BUREAU OF THE CENSUS
STATISTICAL RESEARCH DIVISION REPORT SERIES
SRD Research Report Number: Census/SRD/RR-87/23

Joint Triangulations and Triangulation Maps

by

Alan Saalfeld
Statistical Research Division
Bureau of the Census

Recommended by: Tom O'Reagan

Report completed: September, 1986

Report issued: October 1, 1987

JOINT TRIANGULATIONS AND TRIANGULATION MAPS

Alan Saalfeld
Statistical Research Division
Bureau of the Census
Washington, DC 20233

ABSTRACT

In rubber-sheeting applications in cartography, it is useful to seek piecewise-linear homeomorphisms (PLH maps) between rectangular regions which map an arbitrary sequence of n points $\{p_1, p_2, \ldots, p_n\}$ from the interior of one rectangle to a corresponding sequence $\{q_1, q_2, \ldots, q_n\}$ of n points in the interior of the second region. This paper proves that it is always possible to find such PLH maps and describes them in terms of a joint triangulation of the domain and the range rectangular regions.

One naive approach to finding a PLH map is to triangulate (in any fashion) the domain rectangle on its n points and four corners and to define a piecewise affine map on each triangle $\Delta p_{i1}p_{i2}p_{i3}$ to be the unique affine map that sends the three vertices p_{i1}, p_{i2}, p_{i3} of the triangle to the three corresponding vertices q_{i1}, q_{i2}, q_{i3} of the image triangle $\Delta q_{i1}q_{i2}q_{i3}$. Such piecewise affine maps send triangles to triangles, agree on shared edges, and thus extend globally, and will be called triangulation maps. The shortcoming of building transformations in this fashion is that the resulting triangulation map need not be one-to-one, although there is a simple test to determine if such a map is one-to-one (see Theorem 2 below). If the map is one-to-one, then the image triangles will form a triangulation of the range space; and we will have a joint triangulation. If the map is not one-to-one, then there will be folding over of triangles. It may be possible to alleviate this folding by choosing a different triangulation of the n domain points, or it may be the case that no triangulation of the n domain points will work. (See figures 5 and 6 below). We show that it will be possible, in all cases, to rectify the folding by adding appropriate additional triangulation vertex pairs $\{p_{n+1}, p_{n+2}, \ldots, p_{n+m}\}$ and $\{q_{n+1}, q_{n+2}, \ldots, q_{n+m}\}$ and retriangulating (see Theorem 1 below). This paper examines conditions for triangulation maps to be homeomorphisms and explores different ways of modifying triangulations and triangulation maps to make them joint triangulations and homeomorphisms.

The paper concludes with a section on alternative constructive approaches to the open problem of finding joint triangulations on the original sequences of vertex pairs without augmenting those sequences of pairs.

The existence proofs in this paper do not solve computational geometry problems per se; instead they permit us to formulate new computational geometry problems. The problems we pose are of interest to us because of a particular application in automated cartography.

1. BASIC CONCEPTS AND DEFINITIONS

1.1. Refinements and Triangulation Maps

A refinement T' of a triangulation T is a triangulation, each of whose triangles is contained in a triangle of T. It is easily shown that a triangulation T' is a refinement of a triangulation T if and only if each edge of T is a union of one or more edges of T'. Refinements play an important role in the definition of triangulation maps; refinements provide sufficient flexibility in the triangulating set to guarantee that triangulation maps, as we define them, will be closed under composition.

A mapping Ψ from R_1 to R_2 is a triangulation map if there exists a triangulation T_1 of R_1 such that Ψ restricted to every triangle of T_1 is affine.

Proposition. The composition of two triangulation maps is a triangulation map.

$$T_1 \dashrightarrow^{\phi} T_1'$$

$$R_1 \longrightarrow R_2 \longrightarrow R_3$$

$$T_2 \dashrightarrow^{\Psi} T_2'$$

Proof that $\Psi \circ \phi$ is a triangulation map involves choosing a suitable triangulation refinement of the polygon overlay formed by the edges of T_1 and $\phi^{-1}(T_2)$.

1.2. PLH Maps and Joint Triangulations

If a triangulation map is a homeomorphism, then it is a PLH map, and the two associated triangulations are a joint triangulation.

