Terrain Reconstruction from Contours by Skeleton Retraction

D. Thibault [1] C. M. Gold [1]

[1] Centre for Research in Geomatics
Laval University, Quebec City, Qc, Canada
Tel: +1 (418) 656-3308, Fax: +1 (418) 656-3607
E-mail: david.thibault@scg.ulaval.ca / christopher.gold@scg.ulaval.ca

KEY WORDS: Terrain modeling, Delaunay/Voronoi diagrams, Contour lines, Skeleton, Quad-Edge data structure, Generalization

ABSTRACT: Generating terrain models from contour input is still an important process. Most methods have been unsatisfactory, as they either do not preserve the form of minor ridges and valleys, or else they are poor at modeling slopes. A method is described here, based on curve extraction and generalization techniques, that is guaranteed to preserve the topological relationships between curve segments. Based on this approach, which uses the properties of the Voronoi diagram and Delaunay triangulation, it is possible to make reasonable estimates of slopes for terrain models, and to extract meaningful intermediate points for triangulated irregular networks (TINs).

1. INTRODUCTION

Despite the development of modern satellite systems, a great deal of the world's database of terrain elevation is in the form of contour overlays from traditional mapping agencies. These have the advantage that they were developed using human understanding of the observed landforms, but the disadvantage that they are not easily converted to a useful digital format. Due to the hight cost of direct methods for terrain capture, cartographic documents such as contours maps and profiles are often prefered (Weibel and Heller,1990)

The usual methods for generating a terrain model from contour lines consist either of interpolation onto a grid, or of the construction of a triangulation of points on the contour lines. Carrara et al. (1997) discuss different techniques for generating digital terrain models from contour lines. The errors in gridded data are well known, in particular the difficulty of producing a suitable set of neighbouring data points for the weighted-average function often used. This is particularly pronounced when the data is ill-distributed, as is the case with ships' paths, aircraft flight lines, digitized contours, etc. The use of Voronoi diagrams and "area-stealing" or natural-neighbour interpolation has been shown to adapt well to poor data distributions (Gold, 1989) because the insertion into the mesh of a sampling point generates a well- defined set of neighbours, which may be few or many, and the areas stolen from these neighbouring Voronoi cells act as a well-behaved weighting function. Triangulation methods, the main topic of this article, depend both on the selection of correct triangle edges between points on the contours, and the generation of sufficient additional points to correctly form ridges, valleys, peaks and pits. Garcia et al. (1990) discuss the triangulation of contour input.

2. PREVIOUS WORK

Triangulation methods (TINs) have been popular for some years, and have the advantage of being adaptive to the data distribution. They are usually based on the Delaunay triangulation, the dual of the Voronoi tessellation. These can be extremely effective, but are not always ideal for contour input. They may give a very angular surface, and it is not easy to ensure that they connect adjacent contour lines correctly. If the contours are too close, or the sampling along the contour is inadequate, triangle edges may cross contour lines. If there are re-entrants or promontories in a contour line, then the three triangle vertices may have the same elevation, and are unable to mimic the human perception of a minor ridge or valley as a cause. Robinson (1994) present a method to fix this problem by changing the diagonal of the triangle. Similar problems arise at peaks and pits.

Another issue of note is the estimation (and perception) of slope based on the contour lines. Any method of interpolation that makes a poor choice of neighbours for a weighted average will first of all have a deleterious effect on the perceived slope of the resulting model. Robinson (1994) also discusses various techniques that partially take account of slope information at contours. TINs often have widely variable slope between adjacent triangles, yet it is obvious that the line of maximum slope is perpendicular to each contour line segment. Interpolation techniques, in particular, often generate a zero slope at each data point - giving a terraced or staircase effect to a terrain model derived from contours. Thus current methods are theoretically unsatisfactory although it may be possible to assign slopes to data points prior to interpolation (Gold, 1989).

This paper will attempt to examine both the problem of generating intermediate points for valid triangulations of contours in the regions of ridges, summits, etc., as well as addressing address the question of the valid estimation of slopes at contour lines. This depends on an understanding of the primal/dual relationships of the generating Delaunay/Voronoi structures, and recent methods for extracting curves from unordered input data, such as scanned maps.

