

A One-Step Crust and Skeleton Extraction Algorithm.

Christopher Gold
Chair in Geomatics
Laval University
Quebec City, Quebec, Canada

Jack Snoeyink
Department of Computer Science
University of British Columbia
Vancouver, BC, Canada

Abstract

We wish to extract the topology from scanned maps. In previous work [GNY96] this was done by extracting a skeleton from the Voronoi diagram, but this required vertex labelling and was only useable for polygon maps. We wished to take the crust algorithm of Amenta, Bern and Eppstein [ABE98] and modify it to extract the skeleton from unlabelled vertices. We find that by reducing the algorithm to a local test on the original Voronoi diagram we may extract both a crust and a skeleton simultaneously, using a variant of the Quad-Edge structure of [GS85]. We show that this crust has the properties of the original, and that the resulting skeleton has many practical uses. We illustrate the usefulness of the combined diagram with various cartographic applications.

1 Introduction and Previous Work

Workers in both Computational Geometry and Geomatics have discussed appropriate data structures for managing two-dimensional maps, for example [B86, F87, CDW92, H89, NME90, PC75, W95]. However, it is rarely noted that data input is the largest expense in Geographic Information Systems, and that the spatial data structure has a large influence on the cost of topological data input and structuring. See [G94a, G94b, GRR97] for a discussion of traditional spaghetti digitizing, and the advantages of a space-covering tiling - e.g. the Voronoi diagram - for the creation of topologically connected maps. [OBS92] summarizes many Voronoi applications.

In current practice, the creation of the correct topological connectivity is of considerably greater difficulty than the mere acquisition of the appropriate x-y coordinates. In many cases arcs are digitized by hand, and the polygon construction from these arcs involves many complexities of node determination, arc ordering around nodes, etc. (For a good discussion of current practice see [P94]). [GNY96] discussed simplifying the polygon mapping process with the introduction of simple Voronoi diagrams, and [G98] extended this to the case where the original paper documents were of sufficient quality that topologically correct maps could be generated directly from scanned input.

The Voronoi diagram is implemented using the Quad-Edge structure and an incremental Voronoi

algorithm that closely follows that given in [GS85] among others. In the rapid digitizing procedure of [GNY96], the operator digitizes the interior of each polygon, inserting points around the perimeter, each with the polygon label. The Delaunay triangulation/Voronoi diagram is then constructed, and the triangulation is traversed once, using the visibility order described in [GNY96] to extract the Voronoi boundaries between points with differing labels. The result is a set of arcs (polygon boundaries) with guaranteed topological connectedness.

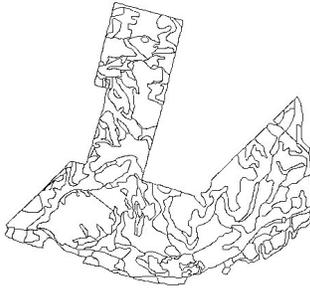


Figure 1: A forest map.

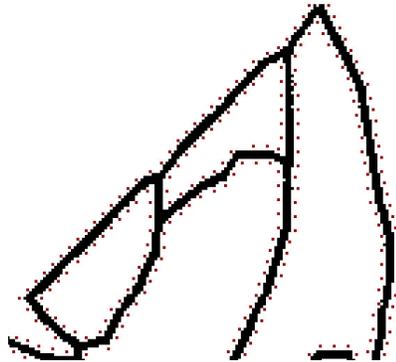


Figure 2: Fringe points.

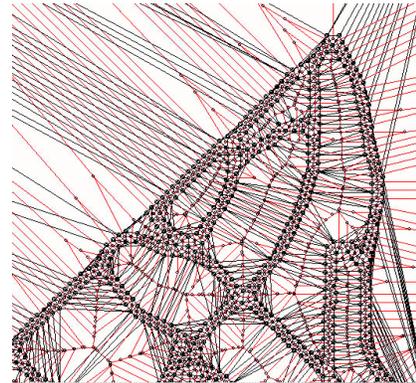


Figure 3: Voronoi/Delaunay tessellation.

The approach involves the generation of labelled “fringe” points just inside each of the polygon boundaries, generating the Voronoi diagram, and extracting only those Voronoi edges that separated fringe points with differing labels (and thus which come from different polygons). The “skeleton” of the black-line boundaries thus forms the topologically connected map. Figures 1 to 5 illustrate this process - the small “dumbbells” represent the Quad-Edge/Quad-Arc topology. The difficulties with this approach are that the fringe points have to be labelled, and that only polygon-type maps can be produced.

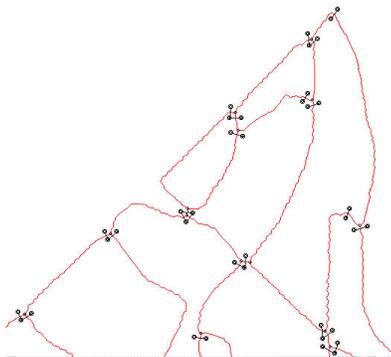


Figure 4: Extracted arcs.

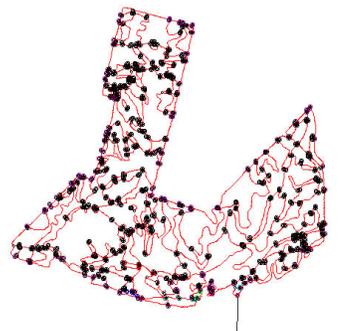


Figure 5: Topologically complete map.

These limitations are overcome due to the work of Amenta, Bern and Eppstein [ABE98], who showed that the “crust” of a curve or polygon boundary can be extracted from unstructured (and unlabelled) input data points if the original curve is sufficiently well sampled. Their intuition was

that, as the vertices of the Voronoi diagram approximate the medial axis (or skeleton) of a set of sample points from a smooth curve (Figure 6, after [ABE98]) then by inserting the original vertices plus the Voronoi vertices into a Delaunay triangulation (Figure 7, after [ABE98]) the circumcircles of this new triangulation approximate empty circles between the original smooth curve and its medial axis. Thus any Delaunay edge connecting a pair of the original sample points forms a portion of the sampled curve - the “crust”. In subsequent papers [AB98, ABK98] they extended this to three dimensions, extracting triangulations of the surface based purely on the x,y,z coordinates of surface sample points.

This solves admirably the problem of extracting the crust - but in our work on scanned maps we wished to extract the skeleton between our rows of fringe points. Experimentation with the crust algorithm showed that, while crust edges (connecting pairs of original sample points) are

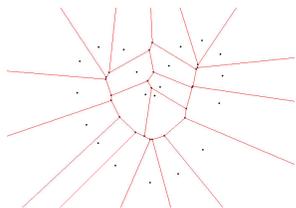


Figure 6: Voronoi diagram of sample points on curve.

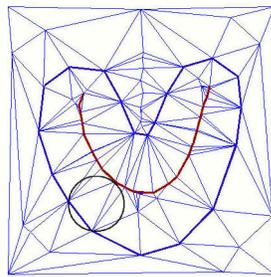


Figure 7: Skeleton and crust segment.

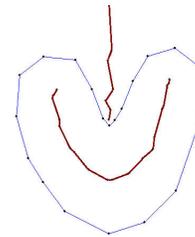


Figure 8: Crust and skeleton extracted together.

extracted correctly, connecting pairs of Voronoi vertices do not necessarily produce a good approximation to the skeleton, as in the construction of the second Delaunay triangulation additional edges are added between Voronoi vertices - see Figures 9 and 10. Our objective was to extract both the crust and the skeleton, in order to process various types of map input.