2. FUNDAMENTAL EXISTENCE THEOREM

The following existence theorem shows that triangulation refinement can be a valuable tool in building piecewise linear homeomorphisms. The proof is by geometric construction:

Theorem 1. Given $\{p_1, p_2, \ldots, p_n\}$ all distinct in the interior of rectangle R_1, and $\{q_1, q_2, \ldots, q_n\}$ all distinct in the interior of rectangle R_2, there is a triangulation $\{\Delta_1, \Delta_2, \ldots, \Delta_\varepsilon\}$ of R_1, and a homeomorphism $\phi: R_1 \longrightarrow R_2$, such that $\phi:\Delta_i$ is affine for $i = 1, 2, \ldots, \varepsilon$, and $\phi(p_i) = q_i$ for $i = 1, 2, \ldots, n$.

The vertex set of the triangulation of R_1 may be chosen to contain the points $\{p_1, p_2, \ldots, p_n\}$, although it will usually not be possible to triangulate on only the set $\{p_1, p_2, \ldots, p_n\}$.

The proof of theorem 1 is by induction. Figure 1 illustrates how to triangulate and define our PLH affine map when $n = 1$.

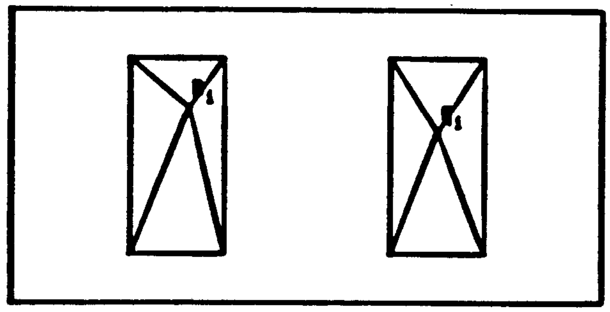

Figure 1. Joint triangulation on one interior point.

Suppose that we can find a desired piecewise map ϕ that maps the first k points, $\{p_1, p_2, \ldots, p_k\}$, onto the corresponding k points, $\{q_1, q_2, \ldots, q_k\}$, of the other set. If p_{k+1} goes to q_{k+1} under ϕ, then we are done. So, suppose otherwise. We now define a series of "local" modifications A_i (also triangulation maps) to be composed with ϕ which will move the image of p_{k+1} across one triangle at a time until it aligns with q_{k+1}.

Let T' be an image triangulation $\Phi(T)$ of a triangulation T of the domain space that includes $\{p_1, p_2, \ldots, p_k\}$, the first k points, in its vertex set. Then T' includes the set $\{q_1, q_2, \ldots, q_k\}$. Let $z = \Phi(p_{k+1})$ lie in triangle t_0; and let $t_0, t_1, \ldots, t_j$ be a sequence of triangles in T' such that t_i is adjacent to t_{i+1}, and q_{k+1} is in t_j.

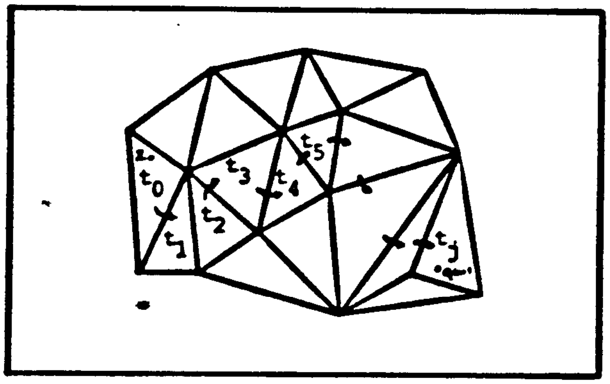

Figure 2. A sequence of adjacent triangles from z to q_{k+1}.

