SPATIAL ORDERING OF VORONOI NETWORKS AND THEIR USE IN TERRAIN DATA BASE MANAGEMENT

Christopher M. Gold

College of Geographic Sciences,
P.O. Box 10, Lawrencetown
N.S., Canada BOS 1MO
BITNET address: COGS@ACADIA

ABSTRACT

"Computational Geometry" traditionally included topics such as curve and surface fitting; more recently it has been concerned with minimizing computational complexity in a global sense. Another aspect of spatial problems still needing attention is the concept of spatial ordering of data so as to minimize data base access, pen movement, etc., as well as to answer questions concerning the adjacency of objects on a map.

If we take point data for terrain modelling as our illustrative problem, and create a Delaunay triangulation as our spatial data base, we may define a limited set of authorized transactions that may access the data base. These include search, insert and delete of data points, and radial search outwards from some initial viewpoint. These transactions (and especially the last) are particularly critical for systematic map accretion procedures and field-of-view problems. Based on these data base considerations maps, block diagrams and field-of-view calculations may be generated in a consistent fashion from similarly-constructed data bases, and a generic procedure for producing block diagrams is given as an illustration of the approach.

INTRODUCTION

While the term "computational geometry" has traditionally included such topics as curve and surface modelling (Faux and Pratt, 1979), more recently the emphasis appears to have been leaning towards a more formal role, e.g. in defining the complexity of a particular geometric problem (e.g. Preparata and Shamos, 1985). Other perspectives may, however, still fall within the broader definition of the term. One of these is the concept of spatial ordering of data - in this case in two dimensions so that individual data points, etc., may be referenced in a consistent and predictable fashion, despite the predilection of traditional computers for handling information in a sequential linear fashion.

This topic is relevant to the computer-based mapping business, since technology is generating ever-more detailed data, covering extremely large areas, and demanding increasingly rapid updating of the appropriate data-base, often only modifying a few local areas at one time. While regional questions, for example the merging of two adjacent data sets, are of great importance, it is becoming increasingly necessary to be able to operate effectively within a very small part of a very large cosmos, without perturbing unnecessarily those regions not relevant to the immediate local issue.

EFFICIENCY OF NETWORK OPERATIONS

Following from these considerations is the philosophy that: A) individual operations on the data (whether interrogation or modification) are local in nature; and B) subsequent operations tend to be near previous ones. This affects significantly the appropriate definition of computational efficiency - in general, operations may be O(n) once the neighbourhood or relevant portion of the data base has been reached. Getting there, however, may be significantly less than half the battle - if the previous operation was nearby. Thus search procedures may sometimes be tolerated that are of less than the theoretical maximum order of efficiency. Preparata and Shamos (1985) indicate that in two dimensions, point Voronoi diagrams may be constructed in O(n log n) time. Experience with triangular networks indicates that operations are O(n) (i.e. local in nature) except for the search through the data structure to find the appropriate local element. Direct walk methods (e.g. Gold et al., 1977) are O(n**1.5) in theory, but in practice they rarely match the worst case. Improved data structures could reduce the search time to O(n log n) if necessary.

An earlier publication (Gold, 1984) examined the problem of terrain modelling or contouring from arbitrarily distributed data points. He broke the problem into five stages: data point entry and retrieval; sample site selection for surface estimation; neighbouring point selection; surface estimation procedures; and display methods. Apart from noting that surface estimation (interpolation) is heavily dependant on neighbouring point selection, our interest in these steps concerns the relation between neighbour selection and data entry or retrieval - in other words, effective utilization of an appropriate data base. For reasons mentioned previously the Voronoi tesselation appears to be a good general purpose measure of neighbourhood relationships. Gold et al.(1977) described a triangulation based data structure for terrain modelling. Lawson (1977) has described a criterion that is equivalent to Delaunay triangulation. Both workers used a technique of switching the diagonals of a quadrilateral formed by two adjacent triangles in order to improve the triangulation. An interesting sidelight on this approach is that using the optimization criterion of Gold et al.(1977), which does not produce a global optimum, approximately 5.7 diagonal switches were required to (locally) optimize one data point, ignoring boundary conditions. For the Lawson case, which is equivalent to the (global) Delaunay criterion (i.e. the triangulation which is the dual graph of the Voronoi tesselation) precisely 6 switches were needed on the average. An excellent summary of the computational geometry approach was given by O'Rourke (1984). It is the earlier work that interests us here - in particular the view of a Delaunay triangulation as a data base.

