

A Unified Approach for Simplifying Polygonal and Spline Models *

M. Gopi

D. Manocha

Department of Computer Science
University of North Carolina at Chapel Hill
{gopi,dm}@cs.unc.edu

Abstract

We present a new approach for simplifying models composed of polygons or spline patches. Given an input model, the algorithm computes a new representation of the model in terms of triangular Bézier patches. It performs a series of geometric operations, consisting of patch merging and swapping diagonals, and makes use of patch connectivity information to generate C-LODs (curved levels-of-detail). Each C-LOD is represented using cubic triangular Bézier patches. The C-LOD's provide a compact representation for storing the model. The algorithm tries to minimize the surface deviation error and maintains continuity at patch boundaries. Given the C-LOD's, the algorithm can generate their polygonal approximations using static and dynamic tessellation schemes. It has been implemented and we highlight its performance on a number of polygonal and spline models.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: model simplification, levels-of-detail, surface approximation, spline patches, surface fitting, dynamic tessellation

1 Introduction

In the last few years, the problem of simplifying geometric models has received considerable attention in computational geometry, computer graphics and geometric modeling. Most of the literature has focussed on the simplification of polygonal models. A number of algorithms based on vertex removal, edge collapse, face collapse and vertex clustering have been proposed. Given a polygonal model, these algorithms produce levels-of-detail (LOD) of the original model. Different algorithms vary based on the error metrics used for surface approximation, the underlying representations used for the simplified model, or whether or not they preserve the topology of the original model.

In many applications, models are defined using rational parametric spline surfaces. These include Bézier patches and non-uniform rational B-spline surfaces (NURBS). Most current systems use tensor-product surfaces, though triangular surfaces have been gaining importance as well. Models composed of tens of thousands (or even more) of such surfaces are commonly used in CAD/CAM, surface fitting and scientific visualization applications. Because current graphics systems are optimized to render triangles, a number of methods based on static and dynamic tessellation have been proposed in the literature to generate polygonal approximations of spline models. However, these algorithms generate at least two triangles for each tensor-product patch and one triangle for each triangular patch. As a result, in such cases the lowest level-of-detail consists of tens of thousands of triangles. One possibility is to generate a polygonal approximation of the spline model and generate

its LODs using polygon simplification algorithms. In practice, this approach has two *drawbacks*. First, it leads to data proliferation. Representing a polygonal approximation of the spline model and its LOD takes considerably more space as compared to the original spline model. This becomes a major issue in the representation of very large CAD environments (e.g. for an automobile or an entire submarine composed of hundreds of thousands of spline patches). The polygonal approximation of the entire model may not fit into the main memory. Transmitting such models over the network is also time consuming. The second drawback relates to using static tessellations and a few discrete polygonal approximations. In particular, Kumar et al. [KML95, KMZH97] have highlighted a number of advantages of algorithms based on dynamic tessellation over static tessellation. Besides reduced memory requirements, these include generating varying tessellations for different patches, gradual and smooth switching between two triangulations based on incremental techniques and use of visibility culling algorithms, including view frustum culling and back-patch culling to increase the frame-rate. To make use of dynamic tessellation algorithms, we will like to represent the LODs using spline patches.

1.1 Main Contribution

In this paper we present a new algorithm to simplify polygonal and spline models. The algorithm takes any such model as an input, and generates a new representation in terms of triangular Bézier patches, known as the *C-Model*. Along with the boundary description in terms of triangular patches, it also stores topology and connectivity information as a graph. Given the C-Model, it performs a series of patch merging, vertex removal and swap diagonal operations to generate curved levels-of-detail (C-LOD). The algorithm attempts to minimize the deviation error. The graph representation is used to identify vertices on which to perform the geometric operations and update the topology information. Eventually it produces a series of C-LOD, each represented in terms of cubic triangular patches.

Some of the main advantages of this approach are:

- **Generality:** The algorithm can handle most geometric models represented using surface boundary. It has the ability to simplify all B-rep models that can be exactly represented or approximated using triangular spline patches.
- **A Unified Representation:** Given a model composed of polygons or spline patches, it computes a uniform representation in terms of triangular Bézier patches, the C-Model.
- **Simplifying Spline Models:** A new algorithm for simplifying spline models and computing C-LODs.
- **Efficiency:** The algorithm can handle large objects composed of thousands of patches or polygons and simplify them in a few minutes.
- **Reduced Memory Requirements:** For spline models, the memory required to store C-LODs is at most two times the size of input models.

The algorithm has been implemented and applied to a number of polygonal and spline models. Different models and their simplifications are shown in the color plates at the end of the paper. It works quite well in practice.

*Supported in part by a Sloan fellowship, ARO Contract P-34982-MA, NSF CAREER award CCR-9625217, ONR Young Investigator Award, and Intel.

1.2 Paper Organization

The rest of the paper is organized in the following manner. We survey related work in model simplification, surface fitting and spline rendering techniques in Section 2. We give an overview of our approach in Section 3. Section 4 briefly introduces the reader to triangular patches, and presents algorithms to generate a C-model representation. In Sections 5 and 6 we present algorithms to generate C-LODs from the C-model representation. We describe our implementation in Section 7 and highlight the performance of our algorithm on a number of models. Finally, we discuss some open issues in Section 8.

2 Previous Work

There is considerable literature on simplification of polygonal models, surface fitting and data compression, and tessellation of spline models. In this section, we briefly survey the state of the art.

2.1 Simplifying Polygonal Models

Given a polygonal model, a number of algorithms have been proposed to generate its levels-of-detail. These include algorithms based on vertex clustering [RB93], vertex removal [BS96, SZL92, Tur92] and edge collapsing [HDD⁺93, Hop96, GH97, Gue95, CMO97, COM98]. They use different local and global error metrics for simplifying polygonal models. Cohen et al. [CVMe96] and Eck et al. [EDD⁺95] have presented algorithms that preserve the topology of the original object and give a global error bound on surface deviation. [DLW93, EDD⁺95] have presented algorithms for multi-resolution analysis for surfaces of arbitrary topology types.