We now define triangulation maps $\{A_i\}$ on the image space which move only points in the two adjacent triangles t_i and t_{i+1} and leave all other triangles fixed. We may make these maps carry the centroid of t_i onto the centroid of t_{i+1} using the following "diamond trick" refinement:

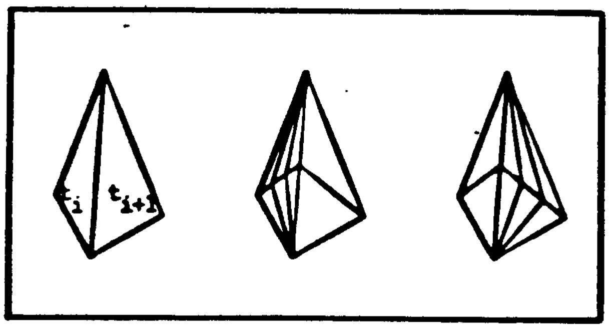

Figure 3. Eight corresponding triangle pairs of diamond trick

The "diamond trick" subdivides the two adjacent triangles into eight subtriangles in two different ways so that the obvious correspondence of subtriangles and their vertices associates the centroid of t_i in the first subdivision with the centroid of t_{i+1} in the second subdivision.

A simple piecewise affine transformation v, involving only a refinement of t_0 and the identity map outside of t_0, moves z onto the centroid of t_0 and another u moves the centroid of t_j onto q_{k+1}:

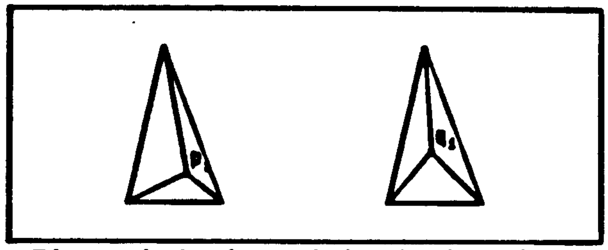

Figure 4. Moving a triangle interior point to the centroid.

Then the composite (triangulation) map: $u \circ A_j \circ A_{j-1} \circ \ldots \circ A_1 \circ A_0 \circ v$ will send p_i to q_i for all $i = 1, 2, \ldots, k+1$.

The above existence proof does not translate very efficiently into a constructive procedure. The composite functions used to "move" the image of p_{k+1} over to q_{k+1} require elaborate triangulation refinements to be realized. An alternative approach to the so-called "extension problem" described above involves searching for a joint triangulation on only the original sets $\{p_1, p_2, \ldots, p_n\}$ and $\{q_1, q_2, \ldots, q_n\}$, and augmenting those sets only when necessary. Determining when augmentation is necessary, how much augmentation is necessary, and how to augment efficiently are much harder problems.

3. THE GENERAL UNAUGMENTED PROBLEM

3.1. The Problem Statement

We restate the open problem in a slightly more general context: Given two arbitrary sequences of n points, $P = \{p_1, p_2, \ldots, p_n\}$ and $Q = \{q_1, q_2, \ldots, q_n\}$, when can we triangulate the convex hull of the first n points and associate with each component triangle $\Delta p_{i1} p_{i2} p_{i3}$ the corresponding triangle $\Delta q_{i1} q_{i2} q_{i3}$ in such a way that the resulting set of triangles in the convex hull of the second set of n points gives a triangulation of the second convex hull? In other words, when and how can we triangulate on the indices of the p_i's and q_i's simultaneously to obtain triangulations of both sets? The goal is a procedure which finds a joint triangulation when it exists and which determines efficiently when no such joint triangulation exists.

3.2. Some Examples

Before discussing progress on the problem, it is illuminating to see two simple examples that show that there are pairs of sequences for which no joint triangulation exists (see Figure 5 below) and also show that there are sequences P and Q for which some triangulations of P yield triangulations of Q and for which other triangulations of P do not (see Figure 6 below).

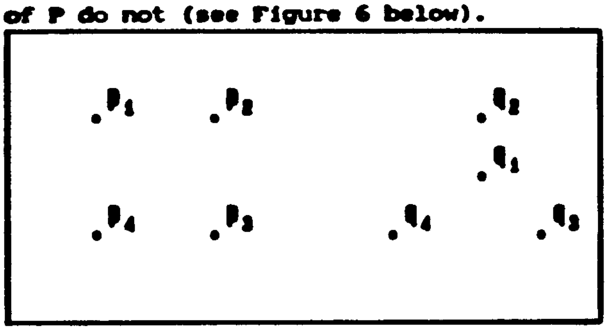

Figure 5. No joint triangulation exists.

