FITTING A TRIANGULATION TO CONTOUR LINES

by Albert H J Christensen

formerly GEONEX Corporation

7836 Muirfield Court
POTOMAC, Maryland 20854
(301) 983 9004

ABSTRACT

This paper presents a technique for creating a triangle mesh that tightly fits a terrain surface represented by a set of digitized contour lines. Basic to the technique is a Medial Axis transformation of a polygon, in this case formed by one or more contour lines. The advantages of using this mesh rather than the well known Delaunay triangulation for computing a gridded Digital Terrain Model (DTM) are discussed, as well as widely used spline interpolation methods. An example illustrates how the Medial Axis relates to the polygon and triangles and thereby facilitates further adjustments to the mesh. More complex adjustments to convert the triangulation into a surface devoid of unnatural features are described. Anomaly-free DTMs can be computed from contours without the supplementary features demanded by interpolation and triangulation procedures in use today. Desk-top computer programs operating on a small area of a scanned contour plate were prepared to test and illustrate the procedures that are outlined.

TWO CONTOUR-TO-GRID METHODS

Converting a given set of contours into a gridded numerical model of elevations, commonly called a Digital Terrain Model (DTM), can be accomplished by two widely different approaches.

Interpolation Method. The better known approach that is called here the 'Interpolation' method consists of the following: Vertical planes passing through each grid point intersect the source contours. Straight lines or planar curves contained in the vertical planes are defined by the intersections, and used to interpolate elevations at the corresponding grid point. There are numerous reports on implementation of this approach and on the nature of the curves used in the process. See references in [8,9].

Triangulation Method. The second, less known method, is called the 'Triangulation' method. The triangulation that constitutes the chief component for converting from contours to grid is performed by an algorithm that selects part or all the points in the source contours and establishes with them a mesh of non-overlapping triangles. From these triangles grid values are computed.

THE DTM IN BETWEEN SOURCE CONTOURS

Before discussing the problems found in interpolated DTMs, the subject of how the DTM is expected to behave in areas devoid of sampling should be examined. Obviously, the replication of source contours from the DTM should be a concern, although by no means the only or the most important one. Until recently, however, and so far as this authour could verify, only those few engaged in the creation or in the inspection of accurate DTMs considered the variations of a DTM away from the sampled areas. A concern so restricted was perhaps due to the lack of a reliable model against which a DTM could be compared.

Recently, two papers [8,9] have been published on evaluations of a number of interpolation techniques. In both papers the source contours are derived from a synthetic surface and converted into gridded DTMs by applying different interpolation methods. The predicted DTM values are then compared to those directly computed from the surface equation. A third paper [14] shows the wide disagreement in areas of low sampling density between derived contours and true contours which were not included in the input data set.

A synthetic surface exhibiting a number of formations similar to those found in topographic surfaces, as in [8], is a very attractive proposition for detecting and measuring DTM undulations. However, for accurate DTMs the maximum deviations allowable are smaller than the 25% of the contour interval which the plots in that paper show as lowest error. The plots in a future article would be perhaps very revealing if the authors would lower the mimimum error to, let's say, 3.5% of the contour interval. This is one of the maximum deviations established for accurate DTMs in flat areas.

DTMs generated to meet such strict specifications have to pass complete and thorough inspections. One of the tests compares a number of grid point values against values sampled from the source document.

Other tests developed for the verification of DTMs are mostly visual. The display of a grid of first and second differences computed from the elevations is an effective test. The differencies tend to highlight areas where the undulations introduced by the splines have propagated in linear or areal patterns, commonly known as 'unnatural' features. Examples are false dams, false depressions and bumps. Also clearly shown are patterns created by the symmetric distribution of intersecting planes, especially strong when only two planes are used.

Not surprisingly, the occurrence of unnatural features is higher where the spatial coherence of adjacent contours is lower, as in flat areas. Under strict specifications such occurrances must be avoided, which interpolation methods can accomplish if additional linear data is available. Examples are 'fabricated' contours, added to the source contours in flat areas. Other additional lines are used by programs that follow the interpolation of the grid points. Such are the drainage lines, with which the programs perform two functions. First, they introduce breaks in the DTM, and second, they remove any false dams accross the drainage lines. Supplementary drainage and other terrain features are created in low coherence areas, usually in correspondence with strong contour sinuosities, and processed together with the natural drainage lines.

