BUREAU OF THE CENSUS
STATISTICAL RESEARCH DIVISION REPORT SERIES
SRD Research Report Number: Census/SRD/RR-87/32

Triangulated Data Structures For Map Merging And Other Applications in Geographic Information Systems

by

Alan Saalfeld
Statistical Research Division
Bureau of the Census
Washington, D.C. 20233

Recommended by: R. T. O'Reagan

Report completed: September 1, 1986

Report issued: October 20, 1987

TRIANGULATED DATA STRUCTURES FOR MAP MERGING AND OTHER APPLICATIONS IN GEOGRAPHIC INFORMATION SYSTEMS

Alan Saalfeld

Statistical Research Division
Bureau of the Census
Washington, DC 20233

ABSTRACT

The partitioning of a single space or map region into triangles or triangular regions has proved to be very useful in numerous diverse applications of geographic information systems. Triangulated irregular networks (TINs) have been successfully used in applications in elevation modeling and representation, in network analysis and routing problems, in land use and hydrologic studies, in roadway design and landscape visualization, in nearest-neighbor search and zoning problems, and in many other searching, sorting, and spatial data organization problems.

This paper examines some of the key properties of triangulations in general, the Delaunay triangulation in particular, and especially focuses on those properties which permit easy manipulation and comparison of two data sets that have been organized into similarly triangulated data structures. Tools for "navigating through" triangulated data sets and for decomposing the triangles or for building additional structures upon the triangles are shown to behave well under transformations of data sets. These transformations include joint triangulations, the *simultaneous* decomposition of *two* spaces into triangular regions, which has the potential for providing additional TIN applications in the two-space case.

The growing need to overlay two maps (of near-identical or different coverages) and to match or compare point features on those two maps will make joint triangulation techniques more important and more useful every day. Compatible TIN coverages on both spaces permit straightforward comparisons between the spaces and provide a sound mathematical foundation for those comparisons.

1. INTRODUCTION

The triangulation of a planar domain—partitioning it into a finite family of triangular regions—provides a key tool for automating the interplay among (1) finite sets, (2) finite combinatorial topology, and (3) two-dimensional, infinite-set, continuous geometry/topology. A triangulation starts with a finite set of points (which will become triangle vertices), then builds a finite collection of line segments or triangle edges (recorded as point pairs), another finite collection of triangles (recorded or stored as point triples), and finite collections (implicit or recorded explicitly) of triangle adjacency relations and triangle edge inclusion relations. The triangulation begins and ends with finite sets; however, the resulting finite collection of triangles, edges, and vertices account for all of the infinitely many points of the continuous two-space that has been triangulated. A triangulation partitions* infinitely many points of space into finitely many managable sets of points, edges, and triangles, each of which requires only fixed storage to describe completely. These finitely many sets are all well behaved, well studied, and well understood, and their topological and geometric structure lends itself nicely to computerization.

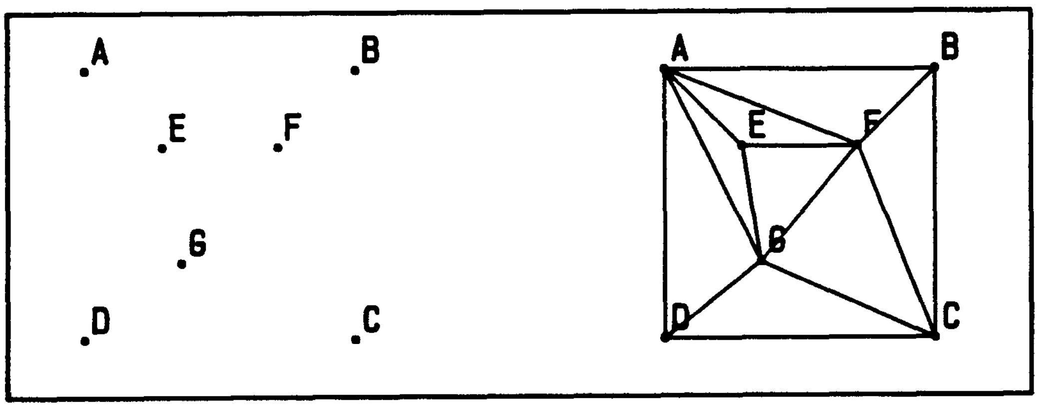

Figure 1. A set of points and a triangulation of those points

The resulting lists used for recording the triangulation are:

