

V. Surazhsky · C. Gotsman

## High quality compatible triangulations

Received: 18 September 2002 / Accepted: 3 October 2003 / Published online: 28 April 2004  
© Springer-Verlag London Limited 2004

**Abstract** Compatible meshes are isomorphic meshings of the interiors of two polygons having a correspondence between their vertices. Compatible meshing may be used for constructing sweeps, suitable for finite element analysis, between two base polygons. They may also be used for meshing a given sequence of polygons forming a sweep. We present a method to compute compatible triangulations of planar polygons, sometimes requiring extra (Steiner) vertices. Experimental results show that for typical real-life inputs, the number of Steiner vertices introduced is very small. However, having a small number of Steiner vertices, these compatible triangulations are usually not of high quality, i.e. they do not have well-shaped triangles. We show how to increase the quality of these triangulations by adding Steiner vertices in a compatible manner, using remeshing and mesh smoothing techniques. The total scheme results in high-quality compatible meshes with a small number of triangles. These meshes may then be morphed to obtain the intermediate triangulated sections of a sweep, if needed.

**Keywords** Compatible triangulations · Low discrepancy point sets · Morphing · Remeshing · Sweep meshes

### 1 Introduction

In CAE, swept volumes, sometimes called two and one half dimensional volumes, are frequently constructed between two base polygons given with a correspondence between their vertices. To discretize a swept volume for finite element analysis, it is necessary to mesh the interiors of the sequence of polygonal cross-sections forming the sweep, usually introducing interior (Steiner) vertices, in a manner such that the mesh is isomorphic, valid and well-shaped within all the polygons. This mesh is said to be compatible with all the polygons. See Fig. 1 for an example. The result is a set of prisms defining the sweep, whose edges are the so-called “ribs” of the sweep [1].


In the case where only the two base polygons of the sweep are given, it is possible to automatically generate the intermediate polygons by a process known as morphing. The morphing problem, in general, is to smoothly transform one given polygon, the source, into another, the target, over time. Constructing the sweep volume may be considered a morphing problem by thinking of the sweep axis as the time axis of the morph. Morphing has been the subject of much research over recent years and has wide practical use in areas such as computer graphics, animation and modelling.

The naïve approach to the morphing problem is to decide that the polygon vertex trajectories are straight lines, where every feature of the shape travels with a constant velocity along the line towards the corresponding feature of the target during the morph. However, this simple approach can lead to undesirable results. The intermediate shapes can vanish, i.e., degenerate into a single point, or self-intersect even though the source and target are simple. Even if the linear morph is free of self-intersections and degeneracies, its intermediate shapes may have areas or distances between features far from those of the source and target, resulting in a “misbehaved” morph. See the top row of Fig. 2. Most of the research on solving the trajectory problem for morphing concentrates on trying to eliminate self-

---

V. Surazhsky (✉)  
Department of Informatics,  
Oslo University,  
Oslo, Norway  
E-mail: vitus@ifi.uio.no

C. Gotsman  
Center for Graphics and Geometric Computing,  
Department of Computer Science,  
Technion—Israel Institute of Technology,  
32000 Haifa, Israel




**Fig. 1a–e** The concept of compatible triangulations of corresponding polygons. Vertex correspondence is denoted by digits. **a, b** Non-compatible triangulation; **c, d** compatible triangulation; **e** the sweep with bases (c) and (d)

intersections and preserving the geometrical properties of the intermediate shapes. Numerous existing methods achieve good results for many inputs, (e.g. [2,3,4]), yet, only the methods that use compatible triangulations are able to guarantee any properties of the resulting morph.

In order to perform finite element analysis on a swept volume—a sequence of corresponding simple polygonal cross sections—it is necessary to mesh the polygon interiors in a compatible manner. In this work we concentrate on compatible triangulations. Compatible meshing is not always possible unless Steiner vertices are introduced into the interior of the polygons. The main challenge is then to minimize the number of Steiner vertices to the least needed to achieve compatibility. Unfortunately, this can be as much as  $\Omega(n^2)$ , where  $n$  is the number of vertices of the polygons. In the first work

**Fig. 2** Morphing of two corresponding polygons. *Leftmost* polygon is source and *rightmost* is target. Note that the correspondence implies some rotation during the morph. *Top row*: the linear morph resulting in self-intersecting intermediate polygons. *Bottom row*: the morph of the polygons generated by embedding the source and target into compatible triangulations and applying the method of [12] or [13] guarantees that the intermediate polygons are also simple



on this problem, Aronov et al. [5] provided two constructions that result in quite a large number of Steiner vertices. In their work on polygon morphing, Gotsman and Surazhsky [6] improved Aronov et al.’s so-called “spiderweb” method to significantly reduce the number of Steiner vertices required. Kranakis and Urrutia [7] presented a completely different method in which the number of Steiner vertices introduced depends on the number of reflex vertices of the two polygons. Gupta and Wenger [8] described an algorithm that uses minimal-link polylines in the polygon.