Adding linear features to a contour set is a task that demands a fair amount of training and a good understanding of the entire DTM process. Moreover, it is a manual digitization task and consequently, costly both in labour and in equipment.

THE RULED SURFACE BETWEEN CONTOURS

More interesting than anything in a DTM Quality Control document is the aforementioned selective point verification. A number of elevations at grid points are evaluated, presumably by comparing them with elevations extracted from the source topographic map. How are these elevations computed? Most likely in the same way a topographer of Yesterday interpolated contours. For instance, when he metricized a map. Since the operation was manual, he had to use the simplest procedure that could be carried out with contour lines.

The topographer proceeded according to the assumption traditional in elementary Descriptive Geometry: between contours a terrain surface is ruled and not developable. In other words, along certain straight lines a topographic surface has constant slope. It follows that DTM derived contours ought to be as regularly spaced as possible between source contours or, more formally, that distances between derived contours measured along lines of maximum gradient should be equal.

Needless to say, if a DTM quality control test is modelled on a ruled surface, it makes good sense to design the DTM around the same model.

THE TRIANGULATION APPROACH

The fact that triangulations created from contours have not been implemented as frequently as interpolations may be explained by their degrees of success. This observation does not apply to triangulations of randomly distributed points, such as in meteorology, geology and the like, where triangulations are routinely accepted.

The triangulation of a point set used today for most applications is Delaunay's [11]. Its tendency to yield triangles as well shaped as possible makes it attractive for applications using functions with singularities at very small angles. It is even more attractive because its uniqueness, which in turn makes the task of programming it light. The disadvantages are, first, a special configuration of points that must be considered [11], second, thin, sliver-like triangles along the perimeter of the mesh, which Delaunay algorithms create just to achieve convexity, and third and most important, that if a 'brute force' approach is taken, the processing time may grow beyond realistic possibilities.

Many solutions have been proposed to reduce the growth of the processing time to more manageable limits. Almost all of them exploit the principle of 'Divide and Conquer' lucidly exposed in [1], and ought to be applied to all triangulations of great numbers of scattered data points.

With the exception of those concerned with a distance optimality, all the examples known to this author on triangulations of point sets are Delaunay's. So do the few Contour-to-Grid conversions by triangulation: two commercial Site Engineering packages and the implementations in [10,13.]

Applied to contours, the Delaunay triangulation knowns only of contour points. The fact that the points are connected in the shape of contours is not considered. Consequently, poorly configured triangles may result. A case is that of a triangle edge crossing a contour segment. The triangle edge, now supposedly an element on the terrain surface, may have in correspondance with the contour segment an elevation different to that of the contour. If the crossed contour is higher or lower than both adjacent contours, the error amounts to 100% of the contour interval, see Fig.1. To prevent such configurations all the contour segments should be selected as triangle edges, which is a proposition that invalidates the Delaunay triangulation as applied to the entire set of contour points. See in Figure 1 a catastrophic false dam.

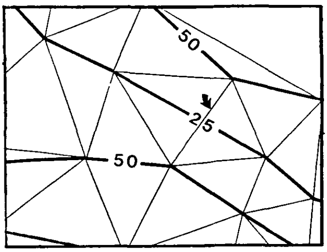

Figure 1. Contour crossing

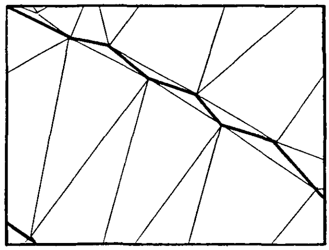

Figure 2. Flattening along contours

A second case is that of a triangle with its three vertices on the same contour. Such triangles cause breaks in the surface and should be avoided if the DTM is to be smooth between contours. Perhaps the most striking result of this poor configuration are horizontal bands of triangles produced along sections of contours, Fig. 2, and beyond the band areas slopes slightly steeper than what they really should be.