Points:
{A, B, C, D, E, F, G}
Segments:
{AB, AD, AE, AF, AG, BC, BF, CD, CF, CG, DG, EF, EG, FG}
Triangles:
{ABF, ADG, AEF, AEG, BCF, CDG, CFG, EFG}

---

*In order to form a partition in the strict mathematical sense, edges do not contain their end points and triangles are open (i.e. they do not include their boundary edges or their vertices).

2. GENERAL TRIANGULATION CONSIDERATIONS

Obtaining a triangulation may be a goal in itself (in order to partition all of space in some useful fashion) or it may be an intermediate goal (as in generating an elevation model). There are many possible triangulations for most point sets. (Just counting all of the possible triangulations for an arbitrary point set is a very hard open problem which has been solved in only a few special cases. Even the "easy" cases, such as "n points are vertices of a convex polygon" have complex solutions.)

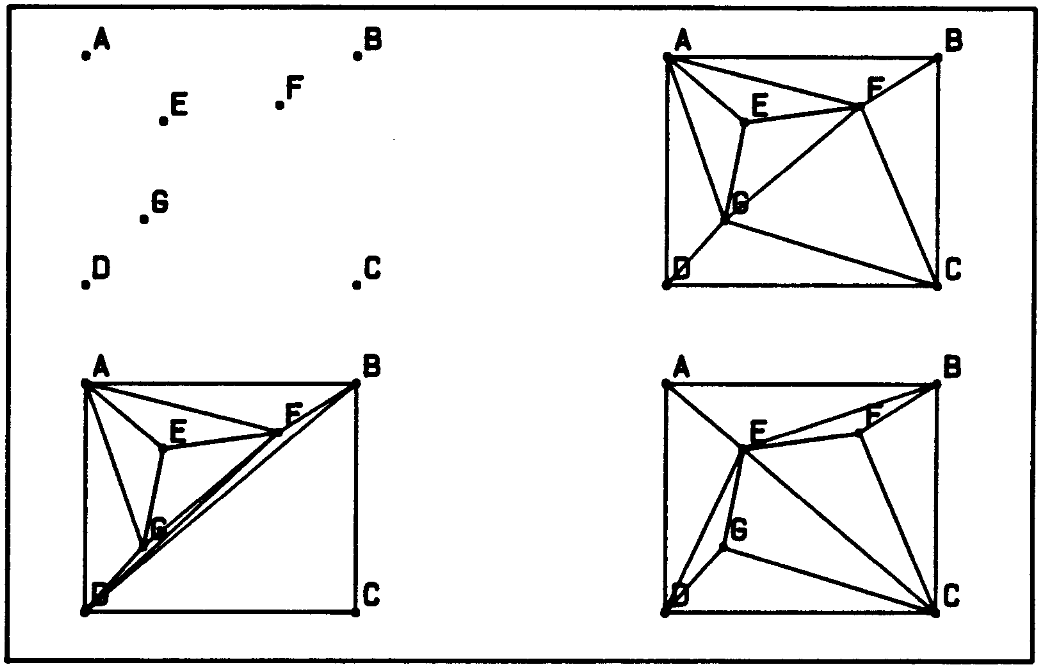

Figure 2. A finite set of points and three distinct triangulations

Some triangulations, such as the Delaunay triangulation (upper right in Figure 2), are more suited for certain applications. Most useful triangulations, including the Delaunay triangulation, can be built in $\Theta(n\log n)$ worst-case time, where n is the number of points on which to triangulate. With appropriate data structures, triangulations can be searched and updated quickly (often in $O(\log n)$ expected time). This paper describes some of those structures. A ***triangulated data structure*** will refer to any abstract data structure specifically organized to facilitate representing and accessing the elements of one or more triangulations and their topological and geometric relations. Some examples of triangulated data structures are Chris Gold's binary adjacency trees [3], David Kirkpatrick's triangulation refinement digraphs [4], and adjacency pointer structures such as DIME or "winged edge" and TIGER, with possible enhancements to exploit special facts such as: every triangle has exactly three vertices and exactly three triangle neighbors.