While compatible triangulations of polygons with a very small number of Steiner vertices are definitely an advantage from a complexity point-of-view, it appears that these triangulations are naturally not well-shaped. They tend to contain long skinny triangles that cannot be adjusted to improve the triangle shape significantly. Hence a major challenge in our application is to introduce as small a number of Steiner vertices as possible, yet obtain two triangulations with decent quality and maintain compatibility of the triangulations throughout the process. We call this process compatible remeshing. This was attempted in the work of Alexa et al. [9], who start off with compatible triangulations of polygons and introduce Steiner vertices in order to improve the quality of the triangulation. They, however, start from a large number of Steiner vertices and thereafter increase this number significantly, in order to achieve triangulation of good quality. This results in compatible triangulations that are overly complex.

The main contribution of this paper is a method to compatibly triangulate two planar polygons using Steiner vertices. Our experiments show that for typical real-life inputs the number of Steiner points introduced is very small. The drawback of our method is that we cannot prove at this time reasonable bounds on the time complexity of the algorithm and on the number of Steiner vertices. Another important contribution of this paper is a new area-based remeshing technique, which dramatically improves the spatial distribution of the vertices over the total mesh area. In this paper we also show how to modify and combine several remeshing techniques to produce compatible triangulations of high quality.

---

## 2 Sweep generation

Meshing for sweep generation has been attempted before by a variety of authors in the meshing community (e.g. [1,10,11]). Their basic approach is to generate a mesh for a subset of the cross-section polygons, usually just one of the sweep bases, and then project this mesh somehow onto the other polygons. Beyond the fact that this certainly does not guarantee that the result will be a valid triangulation, there are also no guarantees for the quality of the triangulation even if it were valid. Our solution, taking into account both sweep bases (and theoretically all intermediate polygons), solves all these problems.

If only the base polygons of the sweep are given, the intermediate polygons, with their corresponding compatible triangulations, may be generated using the morphing methods of Gotsman and Surazhsky [6]. This is done by reducing the problem to that of morphing compatible planar triangulations with a common convex boundary, in which the polygon is embedded, as described by Floater and Gotsman [12] and Surazhsky and Gotsman [13]. Two corresponding point sets admit a compatible triangulation if there exists a triangulation of one point set which, when induced on the second point set by the correspondence, is a legal triangulation there too. The morphing method of Floater and Gotsman [12] is based on the convex representation of triangulations using barycentric coordinates, first introduced by Tutte [14] for graph drawing purposes and later generalized by Floater [15] for parameterization of 3D meshes. This avoids many of the problems associated with morphing and basically guarantees that the triangulation remains valid (i.e. compatible with the source and target) throughout the morph.

To embed the two polygons in a triangulation, first compatibly triangulate the polygon interiors. Then circumscribe the two polygons in a common convex enclosure and compatibly triangulate the two resulting annuli between the polygons and the enclosure [16] (possibly requiring Steiner vertices). This results in two compatible triangulations with a common convex boundary, in which the polygons are embedded. Morphing these triangulations using the methods of [12] or [13] will then result in a valid (compatible) morph of the two polygons. See the bottom row of Fig. 2.

---

## 3 Compatible triangulations

### 3.1 Previous work

As already stated, Kranakis and Urrutia [7] presented two different methods to compatibly triangulate two polygons in which the number of Steiner vertices introduced depends on the number of the polygons' reflex vertices. The first algorithm produces a rather large number,  $O((k+l)^2)$ , of Steiner vertices, where  $k$  and  $l$  are the number of the two polygons' reflex vertices respectively.

The second algorithm introduces, at most,  $O(kl)$  Steiner vertices, but its drawback is that it may add Steiner vertices on the polygon boundaries, which some applications do not allow. Furthermore, enlarging the boundary might prevent this algorithm from being used as a black box in a recursive manner, as the algorithm might not terminate.

Gupta and Wenger [8] described an algorithm, in theory the best so far, which constructs the compatible triangulation based on minimum link paths inside the polygons  $P$  and  $Q$ . A "minimum" link path is a polyline of the minimal number of straight-line segments connecting two given polygon vertices and lying entirely in the polygon interior. The idea behind their algorithm is the following. First, compute an arbitrary triangulation  $T_p$  of  $P$ . Using edges of  $T_p$  it is possible to partition  $P$  into sub-polygons such that the number of links in minimum link paths in those sub-polygons is no more than a small constant (e.g. 5). Then the corresponding partition of  $Q$  is constructed using non-intersecting minimum-link polylines. The vertices of these polylines are the Steiner vertices of the triangulations. These corresponding sub-polygons are then compatibly triangulated, usually requiring a relatively small number of Steiner vertices owing to the properties of the partition.