Other simplification algorithms include decimation techniques based on vertex removal [SZL92, Sch97] and controlled topology modification [ESV97]. All these algorithms generate a few static LODs. Hoppe [Hop96] has introduced *progressive meshes* to incrementally represent various LOD. Based on progressive meshes, view-dependent simplification algorithms have been proposed [XESV97, Hop97]. Hoppe [Hop97] has also applied the resulting algorithm to small (in terms of number of patches) spline models. The algorithm pre-computes a polygonal approximation for a spline model, performs a series of edge-collapse operations on the polygonal model and stores the result as a progressive mesh. At run-time it refines the progressive mesh as a function of viewpoint [Hop97]. While this approach has some nice properties, its main limitations arise from using static tessellation algorithms and having relatively high memory requirements. There exists algorithms for polygonal models to simplify the topology also ([EM98]).

2.2 Surface Fitting

The problem of fitting spline patches to a polygonal model has been extensively studied in computer-aided geometric design. A survey of different techniques has been given in [Die93]. These include algorithms smooth spline surface over irregular meshes [Sar90, Pet95, Loo94, EH96], using n-sided patches for fitting [LD90] and using spline approximation [dBF73, dB74]. More recently, many algorithms have used subdivision surfaces for piecewise smooth reconstruction [HDD⁺94].

3 Overview of our approach

Our approach for simplifying geometric models makes use of surface-fitting algorithms. It can handle polygonal as well as spline models in a unified manner. The relationship between earlier work on simplifying polygonal models (i.e. generating LODs), tessellating spline surfaces, fitting surfaces and computing C-LODs using our algorithm has been shown in Figure 1. Our ultimate goal is to generate good approximations of the original model. As opposed to tessellating the spline models and generating LODs, we generate C-LODs directly. For interactive display, we generate a triangular approximation of these C-LODs using dynamic tessellation algorithms.

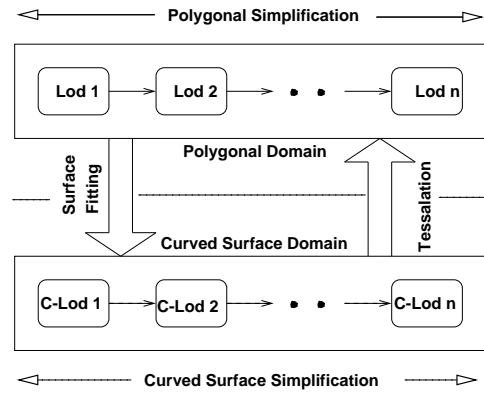


Figure 1: Relationship between Simplification, Surface fitting and Surface Tessellation

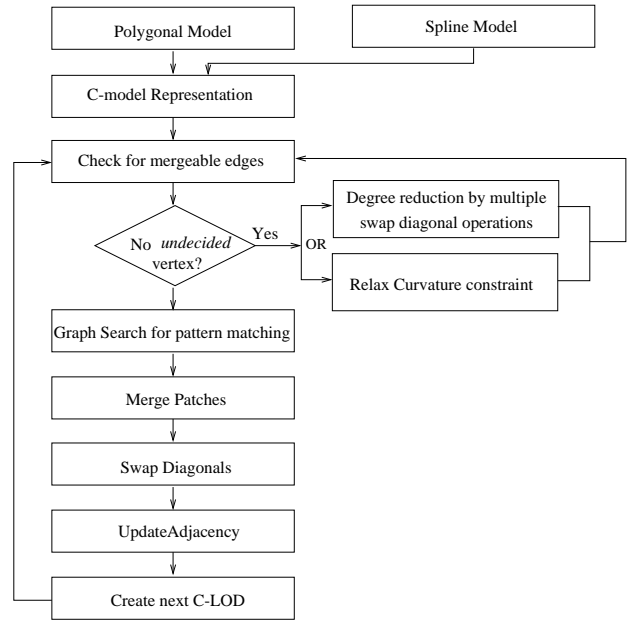


Figure 2: A Flow chart of our simplification system

Our approach towards model simplification starts with generating a uniform representation of the polygonal and spline models in the form of a C-Model representation. The goal of the system is to generate various levels of details (C-LOD) for the C-Model. Our simplification algorithm makes use of vertex removal by merging the patches incident on that vertex. In many ways, these are extensions of simplification algorithms based on vertex removal and face removal for polygonal models. A sequence of such vertex removal operations on a C-Model will generate the next C-LOD. We define a few patch patterns and algorithms to merge the patches forming these patterns. A vertex can be removed only if the set of patches incident on that vertex matches with any of these patterns, so that they can be merged by applying that specific patch pattern merging algorithm. To identify the patterns, and thus all the removable vertices, we use a graph representation and make use of search algorithms. The major components of our system have been highlighted in Figure 2.

- **Representation Conversion:** The input to our system can be a polygonal model, tensor product patch model, or triangular patch model. Any of the above representations is initially converted to a common model representation (C-Model) in

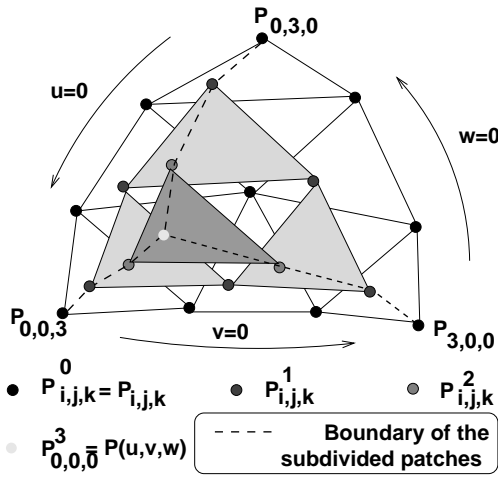


Figure 3: de Casteljau algorithm for subdivision

the form of a triangular patch model, with the complete adjacency information.