Before looking at specific triangulations and some useful data structures in detail, this paper focuses on a specific application of triangulations that will be generalized later. Suppose that a triangulation is given and fixed for the moment. Membership of any point in space in any triangle is well-defined—the point is either in the triangle or not in the triangle—and easily checked. Furthermore, if the point is in the triangle, its position in the triangle with respect to the three vertices can also be precisely defined. This well-defined position permits functions to be extended over an entire triangle when the functions are only defined at the three vertices. Furthermore this extension by triangles agrees on triangle overlap, and thus gives an extension to the whole space. The ability of a triangulation to extend point functions to all of space is a very useful property.

3. FUNCTION EXTENSION PROPERTIES OF TRIANGULATIONS

Elevation models and their classic utilization of triangulation methods typify the function-extension property of triangulation applications: elevations are measured at discrete sites in order to estimate elevations everywhere on a surface. Estimation at all non-measured sites is accomplished by averaging measured values at "nearby" sites; and the triangulation effectively determines (1) which sites are "nearby" and (2) how those "nearby" sites should be weighted to produce the desired average. The triangle to which a non-measured site belongs assigns it three neighboring vertices of "nearby" sites; the relative nearness to each of those sites determines the weight that each vertex elevation should be given. A rule of linear interpolation is easiest to describe and to illustrate: Suppose that the point q belongs to the triangle determined by p_1, p_2, and p_3. Then q can be expressed uniquely as:

$$q = \alpha_1 p_1 + \alpha_2 p_2 + \alpha_3 p_3$$

where $\alpha_1 + \alpha_2 + \alpha_3 = 1$; and α_1, α_2, and α_3 are all non-negative. The unique α_1, α_2, and α_3 are called the convex coordinates of q.

If each p_i has associated elevation e_i, for $i = 1, 2, 3$, then the elevation at q (call it e_q) is given by:

$$e_q = \alpha_1 e_1 + \alpha_2 e_2 + \alpha_3 e_3, \text{ for the same } \alpha_1, \alpha_2, \text{ and } \alpha_3$$

This interpolation computation gives consistent results on shared triangle edges. It also works for vector-valued functions, not just scalar functions such as elevation. In particular, vector-valued functions are interpolated to produce a special class of transformations called triangulation maps [7] and rubber-sheeting transformations [10] of the plane for map conflation [8].

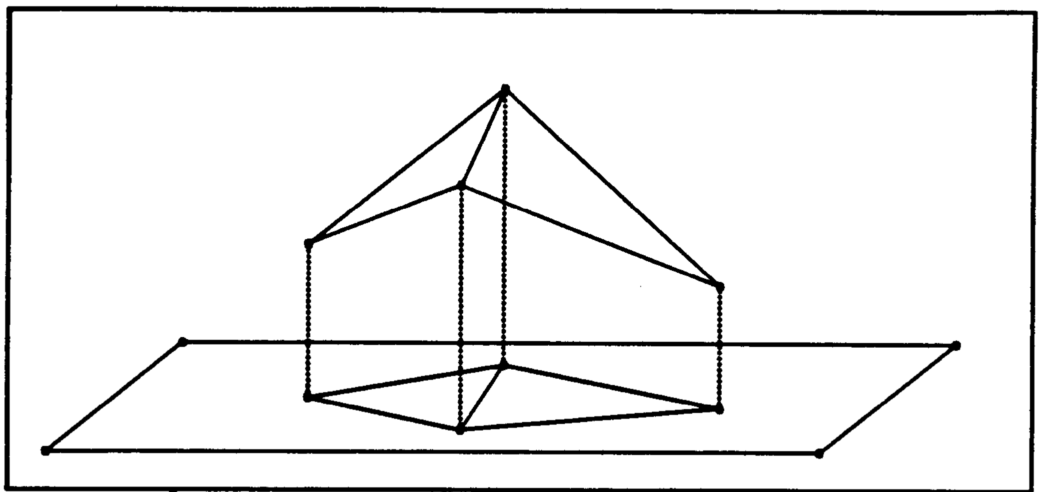

Figure 3. Linear interpolation of elevation over a triangle

4. DESIRABLE PROPERTIES OF DELAUNAY TRIANGULATIONS

Many triangulation packages in operation in geographic information systems, in computer-aided engineering, and elsewhere build Delaunay triangulations rather than some other type of triangulation. This section reviews some of the important properties of the Delaunay triangulation that make it most suitable for numerous applications (see also [6]).