Avoiding the crossing of contours. The contour crossing case could be avoided by performing Delaunay triangulations in between adjacent contours, which means that the entire map would be covered with many triangulations, each of them executed independently of the others. Besides avoiding the contour crossing, this procedure will greatly alleviate the processing time problem, since the contours provide natural boundaries for the application of the 'Divide and Conquer' principle. There is no need for artificial divisions when contours are present. Indeed the triangulation of a point set inside a closed shape is not a novel idea, although this author has not yet seen it applied to topographic surfaces. The field of Pattern Recognition offers one example [7]. Heuristic and optimal triangulations, non-Delaunay, of bands limited by successive planar contours, have been proposed in [4,5,6] for the reconstruction of 3D surfaces. However, these references must not be interpreted as suggesting that such techniques can be applied to topographic surfaces. Terrain surfaces are single value functions of two variables, but they can be far more topologically complex than the true 3D surfaces of the type reconstructed using the reported techniques.

Because of the aforementioned topological complexity, the Delaunay triangulations inside a closed shape, with islands added, is not as easy a proposition for computer programming as the general Delaunay triangulation.

Horizontal Triangle Case. The second objection to the general Delaunay triangulation, triangles with their three vertices on the same contour, is not so easily removed. It will be discussed later.

It must be noted that the critiques in this paper to the general Delaunay triangulation of contour maps ends precisely with the triangulation. After the triangulation has been established, other procedures may be used to reshape it, for instance, exchanging edges so that they would not intersect contours, while others might even fit high order surfaces to the planar triangles for computing grid points. The availability of such follow-up procedures does not negate the conclusions of this paper.

THE MEDIAL AXIS TRANSFORMATION

This author's opportunity for experimenting with some old ideas on how to create a ruled surface from contours arose from the need for a procedure to thicken or widen line features. This need was satisfied by developing a 'Parallel Pairs' procedure, published elsewhere [3], that also suggested possibilites for solving some other problems. One of these problems was the 'Medial Axis Transformation', an operation which turned out to be basic to the Contour-to-Grid solution described in the next sections.

The Medial Axis [7,12] or midline, of a closed shape or polygon, is, rougly, a network of lines whose elements are equidistant from the closest pairs of elements in the shape. The Medial Axis is defined in vector environments. In a raster environment, it corresponds to the skeleton of a shape. In computer operations, the Medial Axis transformation corresponds to the raster 'thinning' or 'skeletonizing' operation with which commercial scanners are often provided. The thinning operation, coupled with the raster-to-vector conversion that comes with commercial scanners, make a fast and robust tool for generating the Medial Axis. There is an abundant literature on thinning, see for instance [2]. The vector mode operation seems to be less popular.

The output of the Parallel Pairs procedure is a set of polygons nested inside the input polygon, see Fig.3, with which the determination of the Medial Axis is accomplished in an efficient way. Exhaustive searches become unnecessary, as reported in [3], because the Parallel Pairs are loaded with pointers that indicate through which points the Medial Axis should be threaded. Pointers extracted from the parallel pairs are loaded into the Medial Axis as well, linking its elements to the equidistant edges of the input polygon.

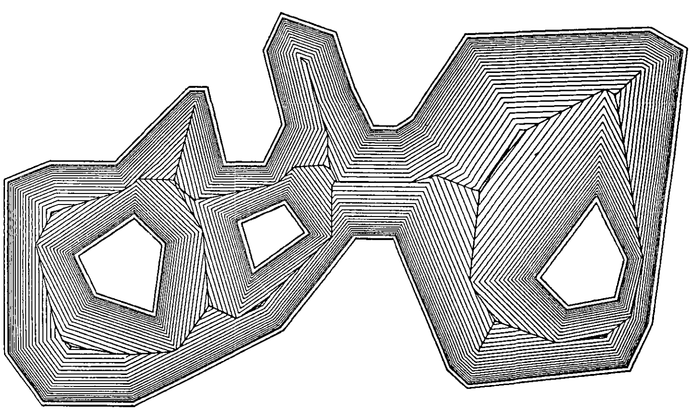

Figure 3. Dense parallel pairs and Medial Axis.

Advantages of determining the Medial Axis in raster rather than in vector mode are simplicity in programming, robustness, and perhaps time performance. Disadvantages are inflexibility, lack of structure and the need for a raster-to-vector conversion that follows the thinning operation. Inflexibility arises from the fixed resolution of a raster system. All the polygons in the file, irrespective of their particular shapes and dimensions, are processed with the same resolution.

Lack of structure refers to the absence of pointers and other features that facilitate further operations. The raster skeleton cannot be related to the polygon edges, at least in today's commercial software. Nor is easy to see how it could possibly be done, when the input is just a raster image. If the skeleton and the polygon together must be processed further, as in the case discussed here, the lack of structure would surely offset any time saved by the raster mode operation.