DATA BASES AND SPATIALLY ORDERED TRANSACTIONS

Perhaps the most convenient form for preserving the triangulation is as a file with one record per triangle, containing pointers to each adjacent triangle and each adjacent data point (vertex). This is conceptually convenient as it separates objects from their spatial relationships. If the object and relationship files (or tables) are to be treated as a data base, only certain authorized transactions may be performed on the underlying triangulation. Based on previous work (Gold et al., 1977, Gold and Maydell, 1978, Gold, 1984) some authorized transactions are:

"Location Search", which walks through the network to find the enclosing triangle for a specified coordinate location;

"Insert", which sub-divides the enclosing triangle into three, therefore updating the relational linkages to accommodate a new data point;

"Switch" or "Optimize" which adjusts individual triangle pairs until the Delaunay (or other) criterion was achieved in the neighbourhood of the new data point;

"Rotational Search" which retrieves the neighbouring points or triangles to an already-entered data point;

"Radial Search" which, given some central reference location, scans outwards from this, retrieving all points, triangles or edges until some terminating criterion is achieved.

Locational Search, Insert and Switch are basic operations required for network generation, and have been described previously. Rotational Search is used when the immediate neighbours are required to some point in the data base - for example, for estimating its slope coefficients. The Radial Search procedure (Gold and Maydell, 1978) permits the handling of spatial data in a front-to-back or centre-to-outside order by treating any triangulation as a binary tree with respect to an arbitrary reference (or viewpoint) location.

Figure 1a illustrates the Delaunay triangulation of a rather well-known test data set (Davis, 1973). The numbers represent the order of triangle processing. A viewpoint has been defined , having average X and very large Y coordinates. The processing order (as a binary tree) ensures that triangles closer to the viewpoint are processed before those further away. Figure 1b illustrates the same data set but with the viewpoint located within the

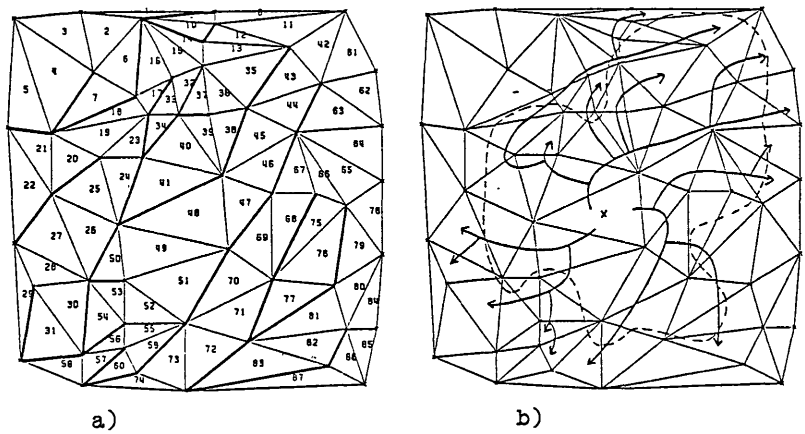

Figure 1.

a) Test data set with triangles, ordered with respect to a viewpoint from the north.

b) The same test data set with a viewpoint interior to the map.

map area. For more details see Gold and Cormack (1986) or Gold and Cormack (in press).

This spatial ordering process permits many useful operations, including: processing a map in contiguous segments; performing line-of-sight, perspective view and hidden line determinations; the extraction of all items likely to be affected in an update of the data base; and appropriate paging of network segments.

APPLICATION TO TERRAIN MODELLING

With the previously described tools at our disposal it is possible to outline an appropriate solution to the generalized perspective view of a topographic surface. The steps are:

1. Insert all data points individually into the database using the Voronoi criterion as described above.

2. Select the viewpoint and field of view.

3. Perform an ordered tree traversal of the triangulated network, starting with data points close to the field of view, and working outwards from there.

4. Reject all triangular facets whose vector normals point "away" from the viewer, as they will have already been hidden by closer portions of the surface.

5. Wherever a closer triangular facet faces the viewer,and its immediately posterior neighbour faces away, a horizon segment has been created. Maintain a radially-ordered linked list of horizon segments. Note that horizon segments will rarely occupy more than a small amount of the active edge of the terrain model at any one time.

6. Where a new forward-facing triangular facet is not behind a current horizon segment it should be drawn. Where it is behind one, it should be determined whether it lies below that horizon or not. If the facet is below the horizon it is not drawn.

7. If the facet is partially above the horizon the visible part of it is drawn, and the appropriate section of the active horizon is deleted. As a consequence of this no facets should occur that are entirely above a currently active horizon segment.

