

ROBUST INCREMENTAL POLYGON TRIANGULATION FOR SURFACE RENDERING

Subodh Kumar¹

Center for Geometric Computing
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA

subodh@cs.jhu.edu

<http://www.cs.jhu.edu/~subodh>

ABSTRACT

This paper presents a simple, robust and practical, yet fast algorithm for triangulation of points on the domain of trimmed Bézier surfaces. These R^2 points are input to this algorithm by a surface sampler. A set of polygons is formed from these samples, which are then triangulated. We also show how to update the triangulation when the samples, and hence the polygons, are updated. The output is a set of triangle strips. The algorithm includes heuristics to avoid long and thin triangles. In addition, it also detects if the sampling of the trimming curve forms any non-simple polygons and corrects the triangulation by adding more samples. The triangulation algorithm is more generally applicable to polygons in a plane. We report an implementation of the algorithm and its performance on extensive surface-model walkthrough.

Keywords: Surface rendering, CAD, Triangulation, Polygon, PSLG, Computational geometry.

1. INTRODUCTION

View-dependent surface triangulation is a popular technique for interactive display and walkthrough of large geometric models like those of ships and submarines. Such real-time inspection in a virtual environment provides the sense of space and is crucial for simulation based design. It can reduce the time and cost of manufacturing by eliminating the need for full-scale mockups. Such view-dependent techniques enable fast rendering by generating more polygons the region close to the viewer and less elsewhere. However, such schemes require fast re-sampling and re-tessellation of surfaces. To efficiently perform such updates, we devise an incremental triangulation algorithm for parametric surfaces that allows addition and deletion of samples. The re-sampling algorithm is presented elsewhere [KML96,Kum96].

Our system uses the Bézier representation as its core primitive. A number of techniques have been proposed for sampling and tessellation of trimmed Bézier patches [RHD89,AES94,PR95,KML96]. Our algorithm performs triangulation by factoring the

domain into a set of rectangles and a planar straight line graph (PSLG), and then triangulating the PSLG. For our purpose, a PSLG is a planar polygon with additional vertices and edges enclosed. Fig. 1(c) shows an example in thick lines. (Note that triangulation of such PSLGs is equivalent to that of general PSLGs.) Our algorithm allows *incremental updates* and is more efficient. Additionally, we have incorporated polygon *simplicity detection* at little extra cost. Consequently, our implementation is *robust*. We also produce few long and thin ('skinny') triangles. Skinny triangles have a small internal angle. Furthermore, unlike most recent algorithms, we directly generate *triangle strips*, which are more efficient to render on current graphics systems than are triangles. Algorithms to generate long strips from triangles exist [EAV99], but they are intended for static models and are rather slow. Several asymptotically efficient polygon triangulation algorithms [FM84,Cha91] are known but most are difficult to implement and they produce generally skinny triangles. It is possible to ensure triangle 'fatness' using additional (Steiner) points [BDE92,Mit93,KU99]. For example, in two

¹ Supported in part by NSF Career grant 9733827, ARO grant DAAH04-96-1-0013 and NSF ERC grant 9731748.

dimensions, a polygon may be tiled using triangles with angles at most $7/8\pi$ using $O(n^2 \log n)$ Steiner points in time $O(n^2 \log^2 n)$ [Mit93]. Unfortunately, these Steiner points greatly increase the number of triangles generated. Moreover, Steiner points are added on polygon edges. In our application such on-edge Steiner points can result in cracks [KML96] and cannot be used. The rest of this paper is organized as follows. In section 2, we describe the application and

2. PROBLEM DESCRIPTION

A tensor product parametric surface, $S(u,v)$, is defined by a vector function over the domain $(u,v) = [0,1] \times [0,1]$. Optionally, a closed sequence of trimming curve, $C_i(t)$, each defined over the domain $t = [0,1]$ may be defined on the domain of S (see Fig. 1). For brevity, we will refer to the sequence as a single trimming curve, C . The trimming curve restricts the domain of S . By convention, we say a trimming curve retains the part of the domain locally to its right. This retained part is what we triangulate. In this paper, we only consider the trimming curves that comprise a single connected component (other than the regular 0-1 domain boundary). A method to reduce multiple components to this case is described in [KM95]. First, sample points on the domain of the surface (forming a grid) and the trimming curve (forming a polygonal chain) are chosen [KML96]. The sampling is performed for every frame in an interactive graphics simulation. Since the sampling density for a surface does not change much in consecutive frames, it is more efficient to modify the triangulation of the previous frame than to re-generate the whole triangulation anew. Triangulation of these samples is the subject of this paper. Thus, we have –