On the contrary, the vector approach offers flexibility and a potential for structure. Its flexibility is found in the wider range of the arithmetic that is used. The Parallel Pairs procedure, as part of a Medial Axis Transformation, increases that flexibility by providing the very significant option for changing offsets in the nesting of polygons. See Fig. 3. It also provides structure in the pointers referred to earlier.

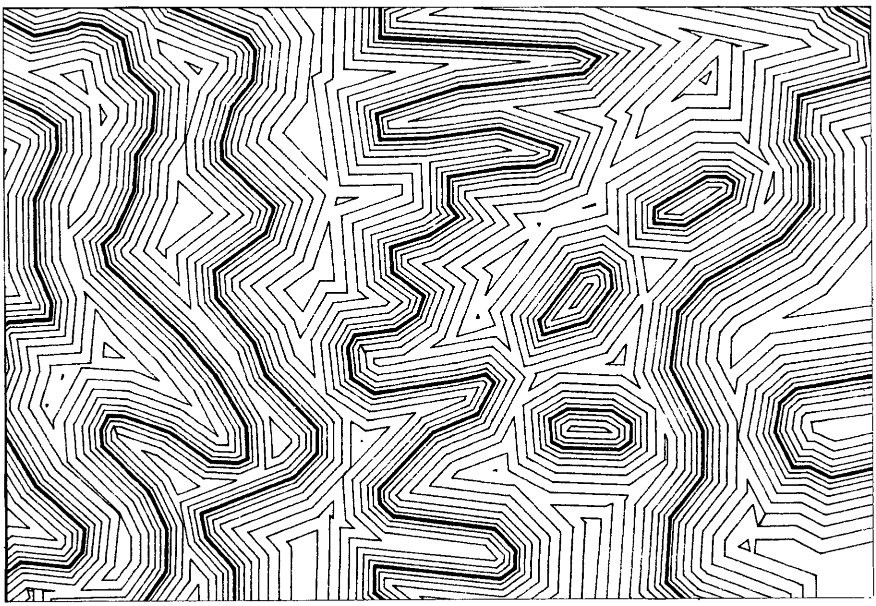

Figure 4. Parallel pairs in a section of a contour sheet

Figure 4 shows a small section of a contour sheet, with nested polygons generated with offsets greater than those used for figure 3. The original was a 1:50000 topographic sheet, scanned at 16 lines/mm resolution. At the original scale the area illustrated here measured

2 by 1.2 cm. The nested polygons were created in the areas bound by adjacent contours, in a process that was run separately for each area. They were then merged and plotted. The smallness and low speed of the desk-top computer used to prepare the software, resulted in data sets that are very limited in complexity.

THE TRIANGULATION OF A POLYGON AND ITS MEDIAL AXIS

As noted earlier, a Delaunay Triangulation, be it applied to disconnected contour points or executed inside a closed contour polygon, in many cases will select the three vertices of a triangle from the same contour, and that these horizontal triangles introduce breaks that do not provide a natural gradient to the surface. This is a problem that could not be ignored.

That problem can be solved by using the Medial Axis because there are always points on the Axis that can be connected to the contour points. The triangulation uses the Medial Axis points to bridge the spans between contours. Furthermore, because it is executed between contours, this triangulation does not cross them. Because each Axis can be given the mean of the contour elevations, the triangles on both sides of the Axis will have the same slope. With this procedure the Medial Axis itself will not turn out to be an unnatural feature.

In the program prepared to test the proposed solution, the vertices of the triangles are selected with a simple rule: the base of a triangle is defined by two consecutive points, either from the contours or from the Axis. If from the contours, then the apex of the triangle is selected from the Axis, and vice-versa. The pointers in the Axis tell the process from which entity, contour or Axis, to select the next base. Executed in this way, the triangulation program is extremely fast. Fig. 5 shows the triangles established in the same small contour shape of Fig.3.