IMPLEMENTATION OF TERRAIN MODEL

In practice, two components are required in addition to the terrain data base software - these are computer graphics routines to permit object rotation and perspective viewing, and an appropriate hidden-line procedure. The transformation routines are available from any computer graphics text. The hidden-line routine that implements steps 5 to 7 may be readily developed. Any hidden-line procedure that maintains a horizon by vector intersection and linked-list maintenance has two properties: it is computationally expensive (hence the importance of eliminating segments as soon as possible, and keeping the active horizon portion short); and it is entirely dependant for its success on the strict preservation of front-to-back ordering of line segments submitted for display. The active horizon consists of those portions of line segments having the largest y coordinate to date (in the screen coordinate system) for any given x coordinate. Any new line segment passed to the routine must be assumed to be further away from the viewer than the previous segments making up the active horizon, thus acting as a clipping window for the new segment. If this is not true, the results make this very obvious - mysterious portions of the final map are blanked out for no apparent reason. Thus this problem is a good example of the strict requirement for ordered spatial processing.

Figure 2 shows the result of submitting a simple test data set to the perspective and hidden-line routines (for the purposes of this presentation lines that would have been hidden are instead drawn lightly). There is, however, one catch. In this example the 28 line segments were ordered manually, taking the viewpoint into consideration - thus Figure 2 shows only that the hidden line routine works. In Figure 3a the interpolation process outlined in Gold (1984)

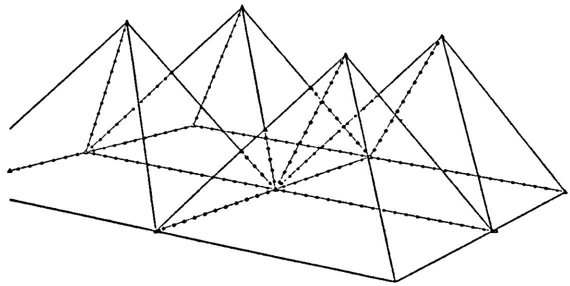

Figure 2.

Four pyramids - line segments viewed in perspective with hidden line removal.

has been used to interpolate elevations at nodes formed by the regular subdivision of the original triangular faces. In this case all line segments were generated by the contouring utilities, and triangles and sub-triangles were ordered away from the viewpoint using the procedures of Gold and Cormack (1986). Since the hidden-line procedure is highly sensitive to line mis-ordering, the spatial ordering procedures previously outlined are clearly effective.

Nevertheless, a problem is evident - perspective depth is not readily discernible: triangular facets do not provide sufficient depth cues. Consequently the surface was re-sampled on a regular grid using the previous interpolation procedure. Note, however, that there is no longer any guarantee that any subsequent surface will match the original data.

This re-sampled, gridded data was then triangulated and displayed as in Figure 2. The result is shown in Figure 3b. Here, however, the desired squares are formed from two triangles. Unlike squares, this gives an unambiguous surface (three points defining a flat plate, not four), but the presence of diagonals reduces the perspective effect. In Figure 3c these diagonals are removed, providing a view in which the perspective cues are satisfactory, even if local surface details are ambiguous. It is salutory to note that the actual information content of Figure 3c is no greater than that of Figure 2. Figure 3d shows a visibility map of the view in Figure 3c.

Thus, unlike some of the better-known discussions of perspective views of grids (e.g. Wright, 1973), it may be