Initial Input: a grid of points and a closed polygonal chain (Fig. 1(b)) on the 2D $u-v$ domain. These input points are uniformly distributed on the domain and the screen-space distance between adjacent points is

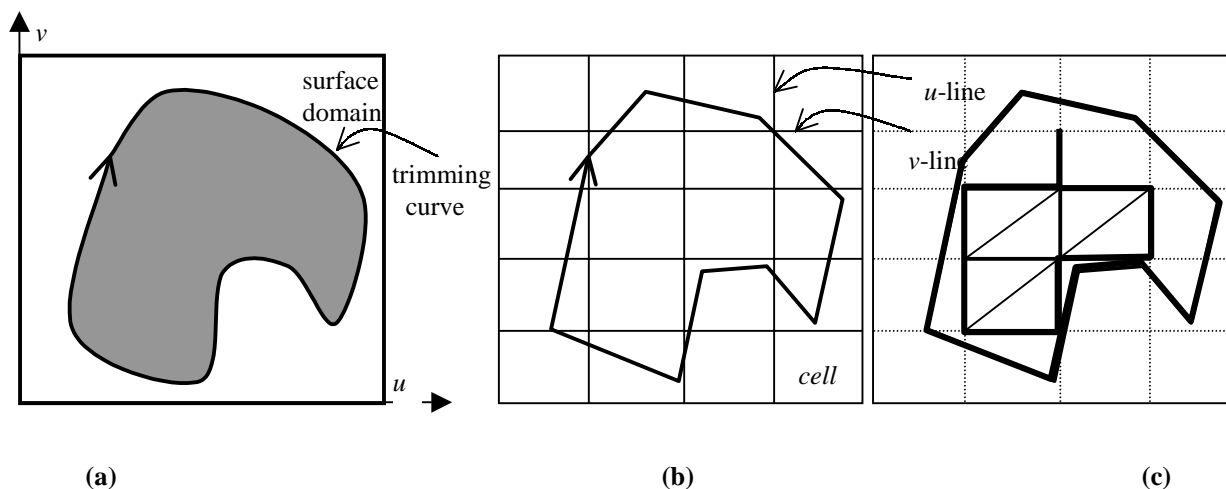


Figure 1: Trimming curve, samples, grids and PSLG on the planar surface domain

bounded. In addition, surface re-sampling results in –
Update Input: add or delete a grid-line or a polygon vertex. While multiple changes on a surface can be

the problem. Section 3 presents the polygon-tracing step that is used to construct a set of polygons. Section 4 discusses polygonal triangulation. In section 5, we describe our data structures for efficient point location and present the incremental algorithm. Finally, we provide the implementation results in section 6 and conclude in section 7.

simultaneously handled, it is easier to describe the algorithm in terms of a single update at a time.

Output: a triangulation such that no triangles lie outside the domain restricted by the trimming curve’s tessellation. In addition, no edge of the triangles generated may be greater than twice the maximum distance between adjacent input samples. Bounded edge length ensures small tessellation-error [Kum96].

Skinny triangles are undesirable for smooth shading. Using simple heuristics, we are able to obtain fat triangles in practice. Furthermore, our algorithm directly produces OpenGL compliant triangle strips, which can be displayed more efficiently than a list of triangles. We avoid triangle fans: we keep the degree of a vertex in the triangulation small in order to lower the cost of incremental updates to the triangulation. Steiner points are not added on the edge of the trimming polygon, since that can lead to cracks in the tessellation.

Note that trimming curves are non self-intersecting. However, a sparse sampling of the trimming curve can still produce self-intersecting polygons. While uncommon, such non-simple polygons do occur and can cause most triangulation algorithms to fail. Instead of assuming that the input polygon is simple, which can require very high sampling density, we check for such cases at little additional cost and thence make the polygon simple by increasing the sampling *only* when necessary.