The resulting compatible triangulations have  $O(M \log n + n \log^2 n)$  triangles, where  $M$  is the number of triangles in the optimal solution. Theoretically, this is good, except that the constant factor is quite large (approximately 40, according to the authors), so it is not very practical for smaller inputs. Moreover, the complexity of compatible triangulations is measured in the number of triangles instead of the number of Steiner vertices. Thus, for example, if the optimal compatible triangulations have no Steiner vertices or only a small (constant) number of Steiner vertices, even a constant factor of the number of triangles of the optimal triangulations will allow adding  $O(n)$  Steiner vertices. This means that the algorithm will result in a large number of Steiner vertices for most real-life polygons, for which several Steiner vertices are usually sufficient.

Another drawback of this algorithm is that it is not symmetric in  $P$  and  $Q$ . The choice of the triangulation  $T_p$  of  $P$  strongly influences the resulting compatible triangulations and the number of Steiner vertices. From a practical point of view, the algorithm involves implementing many state-of-the-art computational geometry algorithms developed over the last two decades. As a consequence, an implementation of the algorithm is currently not available and thus, it is impossible to compare this algorithm with other algorithms generating compatible triangulations.

### 3.2 Our algorithm

Our algorithm is similar in spirit to that of Gupta and Wenger [8], namely, it is based on the idea of partitioning polygons using minimum-link polylines. However, our algorithm is much simpler. Given two


polygons  $P$  and  $Q$  with a correspondence between their vertices, we find a pair of vertices  $u$  and  $v$  with a minimal-link polyline between them in one of the polygons and a corresponding polyline in the second. After the shorter polyline is refined to the same number of vertices as the longer one, the two polylines compatibly partition both of the input polygons into two sub-polygons. The vertices of these polygons are the Steiner vertices. We then apply the algorithm recursively on these two sub-polygons. The process terminates when the input polygons contain only three vertices, namely, the polygons have become triangles.

We still need to show how to find a pair of vertices  $u$  and  $v$  that minimizes the number of links in the partitioning polylines. To achieve this, we employ the method of Suri [17], who showed how to find the minimum-link path between two given vertices in a simple polygon in  $O(n)$  time, where  $n$  is the number of polygon vertices. In a subsequent work, Suri [18] showed how a simple polygon can be preprocessed in  $O(n)$  time in order to query the number of links of the minimum-link path between two given vertices of the polygon in  $O(\log n)$  time. Thus, we can query all possible vertex pairs of the polygon in  $O(n^2 \log n)$  time using this algorithm. We believe that the time complexity for finding the optimal vertex pair(s) of the polygon(s) can be further improved to  $O(n \log n)$  by exploiting the existing preprocessed data structure for the queries, instead of using the query procedure for a specific pair as a black box.

Hence, in this manner we may determine which pair is best to use and then employ the first algorithm to actually compute the paths. Accordingly, in order to find the best path for both polygons we query the two polygons for the minimum-link distance and choose the pair that has the best (minimal) value of the maximum between two distances. Namely, we choose the pair  $(u, v)$  which satisfies:

$$(u, v) = \arg \min_{u, v \in P} \max \{ \text{dist}_P(u, v), \text{dist}_Q(u, v) \} \quad (1)$$

**Fig. 3** Compatible triangulations of two polygons (one on each row), whose vertex correspondence is denoted by digits. The light grey region denotes the current polygon during the recursion. The grey regions are regions already triangulated. Thick segments are minimal-link polylines, which recursively partition the polygons




In practice, this pair is not unique. Therefore, we choose the pair that will partition the polygons into sub-polygons that are as balanced as possible, in order to reduce the overall algorithm complexity. This can be easily done by comparing the indices of the polygon vertices. More formally, if the polygon vertices are  $v_1, \dots, v_n$  and  $m$  is the size of the polygon, we look for:

$$(v_i, v_j) = \arg \max_{v_i, v_j \in P, j} \min \{ j - i, n + i - j \} \quad (2)$$

Note that if it is possible to compatibly triangulate the two polygons without any Steiner vertices, our algorithm will do so in most cases, as opposed to the other algorithms. Since this is the case for many inputs, our algorithm has a significant advantage. However, because of its greedy nature, our algorithm in some rare cases may introduce a few Steiner vertices, when compatible triangulations without Steiner vertices do exist.