1. A Delaunay triangulation on a set of points is a triangulation such that the circumcircle of any triangle (the circle passing through the three triangle vertices) contains no point of the set in its interior. This circle property fully determines a Delaunay triangulation. A Delaunay triangulation always exists.

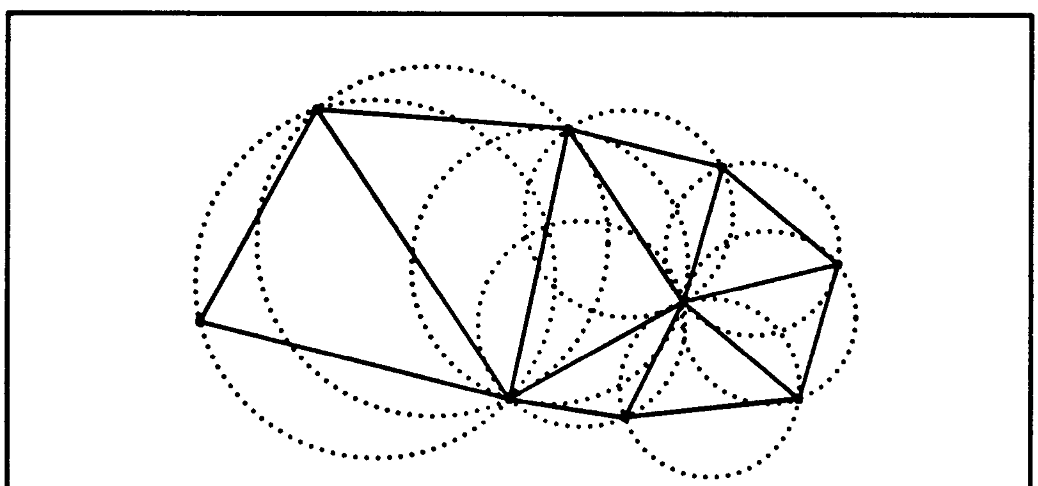

Figure 4. A Delaunay triangulation with circumcircles drawn

2. If no four points of the set are co-circular (i. e. lie on the same circle), then the Delaunay triangulation is unique. A Delaunay triangulation is always unique up to diagonal swapping within polygons whose vertices are co-circular.

3. Furthermore, if no four points are co-circular, then an edge will belong to the (unique) Delaunay triangulation if and only if there exists a disk containing both endpoints of the edge and no other point of the vertex set.

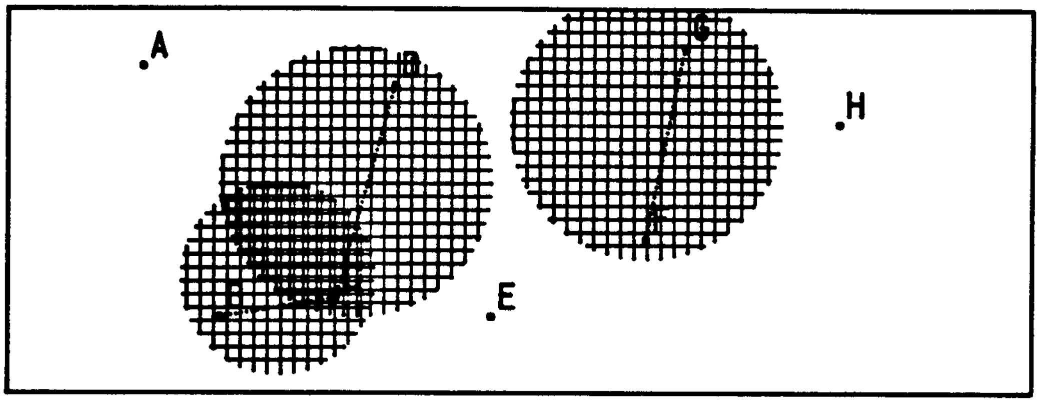

Figure 5. The disk/edge separation property illustrated

4. Because the circle property (1) and disk property (3) fully determine the Delaunay triangulation, the Delaunay triangulation of a point set will remain invariant under any transformation of the point set that preserves circles and circle containment. Rigid motions, scalings, reflections, and combinations of the three movements all preserve circles and circle containment.