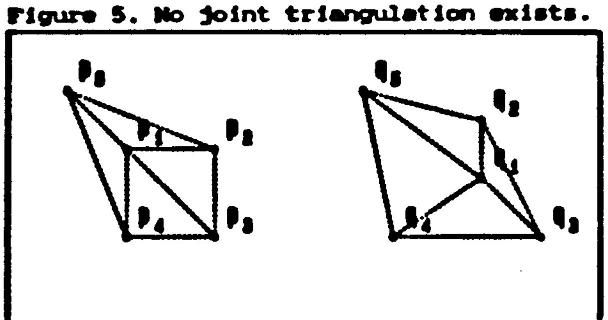

Figure 6. One triangulation of P works.

3.3. Non-constructive Approaches

A non-constructive approach to the problem would determine whether or not a joint triangulation exists without actually finding one. Such an approach would describe verifiable sufficient conditions for existence. Practical sufficient conditions are more difficult to find than necessary conditions; and so far we have no sufficient conditions.

In the remainder of this section, we analyze necessary conditions for the existence of a joint triangulation in order to characterize triangles along the outer edges. In a later section, we sketch an approach which builds a joint triangulation from the outside-in by first finding paired triangles having a vertex on the boundary of the convex hulls.

Suppose that a joint triangulation exists and it is given to us. By examining some of its properties, we can derive necessary conditions for its existence:

1. Suppose the k vertices p_{i1}, p_{i2}, ..., p_{ik} make up the extreme points of the convex hull for P (in cyclic order). Then the k corresponding vertices q_{i1}, q_{i2}, ..., q_{ik} make up the extreme points of the convex hull for Q; and consecutive points are linked by edges. (This necessary condition was violated in Figure 5; hence, there was no joint triangulation.)

2. Let P* be the extreme points in the convex hull of P; and let Q* denote the extreme points in the convex hull of Q. Let R_p (the "ring" of P) be the union of all of those triangles having a vertex in P*; and let R_q (the "ring" of Q) be the union of all of those triangles which have a vertex in Q*. Then the triangles (and their vertices) of R_p must match the triangles (and their vertices) of R_q.

3. Let P' be P − P*; and let Q' be Q − Q*. Then P' and Q' are made up of interior or non-extremal points. Let $E_{P'}$ (the enclosing polygon of P') be the union of the remaining triangles after removing those triangles making up the

ring. Similarly define $E_{Q'}$. Then the triangles of $E_{P'}$ have only interior (non-extreme) points of $p_1, p_2, \ldots, p_n$ as vertices; and the triangles of $E_{Q'}$ have only interior (non-extreme) points of $q_1, q_2, \ldots, q_n$ as vertices. The triangles of $E_{P'}$ match the triangles of $E_{Q'}$; and the boundary of $E_{P'}$ must match the boundary of $E_{Q'}$ point for point and edge for edge.

4. The p_i's and the q_i's in the boundary of $E_{P'}$ and $E_{Q'}$ need not consist of only extreme points of the reduced convex hulls of P' and Q', respectively; nevertheless, points which generate those convex hulls must lie in the boundaries of the enclosing polygons $E_{P'}$ and $E_{Q'}$, for, otherwise, they could not be enclosed by polygons stringing other p_i's or q_i's together. The figures below illustrate the sets P, Q, R_P, R_Q, $E_{P'}$, and $E_{Q'}$. Notice that the intersection of R_P and $E_{P'}$ is a cycle of edges that must match, edge for edge, the intersection of R_Q and $E_{Q'}$.

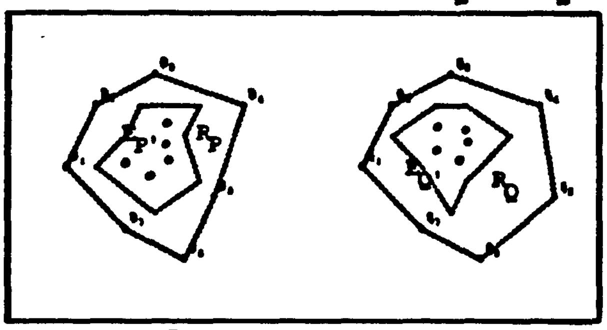

Figure 7. Rings and polygons for matching triangulations

Finally, the corner visibility conditions from vertices on these cycles must agree sufficiently to complete a triangulation of the rings R_P and R_Q. We list some of those conditions below:

5. Suppose we wish to triangulate an annular region using only lines crossing between the inner and outer polygons. It is necessary (figure 8a), but not sufficient (figure 8b), that every vertex of either polygon must be able to see at least one vertex of the other polygon.