We must still show that the algorithm terminates, because when the polygon  $P$  is partitioned into two sub-polygons  $P_1$  and  $P_2$ , theoretically the size of  $P_1$  or  $P_2$  (or both) can be identical to that of  $P$  and, if this repeats, the algorithm can run indefinitely. In general, to prevent such cases we should check that the size of the partitioned polygon stays the same as the size of  $P$  only once during the run of the algorithm. If the size stays the same after two iterations, the algorithm should backtrack and choose another vertex pair for the partition polylines. This, theoretically, may result in exponential time complexity. However, in practice (we have tested the algorithm over numerous, very complex inputs), even the case when the size of the polygons repeats itself twice does not occur. Thus, although we cannot prove it at this time, we believe that for most real-life inputs backtracking is not required and the algorithm complexity is polynomial. By analyzing the algorithm complexity for this case, we achieve a total time complexity of  $O(n^4 \log n)$  for the worst case when  $\Omega(n^2)$  Steiner vertices have been introduced. When no Steiner vertices are needed, the algorithm complexity is  $O(n^3 \log n)$ . However, on the average, the recursive partitioning of the polygon is balanced because of Eq. 2. Thus, the average complexity will be  $O(n^3 \log^2 n)$  and  $O(n^2 \log^2 n)$  for  $\Omega(n^2)$  and zero Steiner vertices respectively.

See Fig. 3 for an illustration of the various stages of the compatible triangulation algorithm.



**Fig. 4** Weighted angle-based smoothing:  $c_i$ s obtained by rotation of  $c$  around  $p_i$  to coincide with the bisector of  $\alpha_i$

## 4 Mesh improvement

Our compatible triangulation algorithm generates a small number of Steiner vertices, at locations which have not necessarily been optimized for mesh quality. It is possible to improve these meshes by smoothing them (moving the vertices), or remeshing them (changing the connectivity). In this section we describe methods for these two operations, which we believe are also of independent interest.

### 4.1 Weighted angle-based smoothing

Zhou and Shimada [19] presented an effective and easy-to-implement angle-based mesh smoothing scheme. They show that the quality of the mesh after angle-based smoothing is much better than after Laplacian smoothing. Moreover, the chance that the scheme will produce inverted (invalid) faces is much less than that in Laplacian smoothing. Unfortunately, this is true mostly for meshes whose vertices have degrees close to the average degree, namely, the mesh connectivity is close to regular. When the mesh has more irregular connectivity, the scheme may fail. In applications involving meshes with very distorted (long and skinny) triangles, a more robust smoothing scheme is critical. We propose a very simple improvement to the original angle-smoothing scheme, which significantly reduces the chances of inverted triangles and improves the quality of the resulting mesh. Furthermore, it has almost the same computational cost per iteration and a lower total computational cost due to better convergence in practice.

The original scheme attempts to make each pair of adjacent angles equal. Given a vertex  $c$  and its neighbours  $p_1, \dots, p_k$ , where  $k$  is the vertex degree, we want to

move  $c$  in order to improve the angles of the triangles incident on  $c$ . Let  $\alpha_i$  be the angle adjacent to  $p_i$  in the polygon  $p_1, \dots, p_k$ . We define  $c_i$  to be the point lying on the bisector of  $\alpha_i$  such that  $\|p_i - c_i\| = \|p_i - c\|$ , namely, the edge  $(p_i, c)$  is rotated around  $p_i$  to coincide with the bisector of  $\alpha_i$  (see Fig. 4). The new position of  $c$  is defined as the average of all  $c_i$  for all the neighbours, namely:

$$c_{\text{new}} = \frac{1}{k} \sum_{i=1}^k c_i \quad (3)$$

We improve this scheme by introducing weights into Eq. 3. For a small angle  $\alpha_i$  it is difficult to guarantee that the resulting  $c_{\text{new}}$  will be placed relatively close to the bisector of  $\alpha_i$ . Since  $\alpha_i$  is itself small, a large deviation of  $c_{\text{new}}$  from the bisector of  $\alpha_i$  will create angles not only much smaller than  $\alpha_i/2$  but even negative (invalid) ones. Thus, the resulting mesh will have poor quality. To prevent this, we modify Eq. 3 in the following way:

$$c_{\text{new}} = \frac{1}{\sum_{i=1}^k 1/\alpha_i^2} \cdot \sum_{i=1}^k \frac{1}{\alpha_i^2} \cdot c_i \quad (4)$$


Namely, the  $c_i$  for small angles  $\alpha_i$  will carry more weight than for large angles. To demonstrate the robustness of our improvement, see Fig. 5, Fig. 6 and Table 1.

Despite the superior results of our weighted angle-based scheme, it still cannot guarantee that the new vertex position forms a valid triangulation. Similarly, the convergence of our scheme as well as the original scheme cannot be guaranteed in cases when the given mesh has invalid (inverted) triangles or when the mesh boundary is far from convex. In these cases, both schemes should be applied in a “smart” manner, namely, verifying that the triangles are still valid, or that the minimum angle of the adjacent triangles has been improved, before a vertex is moved. In some rare cases, both schemes may fail to improve the minimum angle when even Laplacian smoothing may improve it. A “combined” scheme that applies Laplacian smoothing when the angle-based method fails has extremely fast convergence and achieves the best of both worlds.