In our application domain, a significant number of

instances of the algorithm are executed every frame. The size of each instance is relatively small. (We ran several experiments to characterize sampling sizes in typical surface-model walkthroughs. Thousands of

surfaces are re-triangulated per frame. The number of surface samples average about 21, the number of curve samples average about 45 and the number of points in a non-empty cell is around 2.5.) Hence, in addition to simplicity and robustness, the constants of complexity of the triangulation algorithm are crucial to interactive performance. The space requirement per surface is not quite critical, but it is prohibitive to maintain a large data structure for each surface across all frames. To speed up point location operations (required for incremental updates), we use the natural partitioning provided by the grid sampling of the surfaces domain and discard all other auxiliary data structures, which are re-constructed every time they





Figure 2: Polygon-Tracing (PSLG Segments inside the polygon are indicated by dots. Hatched shading indicate the side of the polygon that is to be triangulated. Dotted shading indicates quads.)

are needed. Before we describe the construction of this data structure, we introduce some notation. Without

3. POLYGON-TRACING

The grid provides a natural partition of the domain. We trace the trimming polygon, i.e. process p_i s in polygon order, assigning to each point p_i the grid cell it lies in, $cell(p_i)$. The tracing step generates a set of *quads* (rectangles) and a PSLG that we must then triangulate. During tracing, we also construct auxiliary data useful for triangulation. In addition, we store all intersections of the polygon with u -lines. For each cell, we keep a list of intersections of the polygon with its left boundary in increasing order of their (unique) v -values. We discard all degenerate intersections, i.e. if a polygon segment is collinear with a u -line or a v -line, we remove all corresponding grid points from consideration. We store the following information during polygon-tracing:

For each intersection, I , of the polygon segment, p_i - p_{i+1} , with a u -line, u_j :

Mark I as **MIN**, if $u(p_i) > u(p_{i+1})$. Mark as **MAX** otherwise. ($u(p_i) \neq u(p_{i+1})$ as degenerate intersections are not allowed — only grid points and lines that lie strictly inside the trimming polygon are included in the PSLG.)

loss of generality, assume that the curve is specified clock-wise and the part of the domain to be triangulated is enclosed by the trimming curve.

Definitons: The polygon corresponding to the tessellated trimming curve consists of points p_i (see Fig. 1). The surface grid points are denoted by g_{ij} . We also denote a vertical grid line (a sequence of sample points) by u_i and horizontal lines by v_j . i and j range from 0 to the corresponding sampling sizes. We also refer to the u and v coordinates of a point, p , on the domain as its u -value or $u(p)$ and v -value or $v(p)$, respectively. Each rectangle formed by four adjacent grid points is called a *cell*. The part of the domain between two consecutive grid lines is called a *strip*.

Store $v(I)$, the v -value of the intersection. (Keep multiple intersection with a cell boundary sorted by v -value.)

Store $Maximum(v(I))$ and $minimum(v(I))$ attained by the polygon in the u -strip containing p_i . These bound the quads (rectangles on the domain, Fig. 2).

Store a pointer to p_{i+1} . For each cell's left and right u -boundaries we maintain a linked list of all intersections with that boundary. Note that in the worst case a cell could have $O(n)$ intersections with a polygon with n vertices. However, since the cells and polygons follow the same sampling rule and highly winding trimming curves are tough to generate, the number of intersections of most cell boundaries is small, if not 0 or 1.

If the entire trimming polygon lies within the same cell, no intersections are detected. This case does occur in practice, especially for surfaces with small area on-screen or of degree 1×1 . It means the polygon is contained in one cell. Such (low detail) surfaces need only few samples, hence the number of points on the polygon is small as well. If that is not the case, additional grid lines may be included solely for polygon-tracing, thus ensuring that the number of points in a cell remains small.