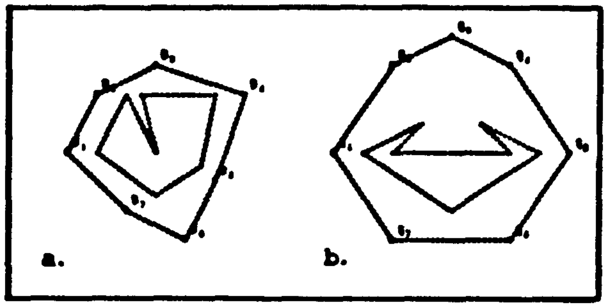

Figure 8. Two cases for which no ring triangulation exists using cross lines. One problem with the second example above is that the visibility conditions, while overlapping, seem to double back. We define the visibility set V_i of a vertex b_i of either polygon to be the collection of all vertices of the other polygon which can be linked to the vertex b_i without passing through either polygon. Visibility constraints of an individual vertex must somehow consider constraints for neighboring vertices simultaneously, since lines drawn from the neighboring vertices can restrict visibility, and since each triangle vertex is linked to (at least) <u>three</u> other vertices (its two polygon neighbors and one vertex opposite).

6. A better criterion (necessary and sufficient) for the existence of an annular triangulation using only cross lines requires that the visibility sets V_i of each vertex b_i of the inner (or outer) polygon each contain a subset U_i of <u>consecutive</u> vertices of the other polygon such that the intersection of U_i and U_{i+1} is a singleton vertex which is the right-most vertex of U_i and the left-most vertex of U_{i+1}, in the cyclic ordering regarded as left-to-right. The U_i's are nothing more than the collection of vertices to be linked to the corresponding b_i's in a legitimate triangulation of the annular region.

4. TESTING TRIANGULATION MAPS FOR HOMEOMORPHISM

Before examining several constructive approaches for joint triangulations, there are some results from applications in automated cartography that are worth mentioning. For those applications, it is desirable to have Delaunay triangulations on one or both of the spaces to be jointly triangulated. (In practice, this reduces the likelihood of folding.) It is also desirable to use as few additional points as possible in a joint triangulation. (In practice, matched pairs are derived from non-geometric external sources.) As a result of these criteria, one usually builds a Delaunay triangulation of the domain space and then simply checks to see if the image triangles fit together in a triangulation. If they do fit, then there will be no further work. If not, then a joint triangulation is sought first by trying alternative triangulations of the domain space. If desired, we can always recover Delaunayhood in the domain space by refinement. The following theorem provides a simple test of given triangulation map to see if it is a homeomorphism:

Theorem 2. Suppose that a triangulation mapping from a convex region maps the boundary of the convex region one-to-one onto a simple polygon. Then the triangulation mapping is bijective if and only if it either preserves orientation on every triangle or it reverses orientation on every triangle.

Let Ψ: C ---> D be a triangulation mapping which sends ∂C, the boundary of C, one-to-one onto the simple polygon S in D.

By connectivity properties, the image triangles must cover the interior of the simple polygon S; for otherwise, the image of the convex region D could be retracted onto the simple polygon; and the composition of the triangulation mapping with the retraction followed by the inverse of the triangulation map (restricted to the boundary) would retract the convex region D onto its boundary, ∂D, which is impossible.

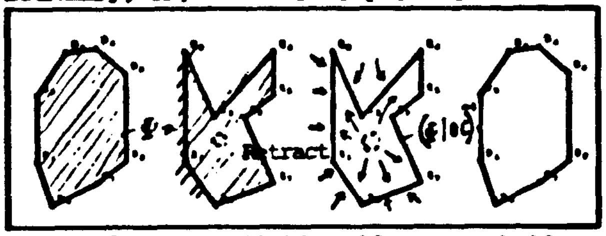

Figure 9. Contradiction if convex hull image does not cover simple polygon.

The sum of the areas of the image triangles must always be at least as great as the area of the simple polygon. We show that the areas are equal if and only if all triangles have the same orientation.

Our proof of the theorem makes use of an area formula for a simple polygon.