2 The One-Step Algorithm

Further consideration led to the examination of the relationship between Voronoi edges and Delaunay edges in the original Voronoi/Delaunay construction. This was made simpler by the use of the Quad-Edge data structure [GS85], where two of the pointers refer to Delaunay vertices, and two to the dual Voronoi vertices. Our intuition was to apply the crust test to individual Quad-Edges on the original diagram, rather than creating a second structure. The circle in Figure 7 above, for example, contains a crust edge but no skeleton edges. Thus, for each Quad-Edge, we wished to determine if the Delaunay edge had a circle that was empty of Voronoi vertices (and hence of a portion of the skeleton), and which intersected the Delaunay edge.

Each Delaunay edge is adjacent to two triangles whose circumcircles are centred at Voronoi vertices. In fact, the Voronoi edge between these circle centres is the dual to the Delaunay edge. The original crust essentially tests each Delaunay edge to make sure that it has a circle that contains the edge, but does not contain any Voronoi vertices. We make this a local test - testing only the two Voronoi vertices that are the endpoints of the dual Voronoi edge.

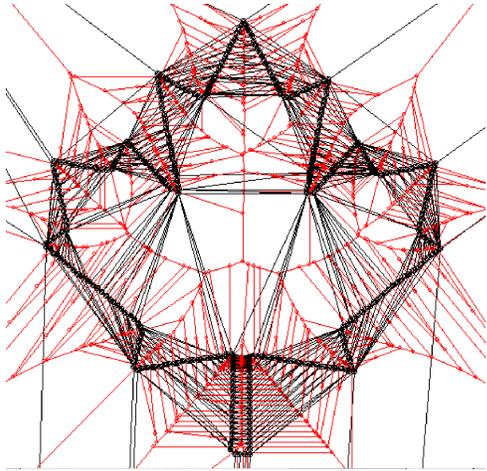


Figure 9: Maple leaf.

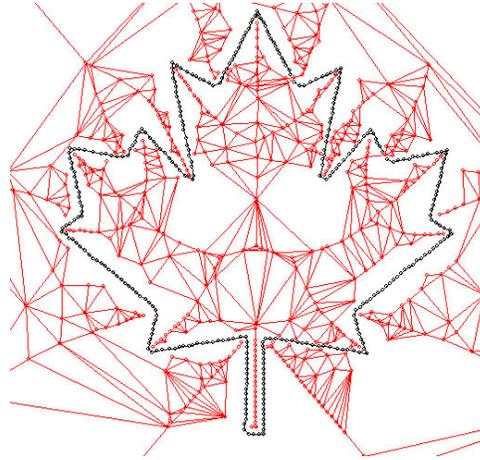


Figure 10: Crust and residual edges - the two-step approach.

For a Delaunay edge (q, r) with dual Voronoi edge (a, b) the test is the standard *InCircle* test [GS85] on these four points. $InCircle(q, r, a, b)$ determines whether point b is outside the circle through (q, r, a) , assuming that this circle is oriented counterclockwise. If so, then that circle shows that edge (q, r) should be included in the crust. On the other hand, if b is inside, then every circle containing (q, r) also contains a or b , and (q, r) is not included in the crust. If the circle through (q, r, a) is oriented clockwise, then the sense of the test is switched. This does not need to be handled as a special case, however, because when it occurs the Voronoi edge (a, b) does not intersect the Delaunay edge (q, r) . The same test will correctly indicate that (q, r) should not be included in the crust.

The standard $InCircle(q, r, a, b)$ test evaluates the sign of a four by four determinant whose row for each point p is $[1, px, py, px*px+py*py]$. An alternate expression of this test is useful if the original data are the Delaunay vertices or sites, and the Voronoi vertex coordinates are derived. Suppose that (p, q, r) and (r, q, s) are the two triangles incident on Delaunay edge (q, r) , and let v be the vector 90 degrees clockwise from $(r-q)$. Then the test

$$(s-q).(s-r) * (p-q).(p-r) \geq - (s-r).v * (p-q).v$$

will be true if and only if the edge (q, r) should be in the crust. This shows that the test is a polynomial of degree 4 in the input variables. Thus, this local test can be evaluated exactly with less precision than the global test, which requires evaluation of polynomials of degree 12 in the input coordinates.

A second idea was that Quad-Edges that failed the crust criterion were part of the skeleton or “anti-crust”. This term was mentioned briefly in the conclusions of [ABE98], citing [RCGH92, BA92, O94]. This is based on the idea that the dual of a crust edge is a Voronoi edge that intersects the crust - and has been rejected. The remaining Voronoi edges form a “tree” structure that extends towards the crust but does not cross it. (Indeed, with the Quad-Edge structure, each leaf of the skeleton points towards a particular crust segment.) The one-step algorithm consists of *assigning* Quad-Edges either to the “crust” or the skeleton (“anti-crust”), instead of constructing

a second diagram. The results for Figure 6 are shown in Figure 8, and the results for Figure 9 are shown in Figures 11 and 12. Mis-assignments occur where the sampling conditions of [ABE98] are not met - especially at acute angles.

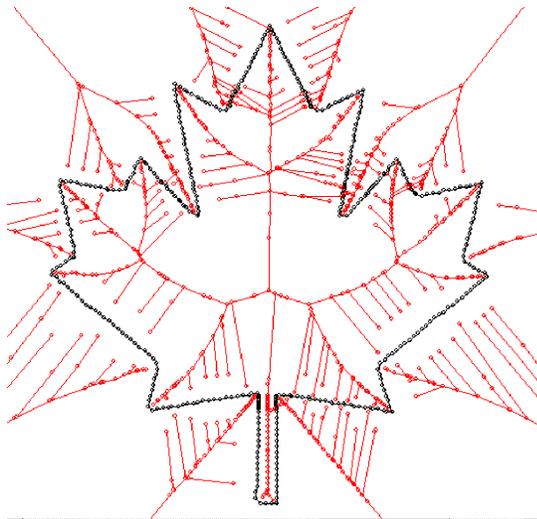


Figure 11: One-step crust and skeleton.

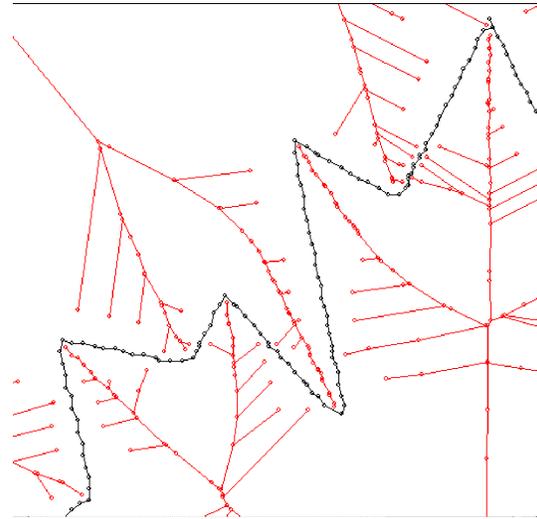


Figure 12: Enlargement of part of Figure 11.

The leaf vertices or "hairs" on this skeleton exist where there are three adjacent sample points whose circumcircle contains no other sample point - either at the end of a major branch of the skeleton, or where there is a minor perturbation of the sample points. These reach out to every minimum of curvature [AS95]. For a true curve, not a sampled curve, they would only occur at the end of major branches, as the skeleton is formed wherever a circle can touch two (not three) points on the curve.