- **Computing mergeable edges:** *Mergeable edges* are those pairs of edges incident on a vertex that are amenable to merging. The conditions that the boundary curves of the patches should satisfy to become *mergeable edge pairs* are explained in Section 6.1. Vertices are tagged *retained* if it does not have any mergeable edge. All other vertices are tagged *undecided*.
- **Pattern Matching Graph Algorithm:** The goal here is to tag the status of all the *undecided* vertices into either *removed* or *retained*. Using the information about the *mergeable edge pairs*, patterns are matched around an *undecided* vertex. If the pattern matching is successful, then the vertex is tagged *removed* and the patches forming the pattern are tagged *merged*. The corners of the patch are tagged *removed*. This process continues till there is no more *undecided* vertex.
- **Patch Mergings:** The patches that were identified for merging by the previous algorithm are merged here.
- **Swap Diagonal operation:** The *swap diagonal* operation, is performed on two adjacent triangular patches forming a rectangular patch structure. The common boundary between these patches is eliminated and a new boundary connecting the other two vertices of the rectangular structure is computed. This is equivalent to switching a diagonal line on a rectangle.
- **Adjacency Update:** After these operations, the modified connectivity and adjacency is updated only for those vertices and patches whose adjacencies have been affected.
- **C-LOD Generation:** This module creates the next C-LOD. This includes system cleanup and various bookkeeping operations to prepare the system for the next iteration.

The merging operation introduced in this section is based on the properties of the triangular patches and the well known de Casteljau subdivision algorithm for triangular patches.

4 Triangular Bézier Patches

The Bézier triangular surface [Far93] of degree n can be written as,

$$\mathbf{P}(u, v, w) = \sum_{(i,j,k), i+j+k=n} \mathbf{P}_{i,j,k} B_{i,j,k}^n(u, v, w) \quad (1)$$

where $\mathbf{P}(u, v, w)$ is a point on the surface, with the constraint $w = 1 - u - v$ and $i + j + k = n$. An additional constraint, $u + v \leq 1$, makes the parametric domain a triangle. The constants $\mathbf{P}_{i,j,k}$ are the control points which form a triangular net (Figure 3). The

blending functions $B_{i,j,k}^n(u, v, w)$ are Bernstein basis functions of degree n given by

$$B_{i,j,k}^n(u, v, w) = \frac{n!}{i!j!k!} u^i v^j w^k \quad (2)$$

We consider only polynomial triangular patches, though the techniques given in this paper can be directly extended to the rational patches also. In the rest of the paper, we refer to Bézier triangular patches as just 'patches'.

4.1 G^1 continuous Triangular Patches

For two patches to be G^1 continuous, the triangles joining the control points along the border of the patches (i) should be pairwise planar (*coplanarity* condition) and (ii) should be an affine transformation of the domain triangles (*affine pairs* condition) (Figure 4). In other words, by condition (ii), in Figure 4, the quadrilaterals formed by the triangles $T1$ and $T2$, and $T3$ and $T4$, should be an affine transformation of the quadrilateral formed by $T5$ and $T6$. These two conditions are necessary and sufficient conditions for the patches to be G^1 continuous [Far93].

The three boundary curves of a triangular patch are Bézier curves formed by the boundary control points of the patch. It is important to note that two adjacent Bézier curves are G^1 continuous (their tangents are same in direction, at the common end point), when the three control points – the common control point, and their adjacent control points on either curve – lie on a straight line. We refer to it as an *edge continuity* condition.

4.2 Subdivision and Merging Algorithms

A triangular patch can be subdivided into, in general, three triangular sub-patches using the de Casteljau algorithm. This is also the basis of the blossoming principle for triangular patches.

Given the control points $\mathbf{P}_{i,j,k}$ and a parametric vector $\mathbf{u} = (u, v, w)$, the de Casteljau algorithm computes, at each iteration, a sequence of sets of control points, using the following equation.

$$\mathbf{P}_{i,j,k}^r(\mathbf{u}) = u\mathbf{P}_{i+1,j,k}^{r-1}(\mathbf{u}) + v\mathbf{P}_{i,j+1,k}^{r-1}(\mathbf{u}) + w\mathbf{P}_{i,j,k+1}^{r-1}(\mathbf{u}) \quad (3)$$

with the condition that $i + j + k = n - r$, where n is the degree of the patch and r is the iteration count. Further, $\mathbf{P}_{i,j,k}^0(u, v, w) = \mathbf{P}_{i,j,k}$. The result computed, $\mathbf{P}_{0,0,0}^n(u, v, w)$, is actually the point on the triangular patch for the parametric value (u, v, w) . As shown in Figure 3, the intermediate points generated by the above algorithm, are actually control points of the sub-patches of the original patch.

If the patch is subdivided at a parameter (u, v, w) , where one parameter is zero, then the de Casteljau subdivision would yield two triangular patches instead of three. Further, the boundary curve that is subdivided into two in this process, actually undergoes de Casteljau subdivision of a Bézier curve.