### 4.2 Area-based remeshing

The idea to use triangle areas as one of the criteria for triangulation optimization is not new. This usually

**Fig. 5 a–c** Comparison of smoothing methods. **a** Laplacian; **b** angle-based [19]; **c** weighted angle-based



**Fig. 6a–d** Comparison between smoothing methods. **a** The original mesh; **b** Laplacian; **c** angle-based [19]; **d** weighted angle-based. See Table 1 for a quantitative comparison

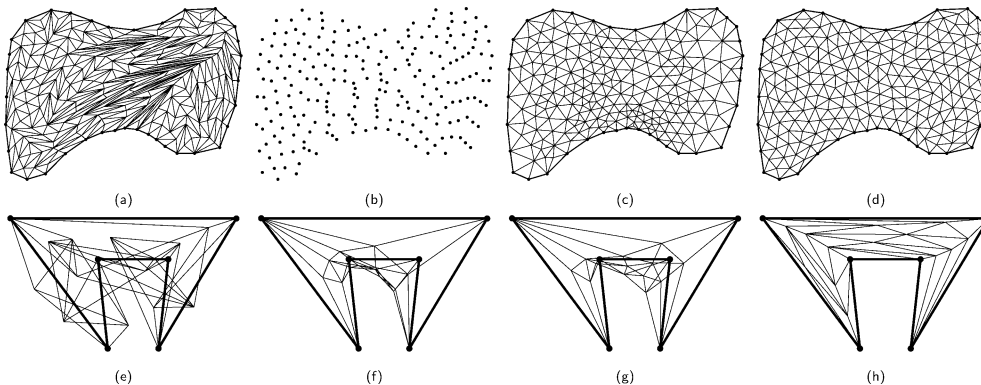


**Table 1** Quantitative comparison between quality of triangulations in Fig. 6. The right three columns correspond to the relative number of faces whose minimal angle is less than a specific value

Scheme	Min $\angle$ ( $^\circ$ ) $< 10^\circ$ (%)	$< 15^\circ$ (%)	$< 20^\circ$ (%)
Laplacian	0.17	2.57	5.31
Angle-based	4.62	0.58	1.66
Weighted angle-based	17.2	0.00	0.00

means trying to form triangles with as uniform an area as possible. However, triangle areas alone cannot be used to obtain meshes of reasonable quality. The reason is that when only the areas are optimized, without taking into account the angles, the resulting mesh can (and in most cases will) have many long and skinny triangles. Only when a mesh has an almost regular connectivity may uniform triangle areas imply well-shaped triangles. Nevertheless, a mesh containing triangles with areas close to equal has one important property: the spatial distribution of the vertices over the total mesh area is very uniform. If we eliminate the edges of the mesh leaving only the vertices, we obtain quite a uniform point distribution, as may be seen in Fig. 7b. In the combinatorial literature (e.g. [20]) this point set is said to have low discrepancy. This usually means that given an arbitrary rectangle of area  $A$  containing  $k$  points of the total  $n$  points, the ratio  $k/n$  is very close to  $A/S$ , when  $S$  is the total area encompassed by the point set.

**Fig. 7a–h** Area-based remeshing. **a** Triangle areas of the mesh from Fig. 2 are equalized; **b** discarding the edges of (a) reveals a uniform vertex distribution; **c** mesh obtained by alternation of angle-based smoothing and weighted angle-improving edge-flips; **d** mesh obtained by alternation of area equalization and edge-flips; **e** a polygon with random distribution of interior vertices; **f** Laplacian smoothing of (e); **g** angle-based smoothing of (e); **h** area equalization of (e)



We propose a remeshing scheme that utilizes this. Given a mesh, we alternate between the area equalization procedure and a series of angle-improving (Delaunay) edge-flips. Edge-flips are performed until improvement is no longer possible. This process results in a mesh that is as close to regular as the ratio between the number of the boundary and interior vertices, together with the geometry of the boundary, allows. It is far superior to the results from an analogous scheme involving angle-based smoothing instead of area equalization. Figure 7c and d compare the two schemes.

To equalize the areas of the mesh triangles, a number of iterations are performed over the mesh. Each iteration moves all the mesh interior vertices sequentially to improve the areas locally. Let  $p = (x, y)$  be an interior mesh vertex and  $p_1, \dots, p_k$  its neighbours.  $(x_i, y_i)$  are the coordinates of  $p_i$ . Denote by  $A_i(x, y)$  the area of triangle  $(p_i, p_{i+1}, p)$ . Note that  $i + 1$  is modulo  $k$ :

$$A_i(x, y) = \frac{1}{2} \begin{vmatrix} x_i & y_i & 1 \\ x_{i+1} & y_{i+1} & 1 \\ x & y & 1 \end{vmatrix} \tag{5}$$

Let  $A$  be the area of the polygon  $(p_1, \dots, p_k)$ , which may be computed as  $\sum_{i=1}^k A_i(0, 0)$ . In order to find the position of  $p$  that equalizes the areas of the adjacent triangles as much as possible, we minimize the following function:

$$(x, y) = \arg \min_{(x, y)} \sum_{i=1}^k \left( A_i(x, y) - \frac{A}{k} \right)^2 \tag{6}$$

This reduces to solving a system of two linear equations in  $x$  and  $y$ . The computational cost of this unique solution is close to that of traditional Laplacian smoothing.