3 Relationship of the Local "One-Step" Algorithm to the Crust Algorithm

Figure 13 illustrates that selecting crust edges based on a local property instead of a global property may allow additional edges to be added to the crust. In this figure, the heavy black edge e is in the locally-defined crust because there are circles that contain e but do not contain the endpoints of the Voronoi edge that is dual to e . On the other hand, e is not in the globally-defined crust because every circle that contains e does include some Voronoi vertex. We can observe, however, that when the sampling condition of Amenta, Bern, and Epstein [ABE98] is satisfied, then the locally-defined crust is identical to the globally-defined crust. In addition, we improve their analysis and relax the sampling conditions.

Amenta et al. define the *Local Feature Size*, $LFS(p)$, for a point p on a smooth curve to be the distance from p to a closest point on the medial axis of the curve. A curve is r -sampled, for some $0 < r < 1$, if, for every point p on the curve, there is a sample point within $r * LFS(p)$ of p . They prove two theorems for the globally-defined crust of an r -sampled curve: First, when $r < 0.40$, every edge between adjacent samples is included in the crust. Second, when $r < 0.252$, no extra edges are included.

The first theorem immediately applies to the locally-defined crust: the local condition is less restrictive than the global, so that any edge that passes the global test will pass the local test. We improve the bound to $r < 0.5$ for the globally-defined crust and to $r < 0.62$ for the locally-defined crust (neither of which are best possible). We then show that the proof of the second theorem applies to the locally-defined crust, and improve the bound to $r < 0.52$. We use relationships between local feature size, Voronoi circles, and the curvature that were observed by Amenta et al [ABE98].

Lemma 1. [ABE98] For any point p on a smooth curve, the two circles tangent to p of radius $LFS(p)$ contain no points of the curve.

Proof: If any circle C tangent to p does contain points of the curve, we can reduce the radius of C and maintain tangency at p until C becomes empty of points of the curve. Then the centre of C is a point on the medial axis of the curve, and $LFS(p) \leq$ the radius of C . QED

Theorem 1. The crust of an r -sampled smooth curve contains an edge between every pair of sample points that are adjacent on the curve, for $r < 0.5$. For the locally-defined crust, this can be improved to $r < 0.62$.

Proof: Consider a pair s, t of adjacent sample points and let p be the point on the curve between them that is equidistant from s and t . Assume that $LFS(p) = 1$; this means that there is a unit circle B' around p that contains no point of the medial axis of the curve. By Lemma 1, it also means that the two circles C and C' of unit radius that are tangent to the curve at p have no other intersections with the curve. These are illustrated in Figure 14.

Suppose, for the sake of deriving a contradiction that some vertex v of the Voronoi diagram lies inside the circle B of radius r around p . Without loss of generality, we can assume that v is above the curve and that s is the closer sample point. If we can show that the Voronoi circle V remains

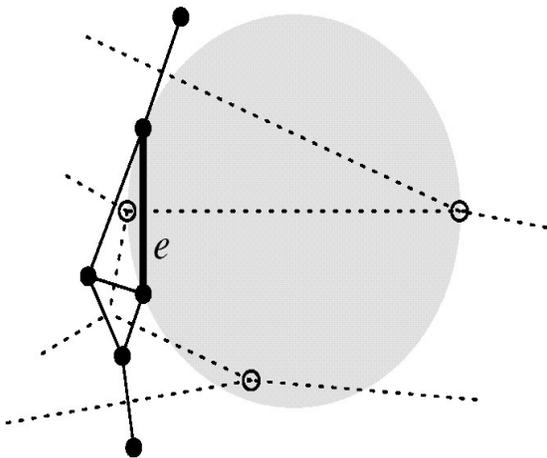


Figure 13: A set of six black points for which the locally-defined crust includes the heavy black edge, but the globally-defined crust does not.

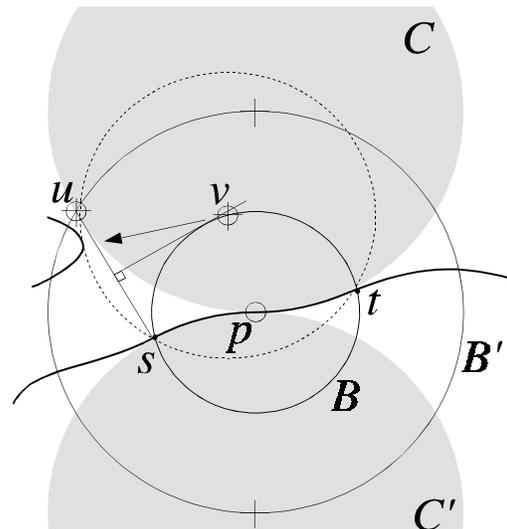


Figure 14: The curve remains outside of the circles C and C' tangent to p , so if there is a Voronoi vertex v in circle B , then there is a point of the medial axis in circle B' .

inside the union $B' \cup C$, then we can observe that there is a point inside B' that is on the medial axis of the curve - simply shrink the circle V about its centre v until it is tangent to the curve at some point q , then move v toward q and continue to shrink the circle V until V becomes tangent at a second point. Since v moves along a line segment whose ends are inside B' , this process terminates with v on a point of the medial axis inside B' .

We use the fact that both s and v lie inside B to show that the Voronoi circle V remains inside the union $B' \cup C$ when $r < 0.5$. For a fixed v , point s can be moved downward to increase the distance $d(s,v)$ until s is at the intersection of circles B and C' . The closest point to v of the boundary of $B' \cup C$ is the point u at intersection of the two circles. The quantity $d = d(s, u)/2$ can be calculated by noting that $s = ((r^4/4 - r^2)^{-1/2}, -r^2/2)$ and $u = (-\sqrt{3}/2, 1/2)$. The circle V first touches the perpendicular bisector of su when $1 - (r + d)^2 = r^2 - (r - d)^2$, and is below for all values of $r < 0.50720648$.

For the locally-defined crust, we need only consider the two Voronoi vertices defined by s , t , and one other point. This allows us to improve the bound to $r < 0.62$ as follows: Consider a Voronoi vertex v , equidistant from s and t , that is inside B' and whose Voronoi circle V remains within $(2 - r)$ of p . Shrink V around v as long as V intersects the curve in more than one connected component. When V can no longer shrink, then it is tangent to the curve at some point q ; as before, move v toward q until V becomes tangent at a second point and v is on the medial axis. We claim that this medial axis point is inside B ; either q is inside B as before, or q is outside but within $(1 - r)$ of B , and the second point of tangency is inside B' , and at least $(1 - r)$ inside the boundary of B . In either case, v ends inside B . QED

Theorem 2. The locally- or globally-defined crust of an r -sampled smooth curve does not contain an edge between non-adjacent sample points, for $r < 0.42$.

Proof: The idea is to show that for non-adjacent sample points s and t to be joined by an edge, we must have $r \geq d(s,p)/LFS(p)$, and that this lower bound is minimized by the configuration illustrated in Figure 15. Thus if we choose $r > 0.4208614 > \sqrt{(3+\sqrt{7})/2}$, we can be assured that s and t are not joined. We sketch the proof, which is best illustrated with a system like Cabrii or Sketchpad.

Let s and t be non-adjacent sample points, and assume, for the sake of deriving a contradiction, that there is a circle B that touches s and t , is empty of sample points, and does not contain the endpoints of the Voronoi edge that is the portion of the bisector of s and t . (Note that if B contains no sample points then there is a Voronoi edge between s and t ; the additional assumption is that its Voronoi vertices lie outside of B .) Since the medial axis of the curve intersects the line segment st , we may further assume that s and t are chosen such that the point of intersection is closer to endpoint s .