We generate the quads and the PSLG using what amounts to a modified sweep line algorithm. We find **MIN-MAX** pairs on each u -line and a corresponding **MIN-MAX** pair on its adjacent u -line. (see Fig. 2(a)). The pairs on each line are available in the sorted order. Note that adjacent **MIN-MIN** or **MAX-MAX** pairs indicate non-simple polygon. The matching of pairs on adjacent line (to obtain a strip of quads) is as simple as matching the i^{th} pair on both lines, except the two cases shown in Figs. 2(b) and 2(c), when polygon chain turns back to intersect the current line

4. PSLG TRIANGULATION

Our triangulation scheme is based on trapezoidation [ZC99,Sei91]. The basic idea of this technique is demonstrated in Figure 3. Trapezoidation of a PSLG (shown in thick solid lines) is obtained by drawing horizontal rays (i.e. dashed lines parallel to the u axis) at each vertex of the graph limited in both directions by the first segment (or vertex) the ray intersects. The PSLG segments and horizontal lines form a set of trapezoids. The diagonals (shown in thin solid lines) of the trapezoids that connect two vertices of the PSLG partition the PSLG into a set of uni-monotone polygons. Uni-monotone polygons consist of a single v -monotone chain and another line-segment. For a discussion and proofs, we refer the reader to [Sei91,FM84]. It can be shown (we omit the proof here) that the line-segments mentioned above, call them monotone segments, are all PSLG segments, and thus small in length for our application. We will exploit this fact while triangulating these monotone polygons.

MONOTONE TRIANGULATION

Simple $O(n)$ algorithms for triangulating monotone polygons have been proposed [GJPT78,FM84] and implemented [NM95]. However, all of them tend to produce triangle fans and skinny triangles, both undesirable properties for our purpose. We propose another algorithm, equally efficient in practice that produces better triangle strips. Our approach is motivated partly by [HM83] and [RR94]. Designed for uni-monotone polygons, our algorithm is much simpler and more efficient. We use a u -tree: a u -tree

instead of the adjacent line. For case (b), we discard (for matching purposes only) the extra pair on the u -line u_i . For case (c), we insert an extra **MIN-MAX** pair on u_i (for matching).

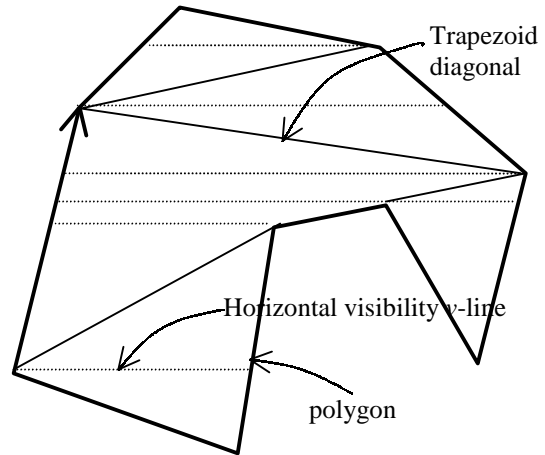

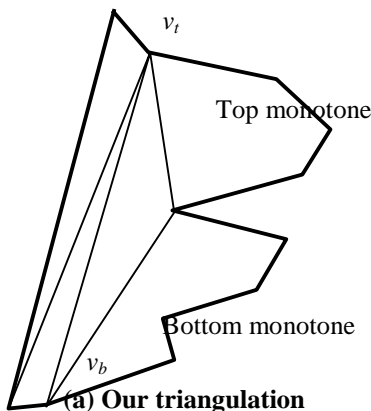


Figure 3: Trapezoidation

maintains all the local u -minima and can be constructed in $O(n)$. The invariant for a u -tree node is as follows: it stores the vertex with the minimum u -value of all its children. All vertices above it (i.e. with higher v -value) are kept in its left sub-tree and all vertices below it are kept in the right sub-tree. Fig. 4(a) illustrates the basic idea of tree construction. A u -tree can be constructed incrementally in a single pass over the n vertices of the polygon in $O(n) + O(k^2)$, where k is the number of local u -minima. Processing the minima in a random order reduces the expected cost to $O(n + k \log k)$. However, k is small in practice (around than 3-4) and processing the vertices in the polygon order is usually sufficient.