5. The Delaunay triangulation is the planar graph dual to the Voronoi diagram which delimits planar regions according to their nearest point in the vertex set. The Voronoi dual is a useful structure for nearest neighbor searches; and it may be obtained from the Delaunay triangulation in $O(n)$ time and vice versa.

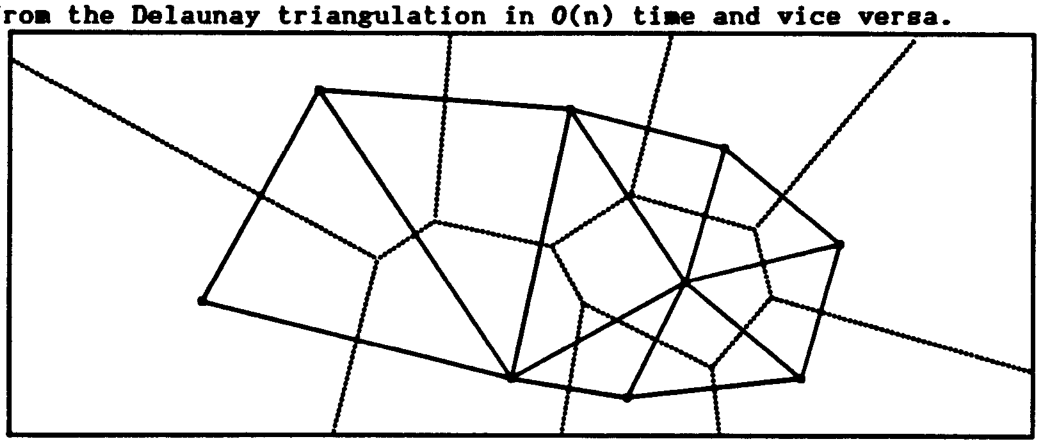

Figure 6. A Delaunay triangulation and its Voronoi dual

6. The Delaunay triangulation maximizes the minimum angle of all the angles that are present in the triangulation. This property is useful if very small angles can cause problems, as in the distortion resulting from piecwise linear homeomorphism of narrow triangles. Although the Delaunay triangulation does not minimize the maximum angle, it does tend to eliminate very large angles simply by enforcing the supplementary relation of three angles of a triangle, namely, they must add to 180 degrees; and if the smaller angles are not too small, then the largest angle cannot be too large. There are applications such as finite element analysis which seek to avoid large angles; and for these applications, the Delaunay triangulation gives good results.

7. The Delaunay triangulation (when unique) produces the lexicographically largest increasing sequence of angles possible in any triangulation. This means that if the angles of the Delaunay triangulation are listed in non-decreasing order:

$$\alpha_1 \leq \alpha_2 \leq \ldots \leq \alpha_i \leq \ldots$$

and if any other triangulation has its angles ordered similarly:

$$\alpha_1 \leq \alpha_2 \leq \ldots \leq \beta_i \leq \ldots$$

where α_i and β_i are the first term for which the sequences differ,

then $\alpha_i > \beta_i$.

This property is a generalization of property (5). It may also be used to extend the definition of Delaunay triangulation to specify further the case of four or more co-circular points.

8. A Delaunay triangulation may be updated locally with addition or deletion of vertices. The local update affects only triangles whose circumcircles contain the update point. An update can be accomplished with an incremental algorithm that adds a vertex in $O(\log n)$ average-case time and removes a vertex in $O(1)$ (constant) expected time when (1) vertices are in general position and (2) appropriate topological linkages are efficiently encoded in the data structure.