Let v and v' be the points in which circle B intersects the perpendicular bisector of st . Since v and v' lie on the Voronoi edge separating s and t , the circles V and V' through s with centres v and v' , respectively, are empty of sample points. Observe in Figure 16 that circles V and V' are

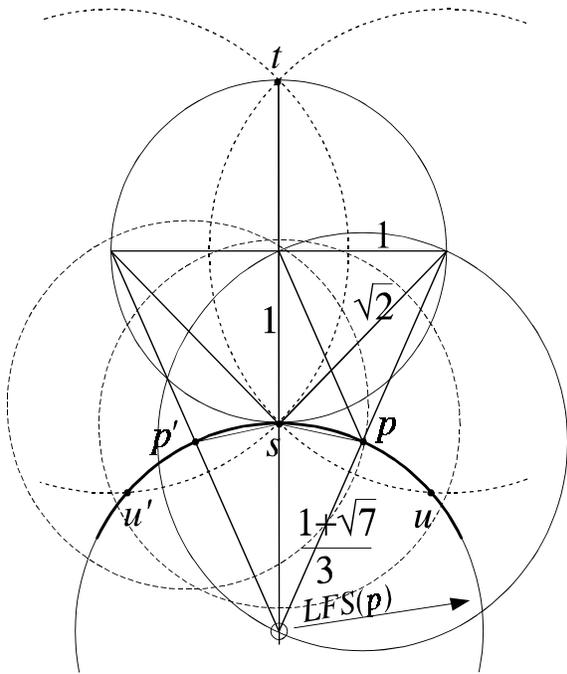


Figure 15: The configuration defining the bound for Theorem 2.

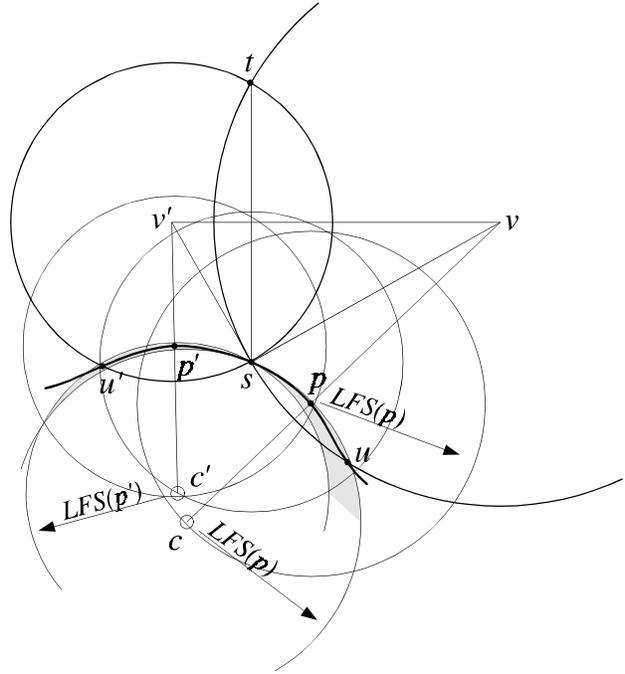


Figure 16: A general configuration for the proof of Theorem 2.

orthogonal. Point s has two adjacent sample points, u and u' , which must therefore be outside of V and V' . Let p be the point on the curve equidistant from s and u , and let p' be equidistant from s and u' .

Consider now two circles of radius $LFS(p)$. The first, centred at p , cannot contain a point of the medial axis - one implication is that $LFS(p) \leq$ the distance from p to the midpoint of st . For the second, start with the circle tangent at p and below the curve, and push it along the bisector of su , until it touches both s and u . Let the centre be at point c . This circle must contain p , and must not intersect the curve between s and u' , or the first circle would contain a point of the medial axis.

If we do the same for circles of radius $LFS(p')$, then we constrain how the curve between u and u' lies - p is inside the circle centred at c and outside the circle centred at c' , and p' is the opposite. Note that this implies that c is right of the line $c'p'$, and c' is right of the line pc .

If $LFS(p') \leq LFS(p)$, then we can move c' along the ray $v'p'$ and decrease the bound $d(s,p')/LFS(p')$ until the circle centred at c' hits p , or $LFS(p') = LFS(p)$. In the former case, there is a point of the medial axis inside the circle of radius $LFS(p)$ centred at p . In the latter, we can move c and c' until they meet on the circle of radius $LFS(p)$ centred at s , and the distances $d(s,p) = d(s,p')$. This occurs when line cs bisects the angle $\angle vcv'$. Calculations show that the bound on the ratio is minimized when st bisects vv' , which is the case in the configuration of Figure 15. QED

If we continue within the framework of [ABE98], in which a smooth closed curve has been r-

sampled according to “local feature size”, then the crust will be the simple polygon joining adjacent sample points on the curve: edges of the Delaunay triangulation that are not selected to be crust edges will complete triangulations of the interior and exterior of this simple polygon.

Consider now the Voronoi edges that are dual to the non-crust edges in the interior of the polygon. These form a binary tree; and since Voronoi edges that intersect crust edges are removed, this tree lies inside the simple polygon. This tree has frequently been used to approximate the medial axis of the polygon; it is not difficult to see that as the sampling density increases, it converges to a diagram that contains the medial axis (and some extra “hair”).

The dual of the exterior triangulation gives an approximation to the exterior medial axis. When multiple closed curves are sampled, this “anti-crust” will always be a set of one-dimensional representations of the topology of the regions between crust curves - each connected component will be a deformation retract of the region that it is in. For example, an annular region will have one cycle with branches (and hair) dangling from it. Thus, we are justified in calling this a “skeleton”.

4 Applications

4.1 Polygon generalization

Following our original objectives, we have processed a variety of polygons and scanned maps with the new algorithm. In the case of the simple maple-leaf polygon of Figure 11, both the crust and the skeleton are extracted essentially correctly. In the case of the crust, sharp corners fail to satisfy the crust sampling criterion and, as expected, an occasional Delaunay edge crosses the tip. In the case of the skeleton, the form is correct but some extraneous branches are generated where perturbations of the boundary data points are treated as incipient salients. This also is a well known situation. Indeed, the “hairs” on this skeleton form an excellent tool for curve generalization or simplification, as they represent minima of curvature, as described above. Simplification of the skeleton by removing individual hairs is achieved by perturbing or removing individual crust points so as to remove these minima of curvature. The resulting curve has a simple skeleton, and may be generalized in the sense of [OI90] or an equivalent.

4.2 Road or river network estimation from boundaries

The skeleton of the polygon of Figure 11 forms an excellent approximation of drainage network development in a homogeneous terrain, given the watershed. While obviously not conforming to the non-homogeneous reality, it may well be a useful approximation. An equivalent application is the design of the road network for access to all parts of a homogeneous terrain - perhaps for forest harvesting.

4.3 Hinterland or watershed estimation from road or river network

Figures 17, 18 and 19 show the reverse operation: a river network is represented as a set of sample points, and the one-step algorithm applied. The crust and the skeleton are recognized with

the one-step algorithm, and the connected components linked with the algorithm described in [GNY96] for rapid digitizing, and in [G98] for scanned maps using the Quad-Edge data structure. The estimated sub-basins may be clearly seen, and extracted, on the basis of the skeleton. This approach has been used successfully in British Columbia for preliminary watershed estimation, with later correction by direct observation. Notice the cases of inadequate sampling according to the [ABE98] criterion, especially at intersections, causing the skeleton to break through the river network. Similar results may be obtained if the construction is based on a transportation network. In both cases the catchment area may be estimated for any point on the network, permitting the development of flow or transportation capacity maps. Connectivity descriptions, such as simple Strahler numbers, may be extracted from the networks of Figures 11 or 18.