We formulate a method using the inverse of de Casteljau algorithm to merge patches. If three patches are obtained by the de Casteljau subdivision of a single patch, then by performing the exact reverse computation, we will be able to get the control points of the original patch. This requires computing the parametric value (u, v, w) at which these three patches were subdivided. We know that the point $\mathbf{P}_{0,0,0}^n$ is on the patch and also on the triangle formed by $\mathbf{P}_{1,0,0}^{n-1}$, $\mathbf{P}_{0,1,0}^{n-1}$, and $\mathbf{P}_{0,0,1}^{n-1}$. The barycentric coordinate of $\mathbf{P}_{0,0,0}^n$, on the above triangle gives the parametric value (u, v, w) . Using this, by reverse calculation, we can compute $\mathbf{P}_{1,1,1}$ of the original patch. The boundary control points of the original patch are same as the boundary control points of the sub-patches. If the original patch was divided into two patches, the computation of (u, v, w) becomes simpler. One of the parameters is zero, and the others can be computed from the inverse de Casteljau algorithm for Bézier curve, as outlined in [Far93]. These techniques are used in our algorithms for patch merging.




Figure 4: Conditions for G1 continuity

5 Generating C-Models

In this section we discuss algorithms to convert a few types of model representations to our C-Model representation. Our C-Model representation consists of a set of cubic Bézier triangular patches along with adjacency and topology information. It also contains other useful input information about patches that is explained in Section 8. We have chosen cubic triangular Bézier patches for various reasons. These include reduced space requirements and fast algorithms for tessellating them into polygons. Furthermore, the cubic patches provide us with sufficient number of control points for surface fitting. Apart from having good mathematical properties, cubic patches have the right number of parameters and unknowns to work with for our algorithm. The following subsections explain various algorithms we use to convert polygons and tensor-product patches into triangular patches.

5.1 Triangular patches for polygonal models

To convert the polygonal model into a C-Model, we start with fitting one triangular patch for every triangle. An algorithm to find the control points for such a triangular patch, is given below.

A cubic patch has ten control points. We use the following naming convention: three 'corner' control points, six 'boundary' control points and one 'center' control point. The goal now is to position these ten control points, to ensure the coplanarity condition, with a reasonably 'good' surface fit over a polygon.

Choosing the corner control points: The three vertices of a triangle become the three corner control points of the patch.

Choosing the boundary control points: Each corner control point with its two adjacent boundary control points, define the tangent plane at that corner. We also know the tangent plane from the normal at each corner vertex. We choose one point on each of the two edges of the given triangle, that are incident on one corner, say A . The projection of these chosen points on the tangent plane at A would give us the two boundary control points adjacent to A . A reasonable choice of a point on the edge would be the one at a distance of one-third of the edge length away from A . The same method is repeated on the other two corners to get all the six boundary control points.

It is important to note here that a vertex has multiple normals, if it is on an edge or crease. We can make use of this information to avoid G^1 continuity in such cases, to maintain the edges and creases. The polygons whose edges define a boundary of the surface, would be treated as boundary polygons. Now we need to choose the only remaining control point, the center control point. It is initialized to be the average of the six boundary control points. The computed position is refined after all the triangles are fitted with patches, making it a two-pass algorithm.

Refining the center control point: Let us consider one boundary curve, say edge A , of the patch. We know from the coplanarity condition that the center control point, two boundary control points of the edge A , and the center control point of the patch adjacent on the edge A , should be coplanar. For three edges we have

three such planes, and the center control point is the intersection of these three planes. Let us construct the plane for the edge A . The cross-product of the line joining the boundary control points of A and the line joining the two center control points of the adjacent patches, defines the normal to the plane we are looking for. One of the boundary control points, with this normal, defines the plane. The three plane intersection point is then computed by solving three linear equations.

Refining the center control point for a boundary patch: Boundary patches are the patches fitted over a boundary polygon. In such cases, in the absence of adjacent patches, refining the center control point becomes an under-constrained problem. Under these circumstances, the center control point found in the first pass is projected on the line (in case of one boundary edge), or onto the plane (in case of two boundary edges).

5.2 Triangular patches for tensor product patch models

This section explains our algorithm to convert a bi-cubic tensor-product Bézier patch into two triangular patches. If the patch degree is less than three, the patch is subjected to degree-elevation to make it bi-cubic. If the tensor-product patch is of a higher degree, then it is subdivided to smaller patches so that each subdivided patch can be approximated by a bi-cubic patch. The non-isoparametric curves of a bi-cubic patch are of degree six. So, the degree of the triangular patches would be six to exactly represent a bi-cubic patch. We approximate each bi-cubic patch with two cubic triangular patches.

The boundary control points of the tensor-product patch become the boundary control points the two triangular patches. The diagonal curve of the tensor-product patch, which is of degree six, is the third and the common boundary curve for the two triangular patches. We have to approximate the diagonal curve with four control points (a cubic curve).

A bi-cubic tensor product patch is given by the equation:

$$P(u, v) = \sum_{i=0}^3 (\sum_{j=0}^3 P_{i,j} B_j(v)) B_i(u)$$

The points computed within the parenthesis are the control points for the cubic iso-parametric curve for a constant v . In the diagonal curve, $u = v$, and hence they are dependent. A simple extension of the algorithm highlighted above gives the approximation of the control points of the diagonal curve, as given below.

$$P_i = \sum_{j=0}^3 P_{i,j} B_j\left(\frac{i}{3}\right), \forall 0 \leq i \leq 3$$

where P_i 's are the new control points. The center control points of the two triangular patches are found in the same way as given in Section 5.1. Similar methods can be devised to convert n-sided patches into triangular patches.

6 Patch Merging Algorithms

In Section 4, we reviewed methods to combine two or three patches into one big patch when the smaller patches were obtained by the de Casteljau subdivision of the big patch. In this section, we extend this idea to merge any combination of G^1 continuous patches. As this would introduce surface deviation error, we try to minimize this error by imposing various constraints on the patches that are merged.

The foremost requirement for the patches to be mergeable is G^1 continuity. In our application, the patches might not be G^1 continuous to start with. Making the triangular patches in a model G^1 continuous, especially to ensure the *affine pairs* condition, is an optimization problem [Man97]. So, in our system, we only ensure the *coplanarity* condition for geometric continuity. In practice, we obtain good results by only using this constraint.