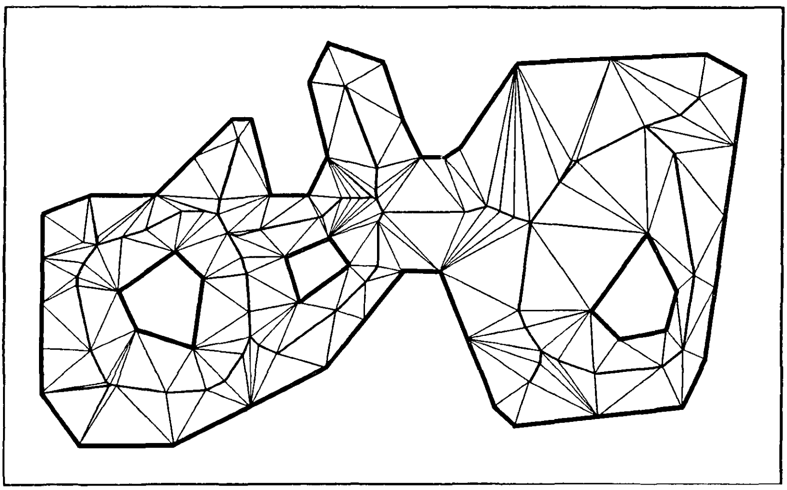

Figure 5. Triangulation of contour shape and Medial Axis

Assigning elevations to the Medial Axis. A quick look at a few shapes and their Medial Axes leads to the following classification. As regards points, only the endpoints of a line, called nodes, are considered. The number of lines incident to a node is called the 'degree of incidence'. These are Graph Theory terms. Lines can be classified as open or closed. A closed line has only one node and this is of degree 2. As Medial Axis, a closed line is a rarity. Open lines are classified here as Main lines and Branches. A Main line has its two node of degrees 2 or higher, or both of degree 1. Branches have one and only one node of degree 1, the dangling endpoint. With this classification is is possible to conclude, in a general way, that Main lines are connected, by means of the triangle edges, to two different contours. Branches to only one contour. See Fig. 5.

One part of the triangle vertices, those on the contours, can only be given the corresponding contour elevation. As for assigning elevations to the rest, on the Medial Axis, it is necessary first to make an assumption on which elevation to give to the Medial Axis.