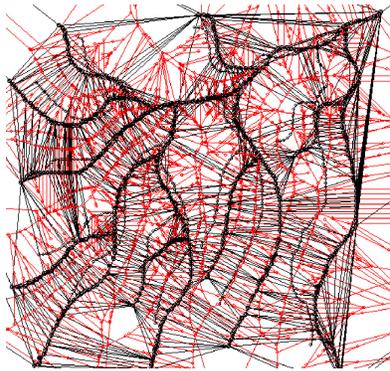


Figure 17: A river network.

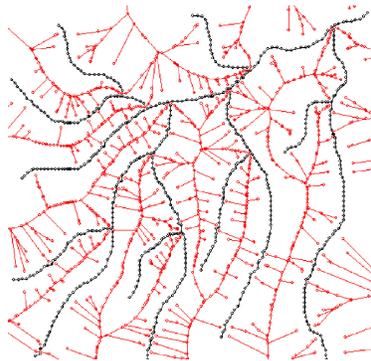


Figure 18: Estimated watersheds.

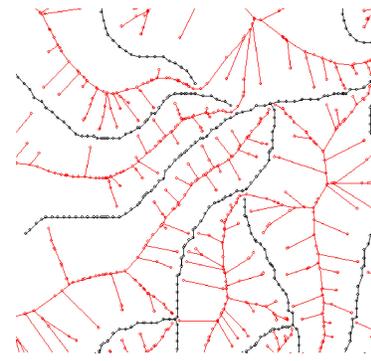


Figure 19: Enlargement of part of Figure 18.

4.4 Scanned maps

Scanned maps, similar to those shown in the first section, may be processed using the crust criterion instead of vertex colouring. Thus edge pixels for each black line are extracted, as previously, but they are not labelled by floodfilling polygons. Instead, the crust criterion is applied to extract the crust of each black/white boundary. As this is purely geometric, it is applicable to closed polygons, connected networks and unconnected black lines, such as text. These crusts separate the black/white portions of the scanned maps. However, the skeletons are generated by the same algorithm - but for the white portions of the map as well as for the black portions. If the black line work forms a connected graph of a polygon set, or a network, then the skeleton of this black region forms the connected topology representing the line work - as in [GNY96] and Figures 1 to 5. The skeletons of the white regions express the relationships between disconnected black line work, or else express the form of the white shapes, as in [B67]. The crusts may or may not be preserved. As in Figure 19, the crust may break if the sampling criterion of [ABE98] is not preserved - especially at sharp junctions. This will cause the “white” skeleton to break through the crust and connect with the “black” skeleton.

Several alternatives exist in this case: we may revert to the labelled vertex algorithm; we may improve the sampling; or we may take advantage of the fact that, as all Voronoi vertices fall on one side or the other of the line of the crust, they may each be labelled as “black” or “white” by reference back to the original image. Links are then broken between black and white vertices.

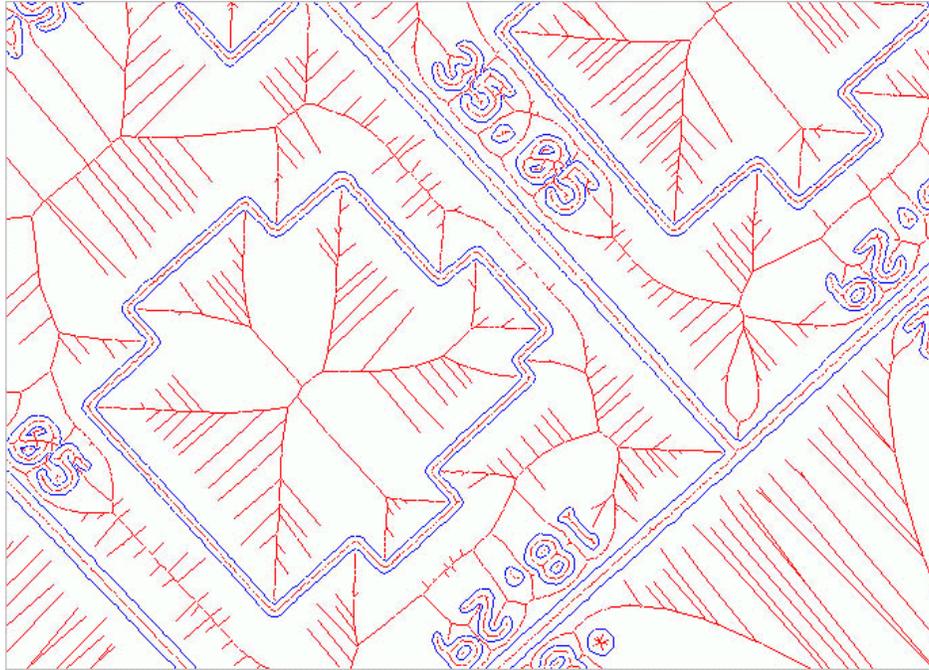


Figure 20: Crust and skeleton of part of a scanned cadastral map.

As shown in Figures 20 and 21, the skeleton generation algorithm allows for the extraction of the topology of the line work as well as for the detection and placement of text. Burge and Monagan [BM95a,b] also worked on the extraction of text from scanned maps, without the topology

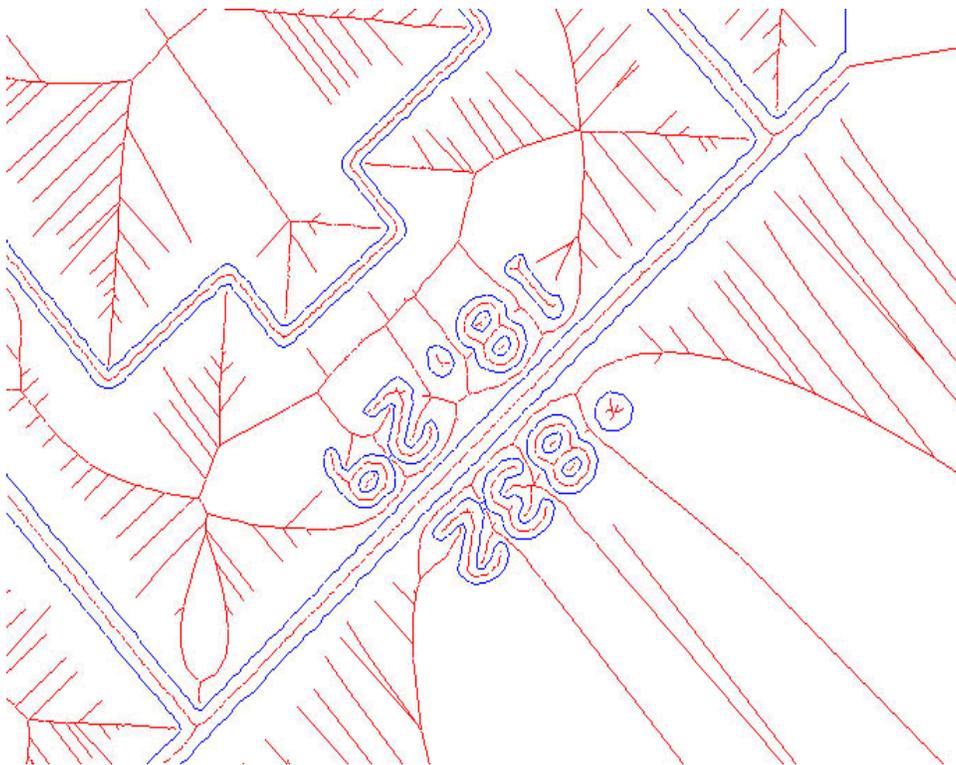


Figure 21: An enlargement of part of the scanned map of Figure 20.

emphasis. For scanned maps the crust is the boundary between black and white pixels, whereas there is a skeleton for the white regions (“polygons”) as well as for the black regions (line work). These last form the centrelines of the scanned line work: skeletonization in the Euclidean, rather than the raster, sense, and form a topologically complete graph of the original input map. It should be noted that the algorithm using labelled fringe points may do better for a simple polygon map, if sampling does not achieve the crust criterion, as more information is supplied about the desired boundary connectivity.