It turns out that a valid mesh can be obtained by equalizing the areas of the mesh triangles, even in cases

such as a highly non-convex boundary. This contrasts with other methods, including the smart Laplacian [21] and both angle-based smoothing methods, which fail (see Fig. 7e–h).

## 5 Compatible remeshing

We now show how to combine the two methods introduced in Sect. 3, along with a refinement procedure (introducing new interior Steiner vertices), to produce high-quality compatible triangulations of two given polygons with a correspondence between their vertices. Compatible triangulations created using the method introduced in Sect. 2.2 usually have a small number of Steiner vertices, but their quality is unlikely to be acceptable. Therefore, remeshing techniques must be applied to improve the quality. The main difficulty with using existing remeshing techniques is that the remeshing criteria that are suitable for a single mesh may fail when applied to two triangulations in parallel.

Our compatible remeshing technique is similar to that of Alexa et al. [9]. The outline of the algorithm appears in Fig. 9. We use a series of simultaneous edge-flips, mesh smoothing and mesh refinement by edge splitting. The parameter  $k_{\text{split}}$  dictates the rate at which new Steiner vertices are introduced. In addition, we control vertex distribution by applying a single iteration of the area equalization technique described in Sect. 3.2 after  $k_{\text{split}}$  refinement operations.

While the criteria for operations in the algorithm in Fig. 9 are rather straightforward for a single mesh, applying them simultaneously on two triangulations requires more precise control. The corresponding properties of triangles within the two meshes may often contradict each other. If care is not exercised, the process of compatible remeshing may not lead to overall

improvement in quality of both meshes. As a result a large number of vertices can be added to the meshes without any gain in quality.

The following empirical criteria, based on their analog for a single mesh, have produced the best results on numerous examples:


- *Edge-flips*: Similarly to when constructing Delaunay triangulations, the edge is flipped only if the minimum angle between the angles of both meshes of the triangles adjacent to the edge is improved.
- *Angle-based smoothing*: Both meshes are independently smoothed, applying the technique described in Sect. 3.1 in the “smart” manner, namely, preserving the validity of both meshes.
- *Edge-split refinement*: Our criterion for choosing an edge  $e$  to be split is based both on the edge length denoted by  $|e|$  and the minimal of the four adjacent triangle angles denoted by  $\alpha_{\min}(e)$ . The edge with the maximal “normalized” length in both triangulations ( $T_0$  and  $T_1$ ) is refined:

$$e = \arg \max_{e \in T_0 \cup T_1} \frac{|e|}{(\alpha_{\min}(e))^2} \quad (7)$$

Note that the refinement is performed simultaneously on both triangulations in order to preserve the compatibility. The criterion defined in Eq. 7 produces better experimental results than the aspect ratio-based criterion of [22] or distortion metrics criteria of [23] and [21]. The number of edges to be split in each iteration ( $k_{\text{split}}$ ) determines the trade-off between the number of Steiner vertices together with mesh quality and the algorithm running time.

- *Area equalization*: As noted in Sect. 3.2, area equalization improves the spatial vertex distribution. Because of the refinement operations, some regions of the mesh may have an excess in vertex density. To smooth this out, we apply a single iteration of area equalization (step 4). This area equalization can prevent a further increase in the number of Steiner vertices at later stages, but at the price of slowing down the algorithm. See Fig. 8. On the one hand, the refinement operations change the meshes locally and thus step 1 (or 3) of the algorithm (Fig. 9) converges quickly. On the other hand, the area equalization affects the mesh globally and thus step 5 takes much longer to improve the mesh globally. If a faster

**Fig. 8a–f** High-quality compatible triangulation of letters  $U$  and  $S$ . **a–b** Optimal compatible triangulations generated by the algorithm of Sect. 2.2. No Steiner vertices are required, but the minimum angle of  $U$  is  $10.8^\circ$  and of  $S$  is  $3.4^\circ$ . **c–d** Compatible triangulations generated by the algorithm of Sect. 4, without area equalization. The number of Steiner vertices is 27 and the minimum angles are  $15.4^\circ$  and  $15.7^\circ$  respectively. **e–f** Compatible triangulations generated using area equalization. The number of Steiner vertices is 7, the minimum angles are  $17.1^\circ$  and  $17.6^\circ$ . The time required to generate (c), (d) and (e), (f) was similar



**while** mesh quality has not been achieved **or** the number of Steiner vertices does not exceed threshold

**Step 1.** Alternate between angle-based smoothing and simultaneous angle-improving edge-flips.

**Step 2.** Refine both meshes by  $k_{\text{split}}$  simultaneous edge-splits.

**Step 3.** The same as Step 1.

**Step 4.** Perform a single iteration of area equalization (see Sect. 3.2).

**Step 5.** The same as Step 1.


**Fig. 9** Compatible remeshing algorithm

algorithm is required, step 4 can be applied more rarely, for example every 2–20 iterations, or can only be applied in regions affected by refinement operations of step 2.