Figure 5: Patterns for merging

6.1 Merging Patterns

A set of patches can be merged only if they satisfy few geometric constraints and match with one of the defined patterns, in their local topology. Merging algorithms are defined only for these patch patterns. The merging patterns are illustrated in Figure 5. The dark edges are *mergeable edge pairs* for the common vertex of those edges.

We impose two conditions for two edges (curves) incident on a vertex V to be the *mergeable edge pairs* of V . The first is the *edge continuity* as explained in section 4.1. In other words, the common control point V should be on the line joining the two adjacent control points of the two curves. The second condition is the *restricted incident faces* condition. According to this condition, the edge pair should have either 0, 2, or 3 faces between them in both clockwise and counter-clockwise directions. It is possible to have no face between the edge pairs, if the vertex is a boundary vertex. For boundary vertices, only the boundary edges can be mergeable. The use of *restricted incident faces* condition will be made clear later in this section.

The patterns T and *star*, are handled exactly the same way as described in Section 4.2. The *star* pattern has no mergeable edge pairs. But there are several ways to check whether the patches forming the *star* are mergeable or not. One simple test is to check the deviation of the common control point for all three patches from the plane formed by its three adjacent control points. Another test is to find deviation of the normal of the common vertex, computed for all three patches. The presence of *mergeable edge pairs* is an essential condition for all patterns except the *star* pattern. Although, we can decompose a $2T$ pattern into two T patterns, we use this prototype for implementation convenience.

Unlike the T and *Star* patterns, there is no equivalent subdivision process for the patches merged by a T -in- T (Triangle-in-Triangle) pattern. The three *mergeable edge pairs* in this pattern are merged according to the inverse de Casteljau algorithm for Bézier curves, and the center control point is found by the algorithm described in Section 5.1. This pattern cannot be directly extended to higher degree triangular patches, because they have too many degrees of freedom.

Any merging pattern explained here allows not more than three faces around a vertex to be merged on one side of its *mergeable edge pair*. This explains the *restricted incident faces* condition required by the edge pair to qualify as *mergeable edge pair*.

It can be argued that, like a $2T$ pattern, there can be a $3T$ pattern, which would allow 4 faces between the *mergeable edges*. But, as there is no other pattern with 4 faces on one side of the edge pair, the probability of finding such a pattern is quite low. On the other hand, with 3 faces, we have two $2T$ patterns and a T -in- T pattern, making the probability of finding one pattern quite high. This is one of the reasons to impose the *restricted incident faces* condition.




Figure 6: Swap Diagonal Operation




Figure 7: Degree Reduction using Swap Diagonal Operation

6.2 Swap Diagonal operation

A *swap diagonal* operation is not a patch merging operation. But it is very useful in reducing the degree of a vertex as shown in Figure 7. It is also useful in stopping the propagation of vertex removals by the graph algorithm which is explained in the next section. The algorithm to swap a diagonal is not as straightforward as in the case of two adjacent planar triangles. As we are dealing with curved surfaces, the diagonals are space curves. Our goal is to find the diagonal curve which is planar, or atleast with minimum deviation from a plane.

The *swap diagonal* operation is performed in two steps. First the two patches, say A and B , are subdivided into two patches each, say $\{A_1, A_2, B_1, B_2\}$, at the same point on the common boundary curve, using the de Casteljau algorithm. In the second step, A_1 is merged with its adjacent sub-patch of B , say B_1 , and A_2 with B_2 , using the inverse de Casteljau algorithm. This swaps the diagonal of the rectangular pattern (Figure 6). The main issue is to find the parameter at which the original patches A and B have to be split, so that in the second step while merging, the sub-patches satisfy the *edge continuity* condition.

We use a simple algorithm to approximate this parameter. We fit a plane, S , that approximates the four corner control points of the rectangular structure formed by A and B . The normal to the plane is computed as the cross product of the lines joining the opposite corners of the rectangle. On this plane, the four points are projected. The existing common boundary curve is represented as a straight line on the plane. The parameter at which this line is divided by the other diagonal is taken as the required parametric value for our subdivision step.

7 Graph Algorithm for Pattern Matching

In the previous section we introduced various patterns, or prototypes, of the triangular patches that can be merged. In this section, we describe an algorithm to identify such patterns in the boundary description. Once these patterns have been identified, the patches forming the pattern are merged according to that specific pattern merging rule to yield one single patch. This patch is substituted for




Figure 8: Generalized algorithm for three faces merging

the merged patches, thus reducing the patch count of the model to generate the next C-LOD.

Let us assume that we are given a set of vertices and their connectivity in the form of adjacency. A vertex can be tagged either *removed*, *retained*, or *undecided*, with obvious meanings. An *undecided* vertex can be made into a *retained* vertex or a *removed* vertex. But once they are tagged *retained* or *removed*, their status is not changed for that particular iteration.

As the presence of a *mergeable edge pair* is essential to finding a patch pattern around a vertex of degree greater than three, non-existence of *mergeable edge pair* would make the vertex to be tagged *retained*. As a result of the *restricted incident faces* condition, any vertex with degree (number of edges incident on the vertex) more than six will also be tagged *retained*. If all vertices are tagged *retained*, multiple swap diagonal operations have to be performed around a few vertices, to reduce their degree to less than or equal to six (Figure 7).

So all vertices with degree greater than six and the vertices with no *mergeable edge pairs* are tagged *retained*. All other vertices are marked *undecided*. Let us also assume that for each *undecided* vertex, we are provided with its *mergeable edge pairs*.

The goal is to tag the *undecided* vertices as either *retained* or *removed*. If it is *removed*, then find patterns among the patches incident on that vertex, for patch merges, so that no patch has this vertex as its corner. It is clear to see that in *T* and *star* patterns there is one *removed* vertex, in a *2T* pattern there are two *removed* vertices and in a *T-in-T* pattern there are three *removed* vertices. In all the cases there are three *retained* vertices.