4.5 Processing scanned text

Scanned text may equally well be treated in the same way as in the scanned map example, after processing with an edge-detection filter, as shown in Figures 22 and 23. The interior skeleton of the character can detect the characteristic form of the connected graph, essentially for the cost of constructing the Voronoi diagram. The exterior skeleton may be used to express the relationships between the characters. These relationships are also useful for languages using non-connected components, or diacritical marks or accents. Burge and Monagan’s methods [BM95a,b] involved extracting the text from cadastral maps using Voronoi diagrams - adjacent pixels were given the same label, and the exterior skeletons were used to connect the letters together to form complete words. As in the original work by Blum [B67] the interior skeletons of arbitrary shapes may be used as a shape descriptor.

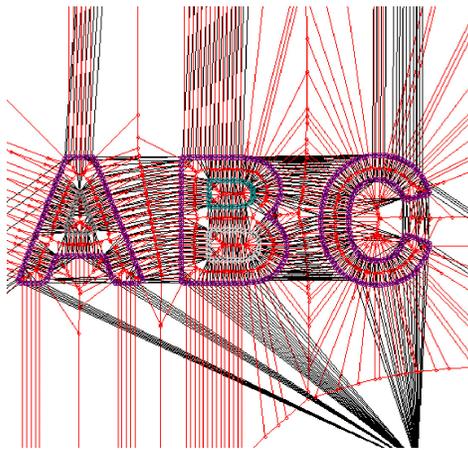


Figure 22: Character outlines.

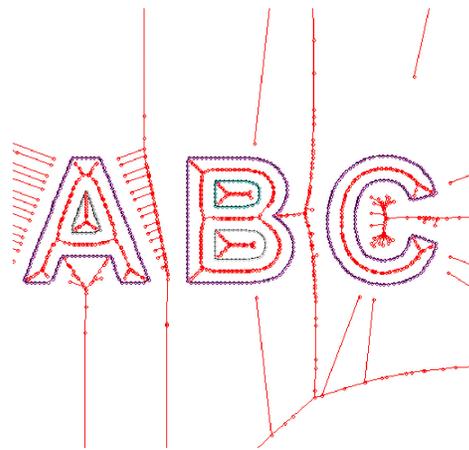


Figure 23: Interior and exterior skeletons.

5 Building the topological structure

5.1 Advantages of the Quad-Edge structure

So far we have been concerned with the assignment of individual Delaunay or Voronoi edges to the crust or the skeleton. Our original motivation was to generate topologically connected graphs from either or both of these, in order to produce topological maps for GIS-type applications.

(This has usually been a very elaborate process.) Thus we needed to extract connected arcs from the individual segments identified by the one-step algorithm, and to leave them connected to each other to form our desired topological map. This removes the need to move to expensive GIS systems to construct topologically connected polygon and network maps.

We needed a structure that was capable of managing both the primal and the dual graph with equal felicity, and that could apply to simple Delaunay/Voronoi structures as well as to polygon arcs. Since we were often working interchangeably with Delaunay/Voronoi representations, the Quad-Edge data structure of [GS85] became an obvious candidate. It has various attractions: it is a method for representing the edge connectedness of any connected graph on a manifold; it is symmetric in its storage of both the primal and the dual edge; it has a complete algebra associated with it, requiring no searches around polygons, nodes, etc.; and it is admirably adapted to an object-oriented programming environment. There are only two operations: Make-Edge to create a new edge, and Splice to connect or disconnect two ends. [G98] shows a half-page of code that implements the basic operations. To us it appears the most elegant structure. One limitation is that it works only for connected graphs. If there are "islands" in a polygon map, for example, and the arcs between the nodes in the data structure graph correspond to the polygon edges in the map, then the presence of the island within some other polygon can not be detected without further processing. The Voronoi diagram, being space-covering, does not have this problem.

Since the algebra for the Quad-Edge structure includes operations for moving from arc to adjacent arc, and for obtaining the vertices, we can navigate within this structure in several ways: We can traverse the whole graph using the classic depth-first or breadth-first searches [S83], which do not require any geometric information, or using the "Visibility Ordering" of a triangulation [GC87, FFNP91, GNY96, BKO97, SS97]. It is also possible to walk from some starting point to some destination point through edges of the triangulation using the "CCW" orientation test on vertices [GCR77, GS85] (although it should be noted that this algorithm is only guaranteed for Delaunay triangulations).

5.2 The Quad-Arc data structure

The original Voronoi structure contains far too many individual points to be used as an archival mechanism, or for internal processing of polygonal maps. It would be desirable as an alternative to remove the large number of irrelevant Voronoi edges without destroying the desired polygonal structure. This would be difficult with an underlying triangulation, but straightforward with the Quad-Edge structure. For our scanned map application, we use a modification of the "Visibility Ordering" traversal implementation of [GNY96] to inspect all edges and remove those defined as "crust" rather than "skeleton". We show a small example; the method has been described in [G98].

Figure 24 shows a small triangulation represented as a Quad-Edge data structure. Figure 25 shows the same structure with two Quad-Edges removed because they were classified as "crust" rather than "skeleton". Figure 26 shows the final result of this process. The graph that remains consists of strings of Quad-Edges ("arcs"), connected at the ends ("nodes"). The graph may be traversed again, using a depth-first search or equivalent, to remove the multiple links while

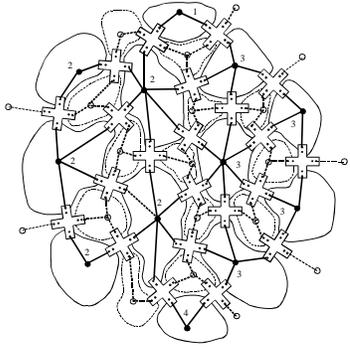


Figure 24: Initial triangulation.

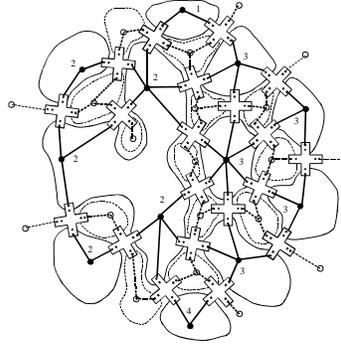


Figure 25: Two edges removed.

preserving the nodes, as shown in Figure 27, so that each of the Quad-Edges indicates an “arc” or polygonal line that joins two nodes rather than indicating a segment. We have called the result of this a Quad-Arc data structure, to emphasize its similarity to the arcs used in a traditional GIS to define the boundary between two polygons. Its topological properties are unchanged from the original Quad-Edge of [GS85].

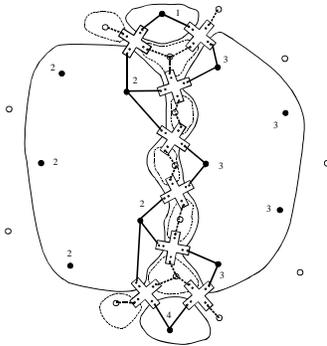


Figure 26: End of edge removal.

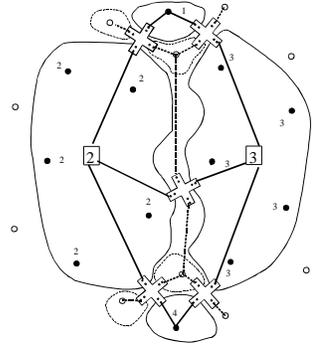


Figure 27: Final Quad-Arc.