Once the u -tree is constructed, we produce the triangle strips in $O(n)$ time as follows:


- Maintain pointers to the current root v_m , current top, v_t , and bottom, v_b , vertices of the chain
- While v_t and v_b both have lower u -value than v_m does:
 - Add the one with lower u -value, say v_t , to the strip, Replace v_t by the next vertex on the polygon




(a) Our triangulation (b) Inevitable Fan (c) Unlikely Skinny Triangles
Figure 4: Triangulating Monotone Polygons

- Otherwise,
 - Add v_t , v_m and v_b to the current strip and output it.
 - Diagonals $v_t - v_m$ and $v_b - v_m$ subdivide the polygon into two v -monotone polygons, the left sub-tree of the current u -tree corresponds to the top polygon and the right sub-tree corresponds to the bottom polygon. Proceed recursively (Fig. 4(a)).


Note that the procedure above uses a diagonal between v_t and v_b only if both lie to the left of the minimum u -valued inflection vertex between them and thus are visible to each other. Due to this advancing front like technique, high degree triangles are less likely to occur. Further, $u(v_t) - u(v_b)$ is small, as the corresponding monotone segment is short (see Fig. 4(c) for an example). Hence skinny triangles of the kind shown in Fig. 4(c) do not occur. However, skinny triangles can be generated due to horizontal trapezoidation, if two trapezoids vertically adjacent to each other are both skinny (see Fig. 6(a)). This is a general shortcoming of our scheme since we avoid skinny triangles only during the second phase: monotone polygon triangulation. While it may be possible to devise an algorithm not based on trapezoidation, we have found the trapezoidation scheme to be very robust. It fails only if the input polygon is non-simple or almost non-simple. Hence, it is more appropriate to implement special cases for such (rare) skinny monotone polygons generated by trapezoidation. In practice, we avoid diagonalizing trapezoids with small height and thus obtain more than a single u -monotone chain. An extension of the monotone triangulation algorithm described above works for this case. Note also that sometimes fans are inevitable as shown in Fig. 4(b) – no other triangulations exist.



(a: Case I)



(b: Case II)



(c)

Figure 5: Simplicity Verification

SIMPLICITY DETECTION

We call a polygon (or a PSLG) simple if no two edges intersect, except possibly at a common vertex. Although it is not common for a polygon input to the triangulation algorithm to be non-simple, if left

undetected, it can cause the triangulation to fail and display to become invalid. However, we pose the simplicity detection in terms of horizontal visibility lines used in the trapezoidation. Subsequently, a minor modification to the trapezoidation algorithm helps us to detect if edges p_i-p_{i+1} and p_j-p_{j+1} of the polygon intersect. If they do, we compute extra samples on the curve between parameter values $t(p_i)$ and $t(p_{i+1})$ and between $t(p_{i+1})$ and $t(p_{j+1})$ and retry. Since polygons are rarely non-simple, the extra cost of re-sampling and iterating is acceptable. Simplicity check is straightforward after realizing that the horizontal line corresponding to some vertex of a non-simple polygon is *inconsistent*. We define inconsistency as follows:

Definitions: A point q is *visible* to point p if the line segment pq does not intersect the given PSLG. Segment s is visible to p if a point on s is visible from p . Point p is on the *interior* with respect to an oriented segment p_i-p_{i+1} visible to p , if it lies on the right hand side of p_i-p_{i+1} . Also on the interior are all points r on the line segment pq , where q is the point on p_i-p_{i+1} such that $v(q) = v(p)$. A point is *inconsistent* if it is on the interior with respect to some edge of the polygon while on the exterior with respect to another. If a point is not *inconsistent* we call it *consistent*. Clearly, each non-simple polygon has inconsistent points on its boundary. In fact:

Theorem: one of the polygon vertices, p_i , must be inconsistent, if the polygon is non-simple.

Proof: Note that a vertex is inconsistent if any point on the horizontal line through it is inconsistent. To sketch a proof, consider two edges p_i-p_{i+1} and p_j-p_{j+1} that intersect. There are four possible cases to consider. Two are shown in Figs. 5(a) and 5(b). The other two are symmetric. Locally interior points are shown shaded for each edge. The black dots show

examples of inconsistent points.

CASE I: If either p_j or p_{j+1} is horizontally visible from the other edge, we have found an inconsistent vertex. Otherwise, there must exist occluding edges and vertices. The vertex with the minimum v -value greater than $v(p)$, p being the point of intersection, must be

visible to both $p_i p_{i+1}$ and $p_j p_{j+1}$ and is hence inconsistent.