With the topology information being coded in the form of *mergeable edge pairs*, we pose the problem of identifying patterns in this topology, as a graph searching problem. The vertices of the model are mapped to the vertices of the graph, and the boundary curves of the triangular patches, to the edges of the graph. We use depth first search algorithm to find patterns.

Let us assume that we are given with a vertex V , that is tagged *removed*, with its *mergeable edge pair*. Let us also assume that a pattern has been matched on one side of the *mergeable edge pair*. Now the problem is to match a pattern on the other side, which we call an *unmatched side*. It is important to note that the two adjacent vertices of V along its *mergeable edge pair*, are *retained* vertices.

There can be either 0,2, or 3 faces on the *unmatched side*, as any other number of faces would contradict the *restricted incident face* condition for the given *mergeable edge pair*. If there is no face on the *unmatched side*, then we are done. If there are two faces then that portion of the geometry can be matched with a *T* pattern. The third vertex is tagged *retained*.

We describe the following generalized algorithm to handle three faces on the *unmatched side*. The immediate neighborhood of any such vertex is *topologically* equivalent to the geometry shown in

Figure 8. The vertex to be removed is denoted by a circle, and *mergeable edge pair* is shown as a thick edge. The shaded vertices are retained vertices. This geometry can be reduced to a *T* pattern by one swap diagonal operation, and the vertex under consideration can be removed, without generating any more removable vertices. In this method, tagging the vertices to be *removed* or *retained* has to be done carefully. The vertices on the either side of the *mergeable edge pair* are always *retained* vertices. Among the other two vertices, one loses an edge and the other gains an edge because of this swap diagonal operation. The vertex that gains an edge should be either *undecided* or *retained* before the swap diagonal, and has to be tagged *retained* after the operation. The vertex that loses an edge can be in any of the three states before the swap diagonal operation. It should be tagged *retained* after the swap diagonal operation, if it was not tagged *removed* before the operation [Gop98]. We will adopt this solution, when none of two *2T* pattern or a *T-in-T* pattern match with the portion of the geometry.

If a *2T* pattern is matched, one vertex is tagged *retained* and one vertex is tagged *removed*. For the removed vertex, a pattern has been matched on one side of its *mergeable edge pair* and there is one *unmatched side*. Similarly a *T-in-T* pattern generates two such removed vertices. The above pattern matching algorithm is applied to the removed vertices recursively.

It can be seen that even if a *2T* or *T-in-T* pattern is identified on the *unmatched side*, the generalized solution can be applied. The advantage of the generalized solution is that it generates no more removable vertices, and is applied to stop the propagation of vertex removals.

The recursion starts with a vertex where both sides of the *mergeable edge pairs* are *unmatched sides*. On one side, we match patch patterns, and in case of a three face pattern, we stop the recursion immediately using the generalized algorithm for handling three faces. Then the second *unmatched side* satisfies the required condition for the above explained algorithm. An example for pattern matching graph algorithm is given in Figure 9, where $V1$ is the first vertex to be removed, and SD denotes a swap diagonal. It can be seen that different sets of patch patterns are possible in the same topology. We attempt to find just one set of patch patterns out of various possibilities.

8 Implementation

In this section, we present implementation details of our simplification system. The algorithms described in this paper have been implemented in C++. We will discuss the system in stages as shown in Figure 2.

C-Model Representation: The polygonal and spline representations are initially converted to the C-model representation. The C-model also has information about the curvature and size of every cubic triangular patch, which are used to decide whether two adjacent patches can be merged or not. A high curvature edge is not merged with any other edge, and a relatively small patch is chosen for merging at the beginning of the pattern matching algorithm.

Computing mergeable edge pair: The *edge continuity* condition for *mergeable edge pairs* is checked by comparing with a tolerance value the angle of deviation of the two line segments connecting the common control point to its adjacent neighbors in the two curves. At higher levels of detail, this tolerance is slightly relaxed. All vertices are *undecided* to start with. If a corner does not have any mergeable edges then it is tagged *retained*.

Pattern Identification by graph algorithm: As explained in Section 7, the graph searching algorithm identifies patches to be merged and the vertices to be removed. The actual merging operation and the vertex removal operation is not performed by the graph searching routine. The patches to be merged are logged in a data structure class called '*ToBeMerged*'. This is an array of sets-of-patches to be merged. Furthermore, whenever a swap diagonal

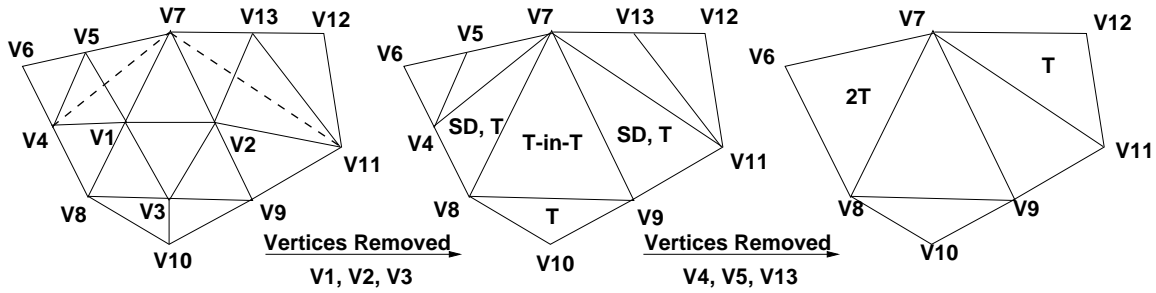


Figure 9: Illustration of the graph algorithm

operation is to be performed, it is logged into another class called *ToBeSwapped*.

Implementing MergePatches: The *MergePatches* function merges every set of patches in the *ToBeMerged* array. The adjacency information can be updated from the data structures.