Let $\{p_1, p_2, \ldots, p_m\}$ be the vertices of the bounding polygon ∂C in counter-clockwise order; and to simplify notation, let $p_0 = p_m$. Let $\{p_1, p_2, \ldots, p_m, p_{(m+1)}, \ldots, p_n\}$ be all of the vertices (including interior vertices) in our domain region being triangulated. Denote the $(2n-m-2)$ triangles of a triangulation T of C by:

$$\{\Delta p_{j1} p_{j2} p_{j3} : j = 1, 2, \ldots, (2n-m-2)\},$$

where we will let $p_{j0} = p_{j3}$ for convenience of notation, and assume the triangle vertices are given in counter-clockwise order. We let $p_i = (x_i, y_i)$, and similarly $p_{jk} = (x_{jk}, y_{jk})$.

Then the area of C is:

$$\text{Area}(C) = \sum_{i=1}^{m} (y_i x_{(i-1)} - x_i y_{(i-1)})/2.$$

Also, each triangle has area:

$$\text{Area}(\Delta p_{j1}p_{j2}p_{j3}) = \sum_{i=1}^{3} (y_{ji}x_{j(i-1)} - x_{ji}y_{j(i-1)})/2.$$

Therefore:

$$\text{Area}(C) = \sum_{j=1}^{(2n-m-2)} \text{Area}(\Delta p_{j1}p_{j2}p_{j3})$$

$$= \sum_{j=1}^{(2n-m-2)} \sum_{i=1}^{3} (y_{ji}x_{j(i-1)} - x_{ji}y_{j(i-1)})/2.$$

Notice that each $(y_{ji}x_{j(i-1)} - x_{ji}y_{j(i-1)})$ term in the double sum corresponds to a particular directed segment, and that each interior triangle segment belongs to exactly two triangles, and the signs of the terms are opposite—hence terms cancel each other for interior segments. Thus one may derive the formula for the simple polygon from the triangle formulas using geometric or algebraic reasoning. Moreover, the algebraic arguments about cancelling of terms hold even if the triangles fold over.

Now suppose that, under a triangulation map, the p_i's map onto q_i's where $q_i = (a_i, b_i)$. If the $(q_1, q_2, \ldots, q_m)$ form the vertices of a simple polygon S in counter-clockwise order; and if, as before, to simplify notation, we let $q_0 = q_m$; then the following derivations also give the area of S:

$$\text{Area}(S) = \sum_{i=1}^{m} (b_i a_{(i-1)} - a_i b_{(i-1)})/2.$$

$$= \sum_{j=1}^{(2n-m-2)} \sum_{i=1}^{3} (b_{ji}a_{j(i-1)} - a_{ji}b_{j(i-1)})/2.$$

$$= \sum_{j=1}^{(2n-m-2)} \text{Signed Area}(\Delta q_{j1}q_{j2}q_{j3})$$

(by the same algebraic arguments given for the p_i's)

where:

$$\text{Signed Area}(\Delta q_{j1}q_{j2}q_{j3}) = \sum_{i=1}^{3} (b_{ji}a_{j(i-1)} - b_{ji}a_{j(i-1)})/2.$$

$$= \pm \text{Area}(\Delta q_{j1}q_{j2}q_{j3})$$

where the sign is (+) if the vertices are ordered counter-clockwise and (-) if the vertices are clockwise in order.

The sum of the signed triangle areas is always exactly equal to the area of the simple polygon and is always less than or equal to the sum of the actual triangle areas. Moreover, the sum of the signed areas can only be equal to the sum of the areas if all signs are positive; and that will happen if and only if no orientations are reversed.

Conversely, if no orientations are reversed, then the true triangle areas will exactly account for all of the area of the simple polygon—hence there can be no overlapping; and the mapping must be one-to-one.

The case where all orientations of triangles are reversed corresponds to the situation in which the boundary orientation is changed to clockwise.

5. PRELIMINARY CONSTRUCTIVE APPROACHES TO THE OPEN PROBLEM

The problem of constructing a joint triangulation remains open. It is not yet known if the problem is NP-complete. Some possible approaches to the problem which have not been fully successful are described below.

5.1. Find all eligible triangles.

The following approach to joint triangulation was suggested by David Mount of the University of Maryland.