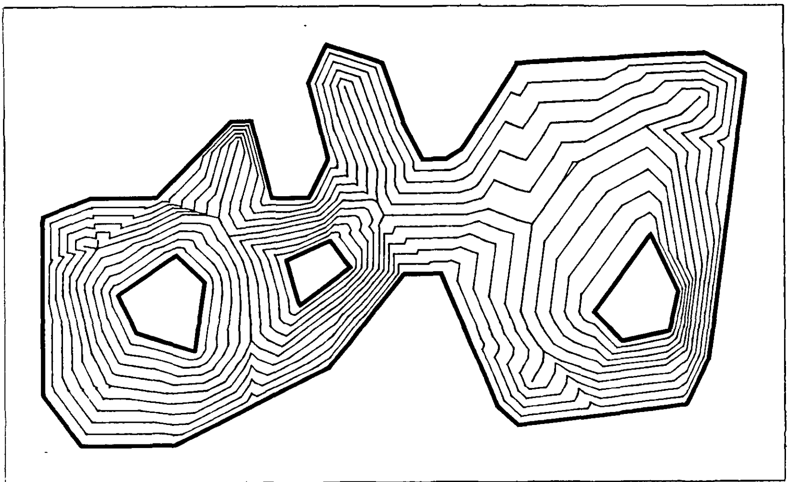

Figure 6. Derived contours for a Medial Axis with constant elevation

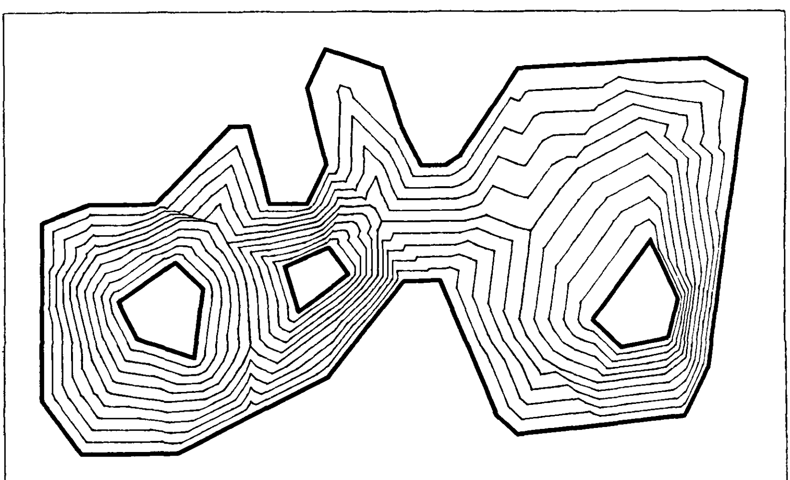

Figure 7. Derived contours for a Medial Axis with adjusted elevations

To avoid turning the Main lines into unnatural features, their points should be given the mean of the elevations of the contours with which each Main line is connected. The same cannot be done with the Branches. If they are given the same mean as the Main lines, the result will be a strong gradient located at the end of the Branch. Fig. 6 shows the contouring of the triangulation executed under this assumption. Notice how the contours are crowded at the end of the various Branches. A more natural look and a better approximation to a ruled surface is achieved by assigning variable elevations to the Branches. Fig. 7 was obtained with an option of the experimental software, which assigns to the Branches' points elevations proportional to the distances measured along the Branch from the non-dangling node. The proportionality is established between the difference in elevation, mean minus contour, and the length of the Branch plus the length of the shortest edge that connect the dangling end of the Branch with the contour. The better quality of the result, compared with the one obtained by giving constant elevation to the Axis, is evidenced in Fig.7: better spaced derived contours along all the Branches.

THE CONTOUR-TO-GRID PROPOSED SOLUTION

The Medial Axis transformation and the simple triangulation that comes after it are just two, if important, steps in the proposed solution. To make the operation of the Parallel Pairs possible, and in general, to improve the time performance of the software, as well to simplify its overall design, the input contours must be preprocessed.

Preprocessing of contours. First, all the contours are assumed distinct and without gaps. Those that reach the map borders must be turned into closed lines. In doing so, the closing lines ought to be such that the nested polygons would have the proper orientation when crossing the map borders. Figure 8, to be inspected together with Figure 4, shows how this was done with a simple program.

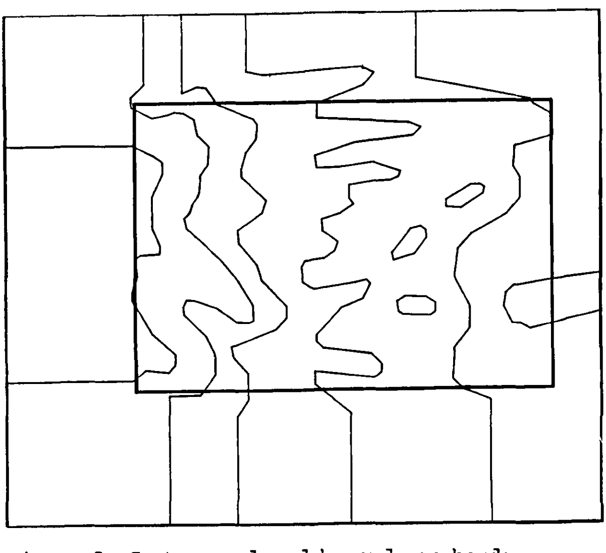

Figure 8. Contours closed beyond map borders

Second, adjacency relationships and containment should be introduced in the source contours and spot heights. These relationships are needed, inter-alia, for assembling pairs of contours into the shapes to be triangulated. The abundant literature on this subject make unnecessary any explanations.

Third. Although not strictly needed, it is convenient to obtain from the contours some measures of size and their proximity to each other.

The last two requirements can be best satisfied by programs operating in raster mode at the time the contours are vectorized. If a commercial vectorizer is being used, it will be necessary to rasterize the contours before running these programs.

After the preprocessing of the contours, the Medial Axis is determined, as already described, for each of the areas enclosed by one or by two succesive contours. This step is followed by the triangulation of those areas, which in turn is followed by the computation of the grid.

The Grid from Triangles. The computation of the grid values from the triangles, if these are considered planar, is a simple operation and does not merit any reference here.

However, if a smoother surface is desired, the triangles may be turned into curved patches that preserve continuity accross edges. There are many ways of defining such patches. The issue, in the view of this author, is not how to do it, but whether or not to do it, and the answer, on technical grounds only, is no. Yet, if some smoothing is still wanted, it will be enough to break the triangles at their half heights, and to assign to the breaks elevations that reflect a curvature along lines of maximum gradient. Of course, in directions normal to these lines, any smoothing would be still more superfluous.