Implementing SwapDiagonal: The *SwapDiagonal* function swaps diagonal of all the related patches logged in the *ToBeSwapped* array. Unlike the *MergePatches* operation, the adjacency is updated immediately after each *SwapDiagonal*, to avoid an inconsistent state in the topological information introduced by the swap diagonal operation. This inconsistent state is due to the increase in the number of edges, hence the faces, of two of the vertices involved in the swap diagonal. Other merging operations will either maintain the edge count or decrease it.

Adjacency Update: Incremental update of adjacencies and connectivity is performed only for those vertices whose adjacency has been affected by the above operations. The adjacency information involves computing the connectivity between the corners and patches, ordering of edges and faces around a corner, and other book-keeping operations for the affected vertices. Every adjacency update is stored in an array in the vertex data structure, and is tagged with the present C-LOD number. If the C-LOD number of the array i is C_i and that of $i + 1$ is C_j , then for any other C-LOD number between C_i and C_j , array i stores the adjacency information.

Resetting the System: This involves tagging all vertices *undecided*, freeing the unused memory in the structures related to removed vertices and merged patches, and resetting the counters and other system variables. After resetting, the system is ready for the next iteration.

8.1 Data structure

In this section, we highlight the data structures used in our implementation. The main classes in the system are *Model*, *Corner*, and *Patch*. From the *Model* class, all C-LODs of the given model can be obtained. It has *Vertices[]*, an array of pointers to the *Corner* class, and *Patches[]*, an array of pointers to original *Patch* class. As no corner is added to the system during the simplification process, *Vertices* is not changed. If the corners are removed, they are tagged as *removed*, and not removed from the data structure. On the other hand, new patches are added to the system by the merging operation. Pointers exist in the *Model* class, only to the patches in the original model. No direct pointer is provided for each generated patch. Patches belonging to a specific C-LOD of a C-Model can be extracted by an interesting *Patch* class traversal using a linked list and is explained in [Gop98].

The algorithms and data structures have been designed to save run-time memory usage. A good balance of speed against memory is achieved in the implementation. In all the functions, the option of parallel implementation is kept open by carefully designed data-access patterns.

(a)	(b)	(c)	(d)	(e)
Phone	Polygonal	43537	1623	235
Bunny	Polygonal	69491	1000	289
Lion	Tensor Patch	4632	402	10.7

Figure 10: Performance of the system: (a)Model (b)Model representation (c)Initial patch count(d) Final simplified patch count (e) Total time taken (secs)

8.2 Generating higher levels of details

The above modules can generate subsequent C-LODs, while there are *undecided* vertices and patch patterns to remove them. The absence of *undecided* vertices may be for two reasons: the degree of the vertices are greater than six, or there is no *mergeable edge pair*. When the simplification process stops because of these conditions, the multiple swap diagonal operations are performed to reduce the degree of a few chosen vertices, as explained before (Figure 7). To make more edge pairs *merge-able*, we move the boundary control points of all the patches, without introducing cracks in the model, to make the boundary curves planar. Apart from increasing the possibility of edges becoming *mergeable*, this process also makes the quality of the merged patches better in terms of surface deviation error. Relaxing the curvature constraint also alleviates this problem.

9 Performance

The algorithms presented in this paper have been applied to various polygonal and spline models. Here we highlight the results on three models, two of which are polygonal models, and one is a tensor product Bézier patch model. The performance numbers for these models are tabulated in Figure 10. The Phone and the Bunny model were given in *'ply'* format with polygon adjacencies. The Lion model was converted to a triangular patch model and adjacencies were generated off-line. The time given in the table, includes the C-Model representation conversion time and the total time to generate various levels of detail until the patch count given in the table is achieved. For the polygonal models, we computed 200 C-LODs and for the lion model we had 14 C-LODs. Every patch in the lion model has been hand-crafted to contribute to represent features of the model, so the simplification achieved in this model, cannot be directly compared with the simplification achieved in polygonal models. There are many more degeneracies to be handled in curved surface models, than in polygonal models. Following are few of the examples for such degenerate conditions: (i) two triangular patches can have same corner vertices and the boundary curves, but still not be the same patch geometrically, (ii) even in a manifold two adjacent corners can be the end points of a multiple number of patches, (iii) a patch can be folded so that two of its corners coincide, and still be a valid patch. Most of the degeneracies can be handled by our system.

The results shown here are from the prototype of our system. All

the timings presented here were measured on an SGI-Onyx with an R10000 processor, 195 MHz clock.

10 Error Analysis

An important component of a simplification algorithm is to compute a tight error bound on the simplified model. In our algorithm no strict error metric had been imposed on the simplification, except in the form of *mergeable edge pairs*. The error metric we are working on is based on the surface deviation from the original model of the resulting patch after merging. Computing exact maximum surface-to-surface distance is rather complex.

The error introduced in patch merging can be easily computed in the following manner. First the new patch obtained by the patch merging operation is subdivided using the de Casteljau algorithm at the same parameter at which it was merged. If there was no error in the merging operation, then these subdivided patches would exactly match original patches. Hence, a good error estimate is to find the maximum distance between the corresponding control points of these subdivided patches from those of the original patches. We can also use a weighted distance function, as the deviation of the corner control points would involve more error than the same deviation of the center control points. This error can be accumulated incrementally over various C-LODs.

We have observed that the swap diagonal operation, which does not contribute to the patch reduction, can introduce deviation error. We are currently working towards a better *swap diagonal* algorithm and its implementation. Some error is also introduced by not satisfying the *affine pairs* condition for G^1 continuity. We are planning to implement the algorithm presented by [Man97]. This is based on local parametric scheme that improves surface shape by empirically proven improved settings of the free parameters.