In the original Voronoi structure the positions of islands or other unconnected features within polygons is clearly defined by the adjacency of the Voronoi cells representing the island with those representing the enclosing polygon. The arc extraction process loses this information because all connecting edges are deleted, and the Quad-Edge algebra applies only to connected graphs. By keeping track of the connected components as edges are deleted, however, it is possible to identify the formation of islands and preserve a dummy arc, flagged as such, to keep the link between the island and the containing polygon.

5.3 Advantages of the Quad-Arc structure

Once the structure is completed, the resulting topological structure may be used to perform polygon shading, neighbour detection, etc., as with any mapping program. Due to the overall simplicity of the algorithm, the method may be used in a wide variety of applications. Current work is on a simple mapping package where hand drawn maps may be scanned into a PC with an

economical scanning unit, and the resulting topologically complete map is available in a very few minutes. This may then be exported to a commercial mapping program or GIS, or else preserved internally for simple mapping projects. We have not studied the editing process in great depth yet, as we envisage that it would be simpler to modify the paper map and re-process it.

The original motivation for this work was to take advantage of our scanned mapping technique. This started with the complete Voronoi structure, which guarantees topological completeness of the processed map, but uses large amounts of storage. Using the Quad-Edge representation of the Voronoi structure permits the elimination of unwanted edges, and the Quad-Edge to Quad-Arc simplification provides even greater storage efficiency. The resulting structure is then equivalent to other polygon-management data structures, in that if there are many intermediate points along the arcs, then the number of pointers contributes little to the storage costs.

Applying line simplification to arcs gives a way to further reduce the amount of data while still preserving the topology. When arcs are simplified, they must be checked for intersections with neighbouring arcs, but these neighbours can be readily identified from the topology.

The resulting data structure cleanly separates the topology from the attributes, in that the pointer structure is entirely within the topological elements: four pointers between the branches of each Quad-Arc in the object-oriented version, four pointers to adjacent Quad-Arcs, and four pointers to attribute information. These are traditionally the two faces and the two nodes associated with the Quad-Edge, but in the Quad-Arc the face pointers point directly to polygon attribute information, perhaps in a database, and the two node pointers may be used to reference the associated arc (chain of x-y coordinates) or other arc information as desired.

Thus only one topology file is required, along with a geometry (arc) file and a polygon attribute storage mechanism. Islands may be integrated with dummy Quad-Arcs. The topology permits any connected graph (not only polygons), thus allowing the maintenance of network graphs. Independent points may be handled in the same way, considering the edges of the Delaunay triangulation to be dummy Quad-Arcs.

Queries in this structure are based on the edge-algebra developed by [GS85]. Finding the boundaries of a polygon for example requires successive calls to *Oprev* (or to *Onext* to go round in a clockwise direction), until one returns to the starting arc. To find adjacent polygons at the same time involves looking at the dual branches of each Quad-Arc at each step. Since travelling from any Quad-Arc to any other involves only a sequence of 'Rot' and 'Next' pointers, these may be encoded as a binary string for any path within the map.

The Quad-Edge or Quad-Arc data structure permits topology maintenance in a mapping system with a minimum of complexity. It fits very easily with current scanned map software using the Voronoi model for data input and topological structuring, and permits straightforward topology maintenance and spatial queries after simplification. Most of all, it is simple to implement in an object-oriented environment, reducing some of the complexity and mystery of cartographic topology.

6 Other applications

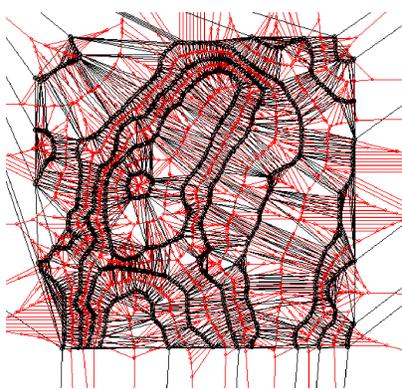


Figure 28: A contour map.

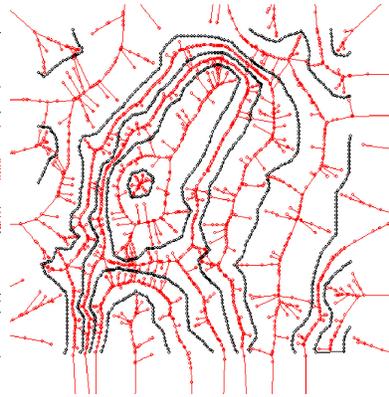


Figure 29: Crust and skeleton.

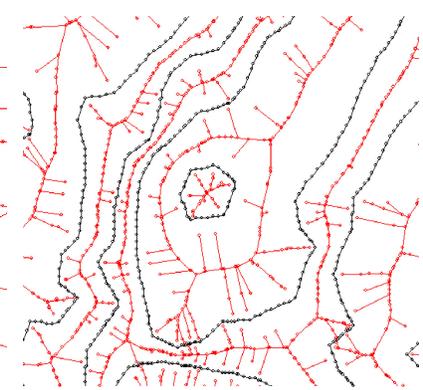


Figure 30: Enlargement of part of Figure 29.

While many other applications of the one-step crust algorithm in mapping come to mind, one more must suffice - but this time without the necessity for Quad-Arc extraction. Another source of map data comes from contour maps, and terrain elevation models are in great demand, especially triangulated elevation models. However, it is not always easy to extract good triangle sets from contours alone - and manually derived contours from stereo images are readily available. The combination of the crust (contours) and the skeleton assists this process considerably, as the skeleton expresses the relationships between contours or between contour segments. Figures 28, 29 and 30 show this well. In general, the skeleton forms a medial contour, but where contour re-entrants occur, branches of the skeleton indicate that the associated dual Delaunay edge would have connected vertices of the same elevation - thus giving an implausible horizontal triangular plate. Various solutions are possible, including the treatment of these branches as gullies, having elevations varying between the medial axis and the associated contour line. This follows closely the interpretation that would be made by the experienced map user. Cleanup of the “hair” could first be performed as described previously. This gives a very powerful tool for interpreting the terrain on the basis of the relationships between contours, and not merely by individual contours themselves. The Voronoi/Delaunay relationships are particularly useful here, as there is zero slope along contour lines defined by Delaunay (crust) edges, and thus the direction of runoff or maximum downhill slope is perpendicular to this - exactly as expressed by the associated dual Voronoi edge. These interrelationships readily permit the reconstruction of a meaningful terrain model, as they mimic the natural processes involved.

7 Conclusions and Future Work

We have shown that a simple one-step algorithm may generate the crust and anti-crust simultaneously, and that these may be extracted to form topologically structured maps. The results are equivalent to those of the crust algorithm of [ABE98]. We have shown how this resolves a variety of issues in map input and analysis, and we expect to address individual applications in more detail in the near future.

In two dimensions the crust is a set of line segments, as is the skeleton. [AB98, ABK98] describe the case for three-dimensional crust extraction, which is similar to that in two dimensions. However, they note that certain Voronoi vertices do not lie on the skeleton, and therefore need to be excluded in their algorithm. Our one-step algorithm also applies here: the choice for a simplex-pair is either a triangular face assigned to the crust, or a line segment assigned to the skeleton. We have not yet completed our exploration of this modification of the crust approach.

The Quad-Arc approach may be of significant use in Geomatics as it provides a simple data structure that suits the needs of the data input process, preserves the map topology, and may be implemented within small-scale PC-based software for topological map querying. The combination of the crust criterion, the one-step algorithm and the Quad-Arc data structure are sufficient to put cartographic topology within the range of simple mapping programs.