CASE II: Consider the segment s , that is horizontally visible from p . If both the end points of s are visible from p , one of these must be inconsistent. Otherwise, there must exist minima and maxima vertices p_k and p_l on either side of the horizontal visibility line through p . (One of these points may lie on s .) If s is oriented upwards, p_k is inconsistent, otherwise, p_l is inconsistent. Strictly speaking, the argument above holds only if any given segment has only one intersection with the rest of the polygon and if no segments are horizontal. However, by using **5. INCREMENTAL UPDATE**

Before we explain the incremental update, we need to consider the data structure for triangulation. Since the number of surface patches can be quite large, we keep the size of cache per patch small. In addition to the triangle strips, for each v -strip on the domain, we maintain a list of triangle strips intersecting that v -strip, i.e., we keep a pointer to the first vertex of each entering triangle strip. In addition, if one of the strip-edges is also an edge of the trimming polygon, we mark it so. This structure is similar to the one proposed in [MSZ96], which allows point locations, in expected $O(n^{1/3})$ time. However, we do not maintain adjacencies between strips and hence cannot “walk” from a random triangle. We directly point to the appropriate strips. It is possible for us to retain pointers between adjacent strips as well, however since the number of strips crossing a cell is typically small, we do not lose much performance by explicitly searching for adjacent strips in any cell.

The types of updates to our triangulation is limited for our application (see [KM95,KML96] for details):

- i) A segment $p_i p_{i+1}$ may be replaced by $p_i q$ and $q p_{i+1}$ (and vice versa), call these update segments.
- ii) A grid line u_i (or v_i) could be added or deleted: call it the update line.

transitivity, the first restriction may be removed and by rotating the input (or topologically sorting it), the second restriction may be removed.

Thus, simplicity detection can be performed while constructing the horizontal visibility line for trapezoidation. The only remaining operation is to verify that the two ends of the visibility line are both in the exterior or both in the interior. (As a special case, for the minimum and maximum v -valued vertices, if the visibility line is locally in the interior, there must exist other segments visible along this visibility line. Fig. 5(c) shows an example with the corresponding smooth curve overlaid.)

The incremental triangulation has the following main steps:

- Re-trace to compute new PSLG
- Delete triangles intersected by the update features (segment or line), thus creating a hole
- Re-triangulate the hole

The first two steps need some explanation. For each update input, we first need to determine if any quads are added or deleted. Thus, we need to re-trace the polygon. However, we need only re-trace the new segments in case (i). No other segments may change. In case (ii), we must retrace all segments intersected by the update line. we can determine the polygon edges that bound the edges intersected by the line. We only need to re-trace between these pairs of bounds.

Any quads introduced by this re-tracing are diagonalized. A quad intersected by the new feature is deleted, if it contains an intersection with the trimming polygon, otherwise it is bisected. Similarly if the deletion of a feature results in the deletion of an intersection, a new quad is generated or two quads are merged. The quads added or deleted in the previous step are also included in the set of update features. Parts of a grid line that result in a quad bisection are deleted.