11 Conclusion and Future Work

In this paper, we have presented a new algorithm for simplifying spline models without tessellating them into polygons. Our approach is general and also applicable to polygonal models. The resulting algorithms have been implemented and demonstrated on different polygonal and spline models. Our system is well designed and suitable for parallel implementation, on shared memory multi-processing environments. As part of future work, we are trying to make the merging operation and the swap diagonal operation, in the simplification process more accurate and bound the error on merging operations. We are also planning to make the patches strictly G^1 continuous, for improved visual quality, and use results from [EM98] to perform topological simplification.

References

- [BS96] C. Bajaj and D. Schikore. Error-bounded reduction of triangle meshes with multivariate data. *SPIE*, 2656:34–45, 1996.
- [COM98] J. Cohen, M. Olano, and D. Manocha. Appearance Preserving Simplification. In *Proc. of ACM SIGGRAPH '98*, To appear.
- [EM98] C. Erikson and D. Manocha. Simplification Culling of Static and Dynamic Scene Graphs. Technical Report, TR98-009, Department of Computer Science, University of North Carolina at Chapel Hill, 1998.
- [CMO97] J. Cohen, D. Manocha, and M. Olano. Simplifying polygonal models using successive mappings. In *Proc. of IEEE Visualization*, pages 395–402, 1997.
- [CVM96] J. Cohen, A. Varshney, D. Manocha, and G. Turk et al. Simplification envelopes. In *Proc. of ACM Siggraph '96*, pages 119–128, 1996.
- [dB74] C. de Boor. Good approximation by splines with variable knots-II. In G. Watson, editor, *Numerical Solutions of Differential Equations*, pages 12–20. Springer-Verlag, New York, 1974.
- [dB73] C. de Boor and G. Fix. Spline approximation by quasi-interpolants. *J Approx. Theory*, 8:19–45, 1973.
- [Die93] P. Dierckx. *Curve and Surface Fitting with Splines*. Oxford University Press, New York, 1993.
- [DLW93] T. Derose, M. Lounsbery, and J. Warren. Multiresolution analysis for surfaces of arbitrary topology type. Technical Report TR 93-10-05, Department of Computer Science, University of Washington, 1993.
- [EDD+95] M. Eck, T. DeRose, et al. Multiresolution analysis of arbitrary meshes. *SIGGRAPH 95 Conference Proceedings*, pages 173–182.
- [EH96] Matthias Eck and Hugues Hoppe. Automatic reconstruction of B-Spline surfaces of arbitrary topological type. *SIGGRAPH 96 Conference Proceedings*.
- [ESV97] J. El-Sana and A. Varshney. Controlled simplification of genus for polygonal models. *Proc. of IEEE Visualization*, pages 403–410, 1997.
- [Far93] G. Farin. *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. Academic Press Inc., 1993.
- [GH97] M. Garland and P. Heckbert. Surface simplification using quadric error bounds. *Proc. of SIGGRAPH '97*, pages 209–216.
- [Gop98] M. Gopi, D. Manocha. A Unified Approach for Simplifying Polygonal and Spline Models. Technical Report TR98-021, Department of Computer Science, University of North Carolina at Chapel Hill, 1998.
- [Gue95] A. Guezic. Surface simplification with variable tolerance. In *Second Annual Intl. Symp. on Medical Robotics and Computer Assisted Surgery (MRCAS '95)*, pages 132–139, November 1995.
- [HDD+93] H. Hoppe, T. Derose, et al. Mesh optimization. In *Proc. of SIGGRAPH '93*, pages 19–26.
- [HDD+94] Hugues Hoppe, Tony DeRose, et al. Piecewise smooth surface reconstruction. *Proc. of SIGGRAPH '94*, pages 295–302.
- [Hop96] Hugues Hoppe. Progressive meshes. In *Proc. of SIGGRAPH 96*, pages 99–108.
- [Hop97] Hugues Hoppe. View dependent refinement of progressive meshes. In *Proc. of SIGGRAPH 97*, pages 189–198.
- [KML95] S. Kumar, D. Manocha, and A. Lastra. Interactive display of large scale nurbs models. In *Proc. of ACM Interactive 3D Graphics Conference*, pages 51–58, 1995.
- [LD90] Charles Loop and Tony DeRose. Generalized B-spline surfaces of arbitrary topology. *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 347–356, August 1990.
- [Loo94] Charles Loop. Smooth spline surfaces over irregular meshes. In *Proc. of SIGGRAPH '94*, pages 303–310.
- [Man97] S. Mann. Surface interpolation with triangular patches. In *Fifth SIAM Conference on Geometric Design*, 1997.
- [Pet95] J. Peters. C1 surface splines. *Siam J. of Numerical Analysis*, 32(2):645–666, 1995.
- [RB93] J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering. In *Modeling in Computer Graphics*, pages 455–465. Springer-Verlag, June–July 1993.
- [KMZH97] S. Kumar and D. Manocha and H. Zhang and K. Hoff. Accelerated walk-through of large spline models. In *Proc. of ACM Symposium on Interactive 3D Graphics*, pages 91–102, 1997.
- [Sar90] Ramon F. Sarraga. Computer modeling of surfaces with arbitrary shapes. *IEEE Computer Graphics and Applications*, 10(2):67–77, March 1990.
- [Sch97] W. Schroeder. A topology modifying progressive decimation algorithm. In *Proceedings of Visualization '97*, pages 205–212, 1997.
- [SZL92] W.J. Schroeder, J.A. Zarge, and W.E. Lorensen. Decimation of triangle meshes. In *Proc. of SIGGRAPH '92*, pages 65–70.
- [Tur92] G. Turk. Re-tiling polygonal surfaces. In *Proc. of ACM SIGGRAPH '92*, pages 55–64.
- [XESV97] J. Xia, J. El-Sana, and A. Varshney. Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):171–183, June 1997.