpolation from one model to another one (a particular case of morphing) is not usually found in these packages.

We hope that satisfying solutions will appear in the next future.



## References

- [1] Airey J.M., Rohlf J. H., Brooks F.P., Jr. *Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments*, ACM Siggraph Special Issue on 1990 Symposium on Interactive 3D Graphics, Vol. 24, N.2, pp. 41-50
- [2] Cruz-Neira C. *Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE*, Computer Graphics (Siggraph 93 Conference Proceedings), May 1994, pp.135-142
- [3] DeHaemer M.J., Zyda M.J. *Simplification of Objects Rendered by Polygonal Approximations*, Computer and Graphics, Vol. 15, No. 2, 1991, pp. 175-184
- [4] Delingette H. *Simplex Meshes: a General Representation for 3D Shape Reconstruction*, INRIA Research Report 2214, 1994
- [5] Funkhouser T.A., Séquin C. H., Teller S.J. *Management of Large Amounts of Data in Interactive Building Walkthroughs*, ACM Siggraph Special Issue on 1992 Symposium on Interactive 3D Graphics, pp. 11-20
- [6] Funkhouser T.A., Séquin C. H. *Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments*, ACM Computer Graphics Proceedings, 1993, pp. 247-254
- [7] Heckbert P.S., Garland M. *Multiresolution Modeling for Fast Rendering*, Proc. Graphics Interface'94, pp. 43-50
- [8] Hoppe H., DeRose T., Duchamp T., McDonald J., Suetzle W. *Mesh Optimization*, ACM Computer Graphics Proceedings, 1993, pp. 19-26
- [9] Krueger W., Froehlich B. *The Responsive Workbench*, IEEE Computer Graphics and Applications, Vol. 14, No. 3, January 1994, pp.12-15
- [10] Latham R. *An Introduction to Model Switching Real Time Graphics*, Vol. 1, No. 1, 1992, pp. 6-13
- [11] Rossignac J. and Borrel P. *Multi-resolution 3D approximations for rendering complex scenes*, Modeling in Computer Graphics, Facidieno B., Kunii T.L. (Eds), Springer-Verlag, 1993, pp. 455-466
- [12] Schauffer G., Sturzlinger W. *Generating Multiple Levels of Detail from Polygonal Geometry Models*, Proceedings of Second Eurographics Workshop on Virtual Environments, 1995
- [13] Schraft R.D., Neugebauer J., Flaig T. and Dainghaus R. *A Fuzzy Controlled Rendering System for Virtual Reality Systems Optimised by Genetic Algorithms*, Proceedings of Second Eurographics Workshop on Virtual Environments, 1995
- [14] Schroeder W.J., Zarge J.A., Lorensen W.E. *Decimation of Triangle Meshes*, ACM Computer Graphics. Vol. 26, No. 2, 1992, pp. 65-70
- [15] Teller S.J., Séquin C.H. *Visibility Preprocessing for Interactive Walkthroughs*, ACM Computer Graphics, Vol. 25, No. 4, 1991, pp. 61-69
- [16] Turk G. *Re-Tiling Polygonal Surfaces*, ACM Computer Graphics, Vol. 26, No. 2, 1992, pp. 55-64