Form the set of all triangles on the set $P = \{p_1, p_2, \ldots, p_n\}$ that satisfy the following necessary conditions for the existence of a joint triangulation:

1. Each component triangle $\Delta p_{i1} p_{i2} p_{i3}$ and its associated triangle $\Delta q_{i1} q_{i2} q_{i3}$ in $Q = \{q_1, q_2, \ldots, q_n\}$ is empty (that is, the triangle contains no other vertices from the same set P or Q.)

2. Each component triangle $\Delta p_{i1} p_{i2} p_{i3}$ and its associated triangle $\Delta q_{i1} q_{i2} q_{i3}$ in the set $\{q_1, q_2, \ldots, q_n\}$ have the same order type or orientation.

Forming this set of eligible triangles can be done in polynomial time. The unresolved issue is selecting a subset of the triangles on one of the sets $\{p_1, p_2, \ldots, p_n\}$ or $\{q_1, q_2, \ldots, q_n\}$ such that the subset triangulates its space. By our result in Theorem 2, the other space will be triangulated as well by the corresponding subset of triangles.

It is known that the general problem TRI of finding a triangulating subset of edges among an arbitrary collection of edges is NP-complete.

It is not known to the author if the general problem of finding a triangulating subset of triangles among an arbitrary collection of triangles is NP-complete.

It is not known whether the collection of triangles arising from the joint triangulation conditions (1) and (2) above will be well-behaved enough to permit special analysis and a polynomial time selection of a triangulating subset.

Finding a triangulation seems to be a two-dimensional analogue to finding a Hamiltonian circuit; and there is no guarantee or even expectation that local adjustment will suffice for finding a triangulation from a group of triangles in the general situation. Nevertheless, there are special properties of the triangles of joint triangulations that permit some local analysis, especially along convex hulls and convex layers.

The next approach has thus far been least fruitful. It also is motivated by Theorem 2.

5.2. "Fix" a triangulation map.

The basic idea of this approach is to start with any triangulation map and to iteratively improve the map based on some measure of distance from homeomorphism, such as the number of triangles whose orientation is reversed, or the number of edge intersections occurring in the image.

This approach also locates the problem areas (where folding occurs) and focuses adjustments to the triangulation of the domain space in those areas.

In order for this approach to succeed, one must guarantee that the solution set is accessible from any triangulation by iterated improvements and that the iterated improvements (such as quadrilateral diagonal swapping) will converge to a global solution provided that one exists; and the improvements will terminate or will permit the recognition of the impossibility of finding a joint triangulation when none exists.

No progress has been made with this approach except to examine, in a few simple cases, the triangulation maps that can be converted to joint triangulations with a single diagonal swap.

The third approach makes use of many of the constraints that must be present in a joint triangulation.

3.3. Build rings from outside-in.

This approach uses the convex layers of both sets P and Q and tries to merge those layers to form ring-like or annular regions which are partial joint triangulations. The ring layers correspond to the depth (from the outside extreme points) of a minimal-depth joint triangulation—a point will be on ring layer m if and only if its graph distance to an extreme point is at least m on any joint triangulation.

The following lemma justifies our formation of annular regions.

Lemma. Suppose that not every point of P is an extreme point of P. Then if there is a joint triangulation of P and Q, then there is also a joint triangulation of P and Q having no edges between extreme points other than the necessary bounding hull edges.

This lemma is interesting in that it promises the existence of joint triangulations of a nice type if they exist at all!

The lemma is easily proved using the convexity of any quadrilateral whose diagonal is a straddling edge.

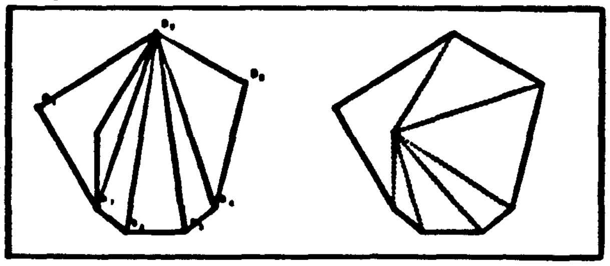

Figure 10. One can always swap with a straddling edge between extreme points.