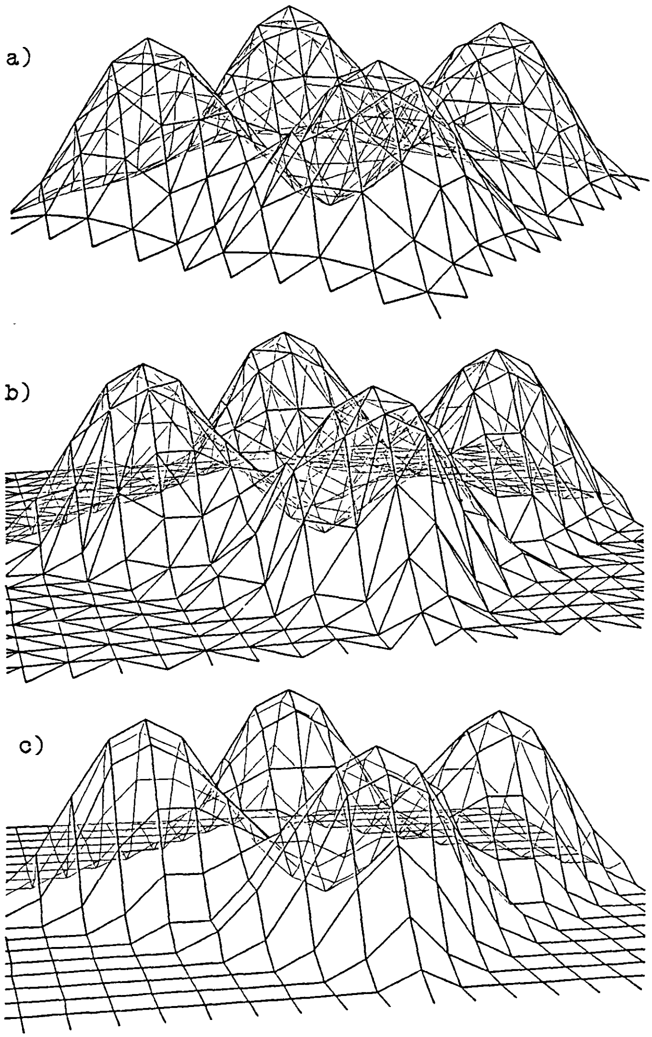

Figure 3.

Four pyramids:

a) broken into sub-triangles, with interpolation;

b) resampled on a grid and triangulated;

c) as in b) but with diagonals removed;

d) visibility plan - light lines not visible from the viewpoint.

d)

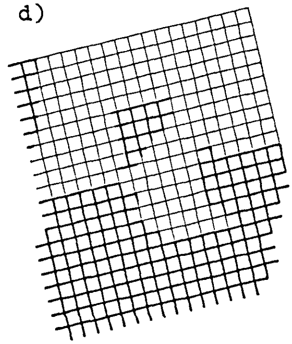

convenient to consider grids to be merely special cases of the triangular irregular network (TIN), and the ability to handle TINs in a spatially-consistent fashion permits consistent grid display. The advantages of grids lie in the depth cues. Their disadvantages are the ambiguity of four points forming a surface patch, and all the problems associated with re-sampling data, where that is required.

Ordered triangulations of Davis' test data set were shown in Figure 1. Figure 4a shows a perspective view of this data as a regular TIN. Figure 4b repeats this with triangular subdivision and interpolation as previously described. Figures 4c and 4d show this data re-sampled on a grid and then displayed again as a TIN, with the attendant ordering advantages. In Figure 4c the diagonals are retained, and in Figure 4d they are removed. In Figure 4e the visibility map is shown - this is a by-product of the hidden-line routine. In all cases the hidden-line procedure validates the spatial ordering processes described here.

SUMMARY

On the basis of the outline above, and the previously described spatial ordering procedures, it should be clear that the terrain display process is fairly efficient, as it accommodates a large number of the available spatial relationships. While vector-display perspective views or block diagrams are the most common application, shaded-surface views are equally applicable, as are shadowed-terrain maps, line-of-sight maps and various military applications.

While the generalized procedure is as given above, the special case of perspective block diagrams requires additional comment. Firstly, most block diagrams do not show all horizon lines since the "fishnet" model only shows the outline of the square grid. Since in fact the four corners of a grid are rarely coplanar, horizon (or outline) definition will necessarily be incomplete. To follow the procedure outlined, the squares must be (at least implicitly) subdivided into triangles prior to processing. Only when the horizon does indeed pass through the square does it matter which way the square is subdivided - and only then need the diagonal be actually plotted.

In conclusion, the following points bear repetition. Firstly, surface networks are a tool whose full potential has not yet been realized. Secondly, networks may be manipulated using local O(n) processes with the exception of the global location search function. While theoretical efficiency of the global search can be improved from O(n**1.5), for many applications this is of only marginal benefit. Thirdly, spatial ordering with respect to an arbitrary reference point is of considerable value for both display and data base manipulation. Finally, the example of terrain modelling illustrates how a well-defined set of spatial ordering procedures permits the development of complete, straightforward and efficient algorithms.