Tops and Depressions. Shown in a contour sheet as empty closed lines, they have preocupied the advocates of triangulations since very early. In most cases, these closed contours do not include spot heights in sufficient numbers and proper distribution to ensure a good reconstruction of the terrain. The results are 'truncated tops' and 'flattened depressions' in the DTM. To produce a correct DTM the user will have to create automatically, or by hand, the right number of spot heights in the right places. The Medial Axis provides an automated solution to this problem. Figure 9 shows the Medial Axis and the resulting triangulation of a top contour.

Assigning elevations to a top contour or to a depression is done either by resorting to the spot heights or to the triangles adjacent to the contour in question. From those triangles slopes can be extracted and then applied to the triangles inside the top contour Or depression. The procedure followed for assigning elevations is very much the same used for contour polygons in the general case.

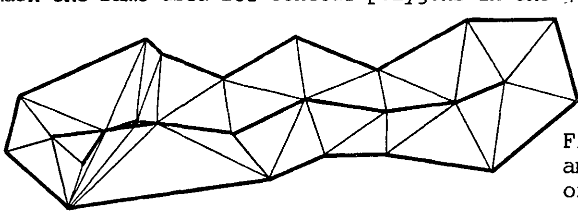

Figure 9. Medial Axis and triangulation of top contour

CONCLUSIONS

The Medial Axis provides the means for triangulating contours in optimal configurations, from which an accurate gridded DTM can easily be computed. This DTM behaves like a ruled surface, and consequently, does not exhibit any of the unnatural features introduced by spline interpolations, nor the breaks and false dams created by Delaunay triangulations. The discussions of these techniques were done with strictly specified DTMs in mind. It is hoped, however, that the precise fit of the triangulation described here will facilitate the introduction of accurate DTMs into fields where profiling and cross-sectioning are still prevalent.

AKNOWLEDGEMENTS

This paper would never have been produced had not the author been provided the encouragement and support of his former employeer, GEONEX Corporation.

REFERENCES

[1] Bentley, J L and Shamos, M I 1978, Divide and Conquer for Linear Expected Time: Information Processing Letters, Vol.7,pp. 87-91

[2] Bookstein, F L 1979, The Line-Skeleton: Computer Graphics and Image Processing, Vol. 11, pp. 123-137

[3] Christensen, A H J, Parallel Pairs in Automated Cartography: to be published in Cartographica, Vol. 24, No. 1

[4] Fuchs, H, Kedem, Z M and Uselton, S P 1977, Optimal Surface Reconstruction from Planar Contours: Comm.of the ACM, Vol.20, No.10

[5] Gannapathy, S and Dennehy, T G 1982, A new General Triangulation Method for Planar Contours: Computer Graphics, Vol.16, No.3, pp.69-74

[6] Keppel, E, 1976, Approximating Complex Surfaces by Triangulation of Contour Lines: IBM Journal of Research and Development, XIX

[7] Lee, D T 1982, Medial Axis Transformation of a Planar Shape:IEEE Transac. on Patt.Analysis and Machine Intell., PAM14,4, pp.363-368

[8] Legates D R, and Willmott, C J 1985, Interpolation of Point Values from Isoline Maps: The American Cartographer, Vol. 13, No.4, pp. 308-323

[9] Loon, J C, and Patias, P G 1985, Digital Terrain Elevation Model Analysis, Report ETL - 0393, U.S. Army Corps of Engineers.

[10] McCullagh, M J 1983, Transformation of Contour Strings to a Rectangular Grid Based Digital Elevation Model: Proc. Euro-Carto II.

[11] McLain, D H, Two Dimensional Interpolation from Random Data: The Computer Journal, Vol. 19, No. 2, pp. 178-181

[12] Shapiro, B, Pisa, J and Skansky, J 1981, Skeleton Generation from x,y Boundary Sequences: Comp. Graph.and Image Proc. Vol.15, pp.136-153

[13] Witzgall, C, Bernal, J and Mandel, B, 1986, On Sampling and Triangulating Large Digitized Contour Data Sets. National Bureau of Standards and U.S. Army Engineer Topographic Laboratories

[14] Yoeli, P. 1986, Computer Executed Production of a Regular grid of Height Points from Digital Contours, The American Cartographer, Vol.13 No.3, pp. 219-229