Using the notation and properties described in section 3.3., we outline our approach to build a suitable joint triangulation of the two spaces by first constructing appropriate triangulations of rings R_P and R_Q. The rings R_P and R_Q (i.e. their triangulations) will be able to be specified provided that we can describe the enclosing polygons $E_{P'}$ and $E_{Q'}$, along with appropriate visibility constraints on their vertices.

In order to determine $E_{P'}$ and $E_{Q'}$, we define a pair of sequences of enclosing polygons $(F_{P'1}, F_{P'2}, \ldots, F_{P't})$ and $(F_{Q'1}, F_{Q'2}, \ldots, F_{Q't})$, such that:

1. Each $F_{P'j}$ and $F_{Q'j}$ is a legitimate enclosing polygon for the sets P' and Q' individually.

2. All the indices of $F_{P'i}$ are contained in $F_{Q'(i+1)}$.

3. All the indices of $F_{Q'i}$ are contained in $F_{P'(i+1)}$.

4. The indices of $F_{P't}$ and $F_{Q't}$ match exactly.

We start by letting $F_{P'1}$ and $F_{Q'1}$ be the first convex layer polygons for each of the interior sets (i.e. the extreme points of the sets P' and Q', in order, respectively).

We then modify the enclosing polygon $F_{P'i}$ to include interior points by examining constraints imposed by the other set's polygon, $F_{Q'i}$. For example, any index appearing in $F_{Q'i}$ must be added to $F_{P'(i+1)}$, and adding this index (in the appropriate position or sequence) will require (1) testing feasibility of addition of the index by looking at mutual visibility conditions, and (2) adding additional indices, if necessary, to $F_{P'(i+1)}$ in order to make it a legitimate enclosing polygon.

Any enclosing polygon for either P' or Q' can be obtained by iterated modification (adding one vertex at a time) starting with the convex hull polygon. Notice that the cyclic order of the vertices in the convex hull polygons cannot be altered by the iterated procedure, since the procedure is merely an insertion procedure to the existing set. However, the points to be added to the enclosing polygon may be added at more than one position, although adding them at different positions will produce added visibility constraints for neighboring points.

The feasibility of going from $F_{P'i}$ and $F_{Q'i}$ to $F_{P'(i+1)}$ and $F_{Q'(i+1)}$ involves merging two lists of indices to maintain cyclic sequencing of the original two lists. Visibility constraints of the original lists assist in this list merging, since the resulting lists must have the combined constraints.

After the annular triangles are added legitimately, what was the interior enclosing polygon will become the outside frame, and the remaining interior vertices will be used to make a new enclosing polygon. Although the new outside frame is no longer necessarily convex, we believe a similar annular construction may be used, and identical visibility criteria may be applied.

It remains to show that selecting a legitimate triangulation for the ring of a particular enclosing polygon will not possibly lead to a dead-end later on, whereas a different ring triangulation would have worked. We are exploring an argument similar to that of Lemma 1 that says if a triangulation exists for which the enclosing rings are "deeper" than produced by the above F sequences, then there also exists a legitimate joint triangulation which corresponds to the F sequences.

BIBLIOGRAPHY

Corbett, J., 1979, Topological Principles in Cartography, Technical Paper 48, Bureau of the Census, Washington, DC.

Goodman, J.E., and Pollack, R., 1983, "Multidimensional Sorting," SIAM J. Comput., Vol 12, No. 3, pp. 484-507.

Guibas, Leonidas and J. Hershberger, D. Leven, M. Sharir, and R. Tarjan, 1986, "Linear Time Algorithms for Visibility and Shortest Path Problems Inside Simple Polygons," Proceedings of the Second Annual Symposium on Computational Geometry, pp 1-13.

Lefschetz, S., 1975, Applications of Algebraic Topology, New York, NY, Springer-Verlag.

Mount, David, 1986, personal communication, College Park, MD.

Preparata, Franco P., and Michael Ian Shamos, 1985, Computational Geometry, An Introduction, Springer-Verlag, New York, NY.

White, M.S. and P. Griffin, 1985, "Piecewise Linear Rubber-Sheet Map Transformation," The American Cartographer, vol 12-2, pp 123-131.

White, M. S., 1983, "Tribulations of Automated Cartography and How Mathematics Helps": Auto-Carto Six, V. 1, pp. 408-418, Ottawa, Canada, 6th International Symposium on Automated Cartography.