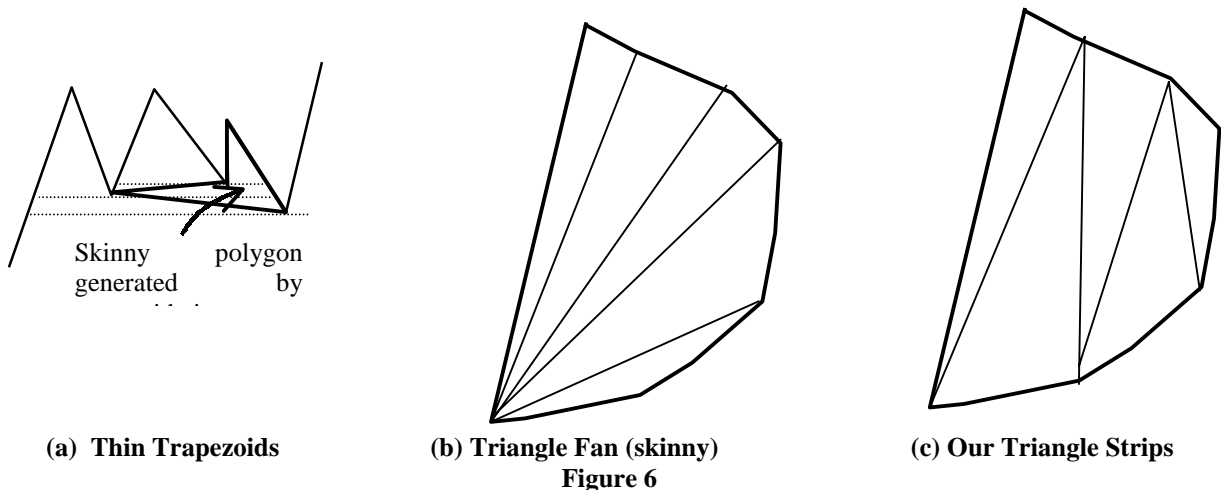


Figure 6

Table 1: Performance of our Triangulation algorithm on an SGI Octane with 195MHz R10k. (Statistics are per frame averaged over more than 5,000 frames. Time is in milliseconds, angle is in degrees.)

Model	Num. Patches	Triangulation time ¹	Minimum angle ¹	Avg. Degree ¹	Triang. time ^{2,3}	Triang. Time ^{2,4}	Min. angle ²	Avg. Degree ²	Tot. frame time ²
Brakehub	560	32	10.5	9.2	33	14	33.5	1.8	24
Torpedo	1201	59	16.1	11.1	62	28	34.2	2.1	42
Pivot	4101	61	14.9	10.8	68	38	31.1	2.2	66
TorpRoom	17032	131	8.4	14.4	143	81	25.3	3.9	110

1. [NM95] implementation: No simplicity check, skinny polygons, generated triangles
2. Includes simplicity check, fat polygons, triangle strip generation
3. First instance of the triangulation
4. Subsequent updates

To locate an update feature in the current triangulation, we start randomly in each triangle-strip crossing the corresponding *v-strip* and walk across triangle boundaries until all triangles adjacent to the feature are located. A strip intersected by an update

6. IMPLEMENTATION AND PERFORMANCE

We augmented the publicly available implementation of Seidel’s triangulation algorithm [NM95] with our algorithm and plugged it into a Bezier surface rendering system. We performed a suite of tests on a variety of Bezier surfaces on a Silicon Graphics Indigo 2 system with Maximum Impact graphics. We compared our results with those of [NM95] by plugging that in. [NM95] uses the algorithm by [FM84] for monotone triangulation. In practice, our implementation results in less than 10% slow-down in one-time triangulation. (Here we compare the times for the first time triangulation performing no incremental triangulation.) For the cost of this slowdown, our system generates more uniform sided triangles, and, more importantly, it never fails – the

7. CONCLUSION

We have presented a simple, robust, efficient and incremental algorithm for triangulating points on a surface. This includes both the generation of the polygons and their triangulation. To contrast our method of polygon generation with that of [KM95], that approach is cell based and attempts to find polygons around each cell. In spite of the higher overhead of this search, [KM95]’s method results in similar sized polygons. Furthermore, we do not need to handle a large number of special cases, nor do we need the clean-up stage, which can be quite slow. In addition to being more efficient, our algorithm directly produces triangle strips and generates better quality triangles. However, due to strip splitting procedure, the strips tend to become fragmented after a while. Currently we perform complete re-triangulation periodically. A slightly more complicated approach could avoid splitting strips by adding extra vertices in

feature is split into two or more (disjoint) pieces. Each intersected triangle is deleted from the strip. We re-triangulate this hole (which is again a PSLG) using the algorithm described in Section 4.

code does not crash and the result is correct even for non-simple curve approximations. Our triangulation is of significantly better quality: the smallest angles goes up from 11.3 degrees to 33.7 degrees on average. The degree of our triangulation is better as well (see Table 1). Furthermore, we generate triangle strips obtaining a rendering speed-up of more than 60% over triangles. A surface-patch needs re-triangulation in less than 20% of the frames on average, and is just re-rendered 80% of the time. The speed-up in triangulation obtained by using an incremental technique over one triangulating from scratch every time sampling changes is almost 90%. Also, note that our implementation is more robust due to the simplicity check. We have performed millions of triangulations without any failure using our implementation.