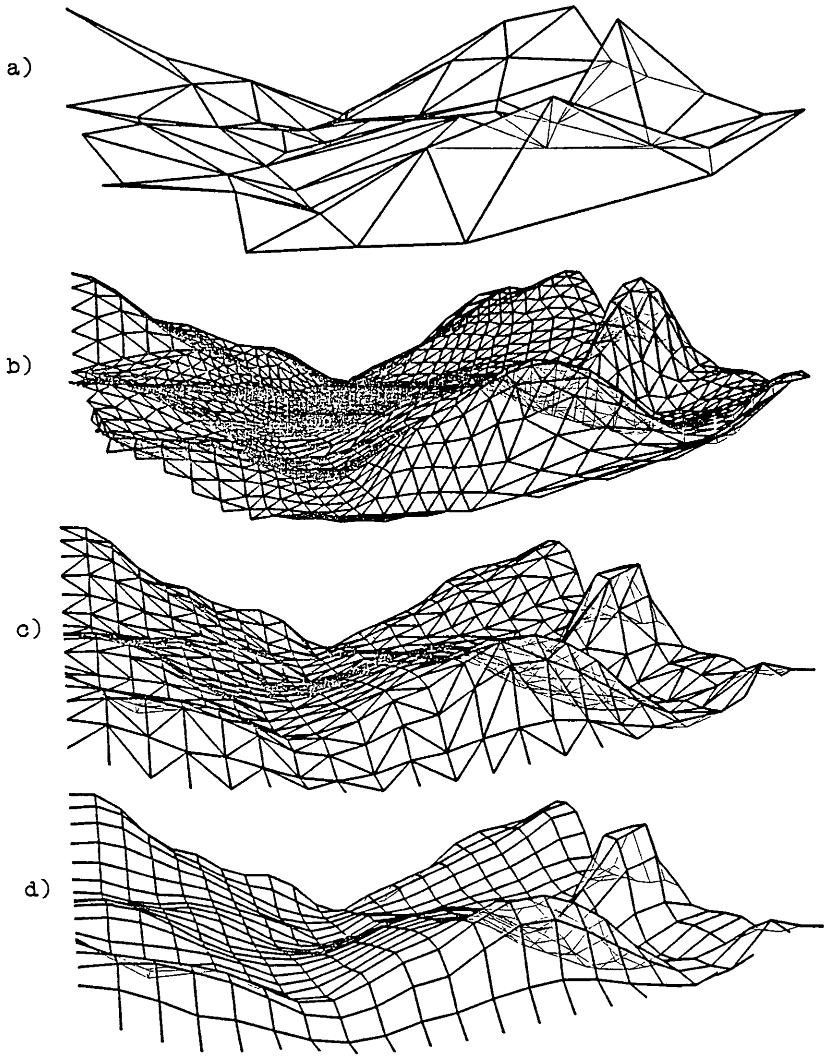

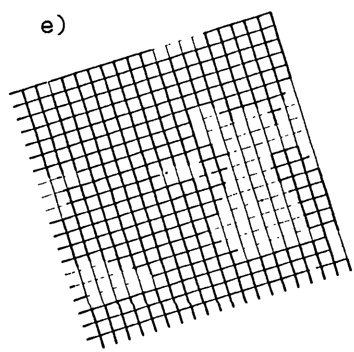

Figure 4. Test data set:

a) as a simple TIN;

b) broken into sub-triangles;

c) re-sampled on a grid and triangulated;

d) as in c) but with diagonals removed;

e) visibility plan - light lines not visible from viewpoint.

REFERENCES

Davis, J.C., 1973, Statistics and Data Analysis in Geology, John Wiley and Sons, New York, p.313.

Faux, I.D., and M.J. Pratt, 1979, Computational Geometry for Design and Manufacturing, Ellis Horwood, Chichester, 331p.

Gold, C.M., 1984, Common Sense Contouring: Cartographica, v. 21 no. 2, pp. 121-129.

Gold, C.M. and S. Cormack, 1986, Spatially Ordered Networks and Topographic Reconstruction: Proceedings, Second International Symposium on Spatial Data Handling, Seattle, July 1986, pp. 74-85.

Gold, C.M. and S. Cormack, (in press), Spatially Ordered Networks and Topographic Reconstruction: International Journal of Geographic Information Systems, v. 1 no. 1, January 1987.

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.11, June 1977, pp. 170-175.

Gold, C.M. and Maydell, U.M., 1978, Triangulation and Spatial Ordering in Computer Cartography: Proceedings, Canadian Cartographic Association Third Annual Meeting, Vancouver, June 1978, pp. 69-81.

Lawson, C.L., 1977, Software for C-1 Surface Interpolation, Jet Propulsion Laboratory Publication 77-30.

O'Rourke, J. 1984, Convex Hulls, Voronoi Diagrams and Terrain Navigation: Proceedings, Spatial Information Technologies for Remote Sensing Today and Tomorrow, Sioux Falls, October 1984, pp.358-361.

Preparata, F.P. and Shamos, M.I., 1985, Computational Geometry an Introduction, Springer Verlag, New York.

Wright, T.J., 1973, A Two Space Solution to the Hidden Line Problem for Plotting Functions of Two Variables: IEEE Transactions on Computers TC22, no. 1, pp. 28-33.