9. A $\Theta(n \log n)$ worst-case divide-and-conquer algorithm exists for building the Delaunay triangulation. The $n \log n$ factor arises from sorting the data, which seems to be a necessary pre-processing step for any triangulation, and especially for the Delaunay triangulation which depends entirely on local behavior of vertices. After pre-processing the vertices to group them locally, the Delaunay triangulation may be found in expected linear $O(n)$ time.

5. STRUCTURES FOR BUILDING AND NAVIGATING TRIANGULATIONS

Christopher Gold once remarked in a talk on triangulations that topological structure is too valuable to throw away. The Bureau of the Census has always affirmed the preeminence of topology as well for its general cell-based map model. The TIN programs developed by Environmental Systems Research Institute generate their triangulations as topological structures, complete with information about adjacency among the nodes, the edges and the triangles themselves. An awareness of the advantages of making topology a part of the triangulated data structure has made advocates of many users. A triangulation has topology; and the prevailing philosophy is to store that topology in an explicit, accessible format along with other triangulation information. As an example of an alternative approach, a triangulation may consist of something as minimal as an edge list—where verifying the fact that the edges actually constitute a triangulation or determining which edges belong to which triangles are left to the user. While topology is clearly constructible at a cost, many of the newly developed applications of triangulations, such as flow analysis, require ready network construction and traversal; and these can only be accomplished efficiently with full topological linkages.

Topology of triangles is somewhat simpler than topology of polygons, although in many ways it is the same. An n-sided polygon may have up to n distinct neighbors, each sharing an edge. A triangle may have only three. In terms of fixed-field data records, this characteristic is helpful. A triangle also has exactly three vertices and three edges, again useful from the fixed-field data record aspect. Geometry of triangles is even more constraining. A triangle is a rigid body in the sense that it is the only polygon that is fully specified by its side lengths. It is also the only polygon whose interior points can be expressed uniquely as a convex combination of its vertices. This last fact makes all ordered triangles affinely equivalent—any triangle with ordered vertices may by mapped onto any other triangle with ordered vertices by an affine map that sends the three vertices of the first triangle to the corresponding vertices of the other triangle. The affine map in question merely uses the same coefficients to form the convex combinations of vertices.

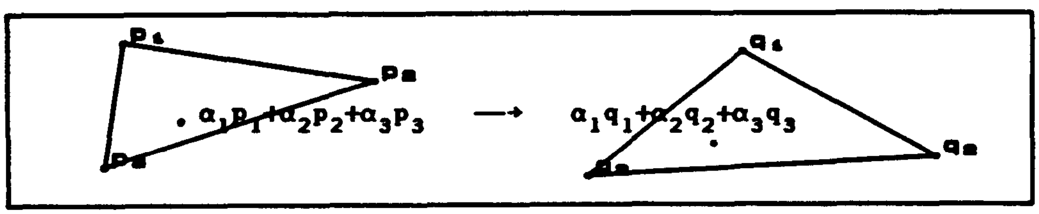

Figure 7. Affine map between triangles

The affine maps between triangles described above induce coherent maps between a triangulated space and a space with distinguished points corresponding to the vertices of the triangulated space. Such a map is called a triangulation map; and the image of the triangles may or may not triangulate the image space. A straightforward test involving preservation of triangle orientation has been devised to determine if the image space is also triangulated [7]; and if both spaces are triangulated by the corresponding triangles, then the triangulation map produces a joint triangulation. A joint triangulation is a topological homeomorphism—a one-to-one bicontinuous association of all points in one planar region with the points of another. Homeomorphic functions do not fold or tear a space when they transform it. A joint triangulation permits a finite description and storage of the homeomorphism relationship involving infinitely many points, just as an ordinary triangulation partitions infinitely many points into finitely many sets that can be manipulated by the computer.

Joint triangulations can be used to produce overlays which are aligned at arbitrarily many points [8]. The remaining points will maintain their topological relationships and produce the rubber-sheet effect desired. Non-linear rubber-sheeting, also popular for alignment tasks, is computationally more complex and may even run a risk of failing to preserve topology. In any case, verifying or proving that topology is preserved by non-linear adjustments is usually harder than doing so in the piecewise linear case. Three key results of a study of joint triangulations are stated here (for proofs see [7]):

1. A triangulation map can be tested for homeomorphism (joint triangulation) in linear time.

2. Although it is not always possible to extend a finite map from n distinct points to n distinct points to a joint triangulation of their convex hulls, it is always possible to augment the associated pairs of n points, giving a finite map of n + m distinct points to n + m distinct points (which sends the original n points to the corresponding original n points) in such a way that one can then find a joint triangulation of the augmented sets.