the middle. Our algorithm also verifies that a polygon is simple at little additional cost. While infrequent, if left undetected, a non-simple polygon can cause a system to crash. While it is possible to extend our incremental algorithm to perform constrained Delaunay triangulation, we believe the cost does not justify the minor benefit. One limitation of our system is its independence from the surface parameterization. We only guarantee triangle quality in the domain. Thus, our triangulations may still be skinny for severely skewed parameterizations. In our experience, most surface models do not suffer from such skews. Our algorithm works well in the common cases.

REFERENCES

- [AES94] S. Abi-Ezzi and S. Subramaniam, Fast Dynamic Tessellation of Trimmed NURBS Surfaces, *Computer Graphics Forum*, 13(3), 107—126, 1994. (Eurographics '94.)
- [BDE92] M. Bern, D. Dobkin and D. Eppstein. Triangulating polygons without large angles. *Proc. 8th Annual ACM Symp. Comput. Geom.*, pages 222—231, 1992.
- [Cha91] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6:485—524, 1991.
- [EAV99] J. El-Sana, J. E. Azanli and A. Varshney. Skip Strips: maintaining triangle strips for view-dependent rendering. *Proc. Visualization*, 1999.
- [FM84] A. Fournier and D. Montuno. Triangulating simple polygons and equivalent problems. *ACM Transactions on Graphics*, 3:153—174, 1984.
- [GJPT78] M. R. Garey, D. S. Johnson, F. P. Preparata, and R. E. Tarjan. Triangulating a simple polygon. *Inform. Process. Lett.*, 7:175—179, 1978.
- [HM83] S. Hertel and K. Mehlhorn. Fast triangulation of simple polygons. *Proc. 4th Internat. Conf. Found. Comput. Theory*, 207—218, 1983.
- [KM95] S. Kumar and D. Manocha. Efficient rendering of trimmed NURBS surfaces. *Computer-Aided Design*, 27(7): 509—521, 1995.
- [KML96] S. Kumar, D. Manocha, and A. Lastra. Interactive display of large NURBS models. *IEEE Transactions on Visualization and Computer Graphics*, 2(4): 323—336, 1996.
- [Kum96] S. Kumar. Interactive Display of Parametric Spline Surfaces. *Ph.D. Thesis*, University of North Carolina, 1996.
- [KU99] E. Kranakis, and J. Urrutia. Isomorphic triangulations with small number of Steiner points. *International Journal of Computational Geometry & Applications*, 9(2):171—180, 1999.
- [Mit93] S. Mitchell. Refining a triangulation of a planar straight-line graph to eliminate large angles. *Proc. 34th Annual IEEE Symposium on Foundation of Computer Science (FOCS)*. Pages 583—591, 1993.
- [MSZ96] E. Mücke, I. Saias, and B. Zhu. Fast randomized point location without preprocessing in two- and three-dimensional Delaunay triangulations. In *Proc. 12th Annual ACM Symp. Comput. Geom.*, pages 274—283, 1996.
- [NM95] A. Narkhede and D. Manocha. Fast polygon triangulation based on Seidel's algorithm. In A. Paeth, editor, *Graphics Gems V*, Academic Press, 1995.
- [PR95] L. Piegl and A. Richard. Tessellating trimmed NURBS surfaces. *Computer Aided Geometric Design*, 27(1):16—26, 1995.
- [RHD89] A. Rockwood, K. Heaton, and T. Davis. Real-time rendering of trimmed surfaces. *ACM Computer Graphics*, 23(3):107—117, 1989. (SIGGRAPH Proceedings).
- [RR94] R. Ronfard and J. Rossignac. Triangulating multiply-connected polygons: A simple, yet efficient algorithm. *Comput. Graphics Forum*, 13(3):C281—C292, 1994.
- [Sei91] R. Seidel. A simple and fast randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Computational Geometry Theory & Applications*, 1(1):51—64, 1991.
- [ZC99] B. Zalik and G. Clapworthy. A universal trapezoidation algorithm for planar polygons. *Computers & Graphics*, 23(3): 353—363, 1999.