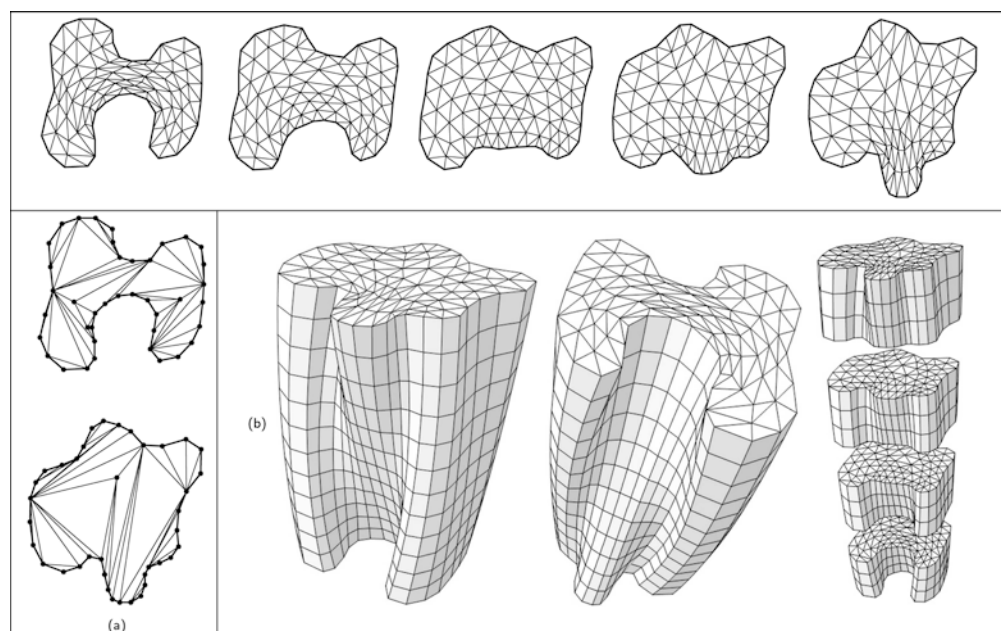
## 6 Experimental results

We have implemented all the algorithms described in this paper and applied them to numerous example inputs. Our inputs consist of two planar polygons, which serve as the source and target (top and bottom) cross-sections of the sweep. These two are compatibly triangulated with sufficient mesh quality (using the algorithm in Fig. 9) and then morphed to create intermediate compatibly triangulated polygons. Especially challenging inputs are when the source and target are significantly different. Figures 10, 11 and 12 show some sample input pairs, the compatible triangulations (usually of low quality) with a small number of Steiner vertices generated by the methods of Sect. 2.2, the remeshed high quality compatible triangulations generated by the methods of Sect. 4 and the intermediate


**Fig. 10a, b** 3D sweep generation. **a** Optimal (no Steiner vertices) compatible triangulation of source and target polygons. *Top row:* high-quality compatible triangulation and intermediates generated by morphing procedure. Minimum angles of the source and target triangulations are  $27.2^\circ$  and  $25.9^\circ$ , respectively. **b** 3D visualization of sweeps from a number of different angles



**Fig. 11a, b** 3D sweep generation. **a** Compatible triangulation of source and target polygons with three Steiner vertices. *Top row:* high quality compatible triangulation and intermediates generated by morphing procedure. Minimum angles of the source and target triangulations are  $15.9^\circ$  and  $15.3^\circ$ , respectively. **b** 3D visualization of sweeps from a number of different angles







**Fig. 12** 3D visualizations of sweeps between letters  $U$  and  $S$  using compatible triangulations from Fig. 8e–f. The optimal compatible triangulations of  $U$  and  $S$  appear in Fig. 8a–b

triangulated cross-sections generated by applying morphing techniques. The latter are shown both as a sequence of 2D cross-sections and as a sliced 3D sweep. For each example, we specify the statistics of the source and target meshes. We found that the angles of the intermediate meshes generated using the techniques of Surazhsky and Gotsman [6,13] were always in between those two, so the mesh quality is preserved throughout the morph.

In terms of runtimes, all these examples required no more than a second or so to run on a Athlon 1.2 GHz PC with 256 MB RAM. Larger inputs, which ultimately involved hundreds of (interior and exterior) Steiner vertices for the mesh and the morph, required no more than 5 s on the same machine.

## 7 Discussion and conclusion

We have shown how to generate compatibly triangulated sweeps with quality adequate for finite-element analysis. Our method is fast, robust and, as opposed to previously published methods, is guaranteed to always produce a valid result.