3. Joint triangulations exist for which one triangulation is Delaunay and the other is not. It is not always possible to find a joint triangulation which is Delaunay in either space. However, by permitting augmentation of the vertex set, one may guarantee the existence of a joint triangulation for which the first triangulation is Delaunay. It may even be possible, with augmentation, to guarantee a joint triangulation involving two Delaunay triangulations, although it is not known at this time.

6. SOME PRACTICAL EXPERIENCES

Census Bureau researchers have found that, in their applications, the following information has proved useful for manipulating a triangulated data structure:

1. TIGER-like topology.

   A. Files of 0-cells, 1-cells, and 2-cells.
   B. 0-cells with coordinates, pointers to "first 1-cell."
   C. 1-cells pointing to "to" and "from" 0-cells, "left" and "right" 2-cells, and four other 1-cells.
   D. 2-cells (triangles) pointing to three 0-cells, three 1-cells, and three neighboring 2-cells.

For direct spatial search queries, structures with $O(\log n)$ access time are available. The Census Bureau stores Peano-key sequences in B-trees for logarithmic access.

2. Vertices in Peano-key order accessed through a B-tree.

For neighborhood searching, coefficients of the straight line equation describing an edge permit a quick test to determine on which side of the line a point lies. The test consists of plugging the point's coordinates into the line formula and computing the sign of the expression.

3. Coefficients of the line equation of each 1-cell.

For building or maintaining a Delaunay triangulation, the circumcircle centers and radii are also useful data points [9]. (For computational purposes, the square of the radius is more easily computed and is often kept in the data structure).

4. Radius (squared) and center of circumcircle, with centers in Peano-key order stored in a B-tree.

Other tree-like structures to record adjacency of triangles or history of evolution in the case of incremental construction have been mentioned in the introductory section. Depending on applications, these additional structures may be well worth the overhead to store and maintain them. For recurring search applications with limited or local update requirements for the triangulation, a triangle-based directory of data points may be an efficient search tool [8].

REFERENCES

1. Environmental Systems Research Institute, 1987, ARC/INFO TIN Users Manual, Surface Modeling and Display, Version 3.2, Redlands, CA.

2. Frank, A., and W. Kuhn, 1986, "Cell Graphs: A Provable Correct Method for Storage of Geometry," Proceedings of the Second International Symposium on Spatial Data Handling, Seattle.

3. Gold, Christopher, and Sean Cormack, 1986, "Spatially Ordered Networks and Topographic Reconstructions," Proceedings of the Second International Symposium on Spatial Data Handling.

4. Kirkpatrick, D., 1983, "Optimal search in planar subdivisions," SIAM Journal of Computing, 12(1), 28-35.

5. Peucker, T. K., R. J. Fowler, J. J. Little, and D. M. Mark, 1976, "Triangulated Irregular Networks for Representing Three-Dimensional Surfaces," Technical Report No. 10, Simon Fraser University, Burnaby, B. C., Canada.

6. Preparata, F. P., and M. I. Shamos, 1985, Computational Geometry: An Introduction, Springer-Verlag, New York, NY.

7. Saalfeld, Alan, 1987, "Joint Triangulations and Triangulation Maps," Proceedings of ACM's Third Annual Symposium on Computational Geometry, Waterloo, Ont., Canada.

8. Saalfeld, Alan, 1987, "Conflation: Automated Map Compilation," Proceedings of ICA's Thirteenth Annual Symposium on Cartography, Morelia, Mexico.

9. Watson, D. F., 1981, "Computing the n-dimensional Delaunay tesselation with applications to Voronoi polytopes," The Computer Journal, 24(2), 167-172.

10. White, M., and P. Griffin, 1985, "Piecewise Linear Rubber-Sheet Map Transformation," The American Cartographer, 12(2), 123-132.