Acknowledgments

Support for this research was received from the Natural Sciences and Engineering Research Council of Canada and the Networks of Centres of Excellence Program, under the IRIS network, and from the Industrial Research Chair in Geomatics at Laval University. David Thibault prepared many of the diagrams and assisted with the programming. Mir Mostafavi, Valerie Ramsumeer, Darka Mioc, Francois Anton, Weiping Yang and Luc Dubois assisted at various stages of the development of the underlying programs.

References

- [AS95] Alt, H. and O. Schwarzkopf, 1995. The Voronoi diagram of curved objects. Proc. 11th Annual ACM Symposium on Computational Geometry, pp. 89-97.
- [AB98] Amenta, N. and M. Bern, 1998. Surface reconstruction by Voronoi filtering. Proceedings, 14th Annual ACM Symposium on Computational Geometry, pages 39-48.
- [ABE98] Amenta, N., M. Bern and D. Eppstein, 1998. The crust and the beta-skeleton: combinatorial curve reconstruction. Graphical Models and Image Processing, 60/2:2, pp. 125-135.
- [ABK98] Amenta, N., M. Bern and M. Kamvysselis, 1998. A new Voronoi-based surface reconstruction algorithm. Proceedings, Siggraph '98, pp. 415-421.
- [B67] Blum, H., 1967. A transformation for extracting new descriptors of shape. In: W. Whaten-Dunn (ed.), Models for the Perception of Speech and Visual Form. MIT Press, Cambridge, Mass., pp. 153-171.
- [BA92] Brandt, J. and V.R. Algazi, 1992. Continuous skeleton computation by Voronoi diagram. Computer Vision, Graphics and Image Processing, v. 55, pp. 329-338.
- [BKO97] de Berg, M., M. van Kreveld, R. van Oostrum, M. Overmars, 1997. Simple traversal of a subdivision without extra storage. International Journal on Geographical Information Systems, v. 11, pp. 359--373.
- [BM95a] M. Burge and G. Monagan, 1995a. Using the Voronoi tessellation for grouping words and multi-part symbols in documents. Proc. Vision Geometry IV, SPIE Vol. 2573, (San Diego, California), pp. 116-123.
- [BM95b] M. Burge and G. Monagan, 1995b. Extracting Words and Multi-part Symbols in

- Graphics Rich Documents. Proc. 8th International Conference on Image Analysis and Processing ICIAP, (San Remo, Italy), LNCS, Springer Verlag, pp. 533-538.
- [B86] Burrough, P.A. 1986. Principles of Geographical Information Systems for Land Resources Assessment. Oxford University Press, New York.
- [CDW92] Chrisman, N.R., J.R. Dougenik and D. White, 1992. Lessons for the design of polygon overlay processing from the Odyssey Whirlpool algorithm. Proc. 5th International Symposium on Spatial Data Handling, pp. 401-409.
- [F87] Frank, A. 1987. Overlay processing in spatial information systems. Proc. Auto-Carto 8, pp. 16-31.
- [FFNP91] De Floriani, L., B. Falcidieno, G. Nagy and C. Pienovi, 1991. On sorting triangles in a Delaunay tessellation. *Algorithmica*, v. 6, pp. 522-532.
- [G98] Gold, C.M., 1998. The Quad-Arc data structure. In: Poiker, T.K. and Chrisman, N.R. (eds.). Proceedings, 8th International Symposium on Spatial Data Handling; Vancouver, BC, 1998. pp. 713-724.
- [G97] Gold, C.M., 1997. Simple topology generation from scanned maps. Proceedings, Auto-Carto 13, ACM/ASPRS, Seattle, April, pp. 337-346.
- [GCR77] Gold, C.M., T.D. Charters and J. Ramsden, 1977. Automated contour mapping using triangular element data structures and an interpolant over each triangular domain. *Computer Graphics*, v. 5, pp. 170-175.
- [GC87] Gold, C.M. and S. Cormack, 1987. Spatially ordered networks and topographic reconstructions. *International Journal of Geographical Information Systems*, v. 1, pp. 137-148.
- [GRR97] Gold, C.M., P.R. Remmele and T. Roos, 1997. Voronoi methods in GIS. In: Van Kreveld, M., Nievergeld, J., Roos, T. and Widmeyer, P. (eds.), "Algorithmic Foundations of GIS. Lecture Notes in Computer Science No. 1340", Springer-Verlag, Berlin, Germany, pp. 21-35.
- [GNY96] Gold, C.M., J. Nantel and W. Yang, 1996. Outside-in: an alternative approach to forest map digitizing. *International Journal of Geographical Information Systems*, v. 10, no. 3, pp. 291-310.
- [G94a] Gold, C.M., 1994a. Dynamic data structures: the interactive map. In: Molenaar, M. and de Hoop, S. (eds.), "Advanced Geographic Data Modelling - Spatial Data Modelling and Query Languages for 2D and 3D Applications", Netherlands Geodetic Commission Publications on Geodesy (New Series), pp. 121-128.
- [G94b] Gold, C.M., 1994b. Three approaches to automated topology, and how computational geometry helps. In: Waugh, T.C. and Healey, R.G. (eds.). Proceedings, Sixth International Symposium on Spatial Data Handling; Edinburgh, Scotland, 1994. pp. 145-158.
- [G92] Gold, C.M., 1992. The meaning of "neighbour". In: "Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Lecture Notes in Computing Science No. 639", Springer-Verlag, Berlin, Germany, pp. 220-235.
- [GS85] Guibas, L. and J. Stolfi, 1985. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *Transactions on Graphics*, v. 4, pp. 74-123.
- [GSS89] Guibas L., D. Salesin and J. Stolfi, 1989. Epsilon geometry: building robust algorithms from imprecise computations. Proc. 5th Annual ACM Symposium on Computational Geometry, pp. 208-217.
- [H89] Herring, J., 1989. A fully integrated geographical information system. Proc. Auto-Carto 9,

- pp. 828-837.
- [NME90] Nagy, G., M. Mukherjee and D. Embly, 1990. Making do with finite numerical precision in spatial data structures. Proc. 4th International Symposium on Spatial Data Handling, pp. 55-65.
- [O94] Ogniewicz, R.L. 1994. Skeleton-space: a multiscale shape description combining region and boundary information. Proceedings of Computer Vision and Pattern Recognition, 1994, pp. 746-751.
- [OI90] Ogniewicz, R. and M. Ilg, 1990. Skeletons with Euclidean metric and correct topology and their application in object recognition and document analysis. Proceedings, 4th International Symposium on Spatial Data Handling, v. 1, pp. 15-24.
- [OBS92] Okabe, A., B. Boots and K. Sugihara, 1992 Spatial Tessellations - Concepts and Applications of Voronoi Diagrams. John Wiley and Sons, Chichester.
- [PC75] Peucker, T.K. and N.R. Chrisman, 1975. Cartographic data structures. American Cartographer, v. 2, pp. 55-69.
- [P94] Pullar, D. 1994. A tractable approach to map overlay. Ph.D. Thesis, University of Maine, Orono.
- [RCGH92] Robinson, G.P., A.C.F. Colchester, L.D. Griffin and D.J. Hawkes, 1992. Integrated skeleton and boundary shape representation for medical image interpretation. Proc., European Conference on Computer Vision, 1992, pp. 725-729.
- [S83] Sedgewick, R., 1983. Algorithms. Addison-Wesley, Reading, Mass., 551 p.
- [SS97] Speckmann, B., J.S. Snoeyink, 1997. Easy Triangle Strips for TIN Terrain Models. Proc. 9th Canadian Conference on Computational Geometry, pp. 239-244.
- [W95] Worboys, M.F., 1995. GIS: A Computing Perspective. Taylor and Francis, London.