Several components of our algorithm, in particular the weighted angle-based smoothing procedure, may be used in their own right in other meshing applications.

The method was designed primarily for parallel planar inputs, but can probably be extended easily to more general cases. A straightforward extension is to generalize all presented algorithms to a sequence of  $n > 2$  input polygons. Another extension is to treat a set of source and target polygons, forming so-called “multi-sweeps” or “barrels” [10]. Compatible hexahedral meshes is another challenging direction for future work.

**Acknowledgements** The work was carried out while the authors were at Technion—Israel Institute of Technology. Thanks to Tatiana Surazhsky and Michael Floater for their contribution to the area-based remeshing method, to Alla Sheffer for helpful discussions on sweeps and to Gill Barequet for helpful discussions

on the implementation of minimum-link path algorithms. This work was partially supported by the Technion Computer Science Software Technology Laboratory (STL) and the Technion Fund for Promotion of Research.

## References

1. Staten ML, Canann SA, Owen SJ (1998) BMSweep: locating interior nodes during sweeping. In: Proceedings of the 7th International Meshing Roundtable, 26–28 October 1998, Dearborn, MI, USA, pp 7–18
2. Shapira M, Rappoport A (1995) Shape blending using the star-skeleton representation. *IEEE Trans Comput Graphics Applic* 15(2):44–51
3. Sederberg TW, Greenwood E (1992) A physically based approach to 2D shape blending. *Comput Graphics (SIGGRAPH '92)* 26:25–34
4. Sederberg TW, Gao P, Wang G, Mu H (1993) 2D shape blending: an intrinsic solution to the vertex path problem. *Comput Graphics (SIGGRAPH '93)* 27:15–18
5. Aronov B, Seidel R, Souvaine DL (1993) On compatible triangulations of simple polygons. *Comput Geom Theor Applic* 3:27–35
6. Gotsman C, Surazhsky V (2001) Guaranteed intersection-free polygon morphing. *Comput Graphics* 25:67–75
7. Kranakis E, Urrutia J (1999) Isomorphic triangulations with small number of Steiner points. *Int J Comput Geom Applic* 9:171–180
8. Gupta H, Wenger R (1997) Constructing piecewise linear homeomorphisms of simple polygons. *J Algorithms* 22:142–157
9. Alexa M, Cohen-Or D, Levin D (2000) As-rigid-as-possible polygon morphing. In: Proceedings of SIGGRAPH 2000, 23–28 July 2000, New Orleans, LA, pp 157–164
10. Blacker T (1996) The Cooper tool. In: Proceedings of the 5th International Meshing Roundtable, 10–11 October 1996, Pittsburgh, PA, pp 13–19
11. Knupp PM (1998) Next generation sweep tool: a method for generating all-hex meshes and two and one-half dimensional geometries. In: Proceedings of the 7th International Meshing Roundtable, 26–28 October 1998, Dearborn, MI, pp 505–514
12. Floater MS, Gotsman C (1999) How to morph tilings injectively. *Comput Appl Math* 101:117–129
13. Surazhsky V, Gostman C (2001) Controllable morphing of compatible planar triangulations. *ACM Trans Graphics* 20:203–231
14. Tutte WT (1963) How to draw a graph. *Proc Lond Math Soc* 13:743–768
15. Floater MS (1997) Parameterization and smooth approximation of surface triangulation. *Comput Aided Geom Design* 14:231–250
16. Babikov M, Souvaine DL, Wenger R (1997) Constructing piecewise linear homeomorphisms of polygons with holes. In: Proceedings of the 9th Canadian Conference on Computational geometry, 11–14 August 1997, Kingston, Ontario, Canada, pp 6–10
17. Suri S (1986) A linear time algorithm for minimum link paths inside a simple polygon. *Comput Vis Graph Image Process* 35:99–110
18. Suri S (1990) On some link distance problems in a simple polygon. *IEEE J Robotics Automat* 6:108–113
19. Zhou T, Shimada K (2000) An angle-based approach to two-dimensional mesh smoothing. In: Proceedings of the 9th International Meshing Roundtable, 2–5 October 2000, New Orleans, LA, pp 373–384
20. Niederreiter H (1992) Random number generation and quasi-Monte Carlo Methods. SIAM
21. Canann SA, Tristano JR, Staten ML (1998) An approach to combined Laplacian and optimization-based smoothing for triangular, quadrilateral, and quad dominant meshes.

- In: Proceedings of the 7th International Meshing Roundtable, 26–28 October 1998, Dearborn, MI, pp 479–496
22. Parthasarathy V, Kodiyalam S (1991) A constrained optimization approach to finite element mesh smoothing. *Finite Elem Anal Design* 9:309–320
  23. Berzins M (1998) Mesh quality: a function of geometry, error estimates or both? In: Proceedings of the 7th International Meshing Roundtable, 26–28 October 1998, Dearborn, MI, pp 229–238