3. CURVE EXTRACTION

Given that the Voronoi/Delaunay construction satisfactorily adapts to varying data distribution, this was taken as our baseline model for preserving topological relationships. Okabe et al. (1992) provide an excellent summary of Voronoi diagrams and applications. Our first problem was then to examine the problem of ensuring that all contour line segments were part of the Delaunay triangulation - or at least to have appropriate sampling rules. This could be considered in terms of curve extraction from unordered input data. Gold et al. (1996) worked on this for polygon maps, and Gold (1997) for scanned maps, where "fringe" points were generated on each side of the boundary, the Delaunay triangulation constructed, and the "skeleton" of Voronoi edges between the two rows of fringe points was extracted. Various workers in computational geometry, most recently Amenta et al. (1998), have worked on the problem of extracting the "crust", or the connected subset of Delaunay edges that forms the curve or polygon boundary. However, our original need was to extract a skeleton, not a crust. Of great importance, as well as defining the crust for a well-sampled curve, they were able to define the sampling conditions that were necessary to guarantee successful curve extraction. Point separation is proportional to the distance from the crust to the medial axis (or skeleton) of the curve

In Gold (1999) a simplified method of the technique of Amenta et al. was described, which depended only on the local testing of associated Voronoi/Delaunay edge pairs. Each pair would be drawn as a Delaunay edge (crust), or as a Voronoi edge (skeleton), but not both. The test simply stated that, for a Delaunay edge to be considered part of the crust, a circle must exist through its two vertices that does not contain any Voronoi vertices. Otherwise the associated Voronoi edge is part of the skeleton. This directly resolved our interest in skeleton generation, but, of even more importance, both the crust and the skeleton could be generated at the same time, and the relationships between them preserved. Figure 1 shows the crust and skeleton for a simple polygon. Each of the crust (boundary) edges drawn is a Delaunay edge, because of the previously-mentioned empty circle condition, and the associated Voronoi edge (passing between the two Delaunay vertices) is suppressed. Similarly, the skeleton (or medial axis) consists of those Delaunay edges that failed the circle test, and in this case the dual Voronoi edge is drawn, and the Delaunay edge suppressed.

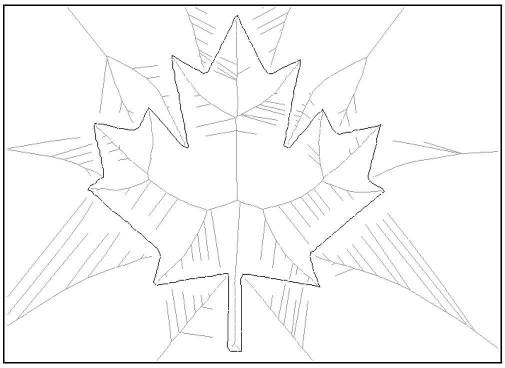

Fig. 1 Crust and skeleton

The resulting form is quite satisfactory (except perhaps at sharp corners where the sampling condition is not met). However, the boundary was generated from an image, and therefore has a certain sampling error. This is reflected in the "hairs" on the skeleton, which are formed by three adjacent co-circular vertices. It would be desirable to trim these. Gold (1999) shows that this can be achieved by "retracting" the leaves of this skeleton, as shown in Figure 2.

4. CURVE GENERALIZATION

The key idea applied here to generalization processes is the concept of *skeleton retraction*. The idea is quite simple - simpler objects have simpler skeletons which means simpler shapes (Blum 1967). (This is also true for the space between objects - retraction here simplifies the common Voronoi boundary between them.) This approach is only made possible because we can preserve the link between the skeleton and the boundary (or "crust"). This paper only describes some of the initial steps in this controlled retraction process - that of retracting each leaf of the skeleton tree to the same location as its parent. This provides a skeleton with fewer branches, and a crust/boundary with fewer minor perturbations. More advanced retraction procedures are possible given certain knowledge rules about the desired shapes of the objects concerned. Some of these issues are discussed in Ogniewicz (1994) and Ogniewicz and Ilg (1990).

The work described here is a continuation of our previous efforts (Gold et al., 1996; Gold, 1997; Gold, 1998) on data input. Briefly, in the first paper we developed a digitizing technique for forest polygons based on digitizing around the interiors of each polygon, giving the points the polygon label, creating the Voronoi diagram, and extracting the Voronoi boundaries between points with different labels. This was rapid, and guaranteed to be topologically correct. In the second paper we extended this to scanned maps, by deriving "fringe" points on the interior of each polygon by image filtering, labelling these points with a floodfill algorithm, and extracting the polygon boundaries as just mentioned. In the third paper we show that the work of Guibas and Stolfi (1985), in developing the Quad-Edge structure, may be extended to connected maps by permitting the simple edge between two vertices to become an arc (the "Quad-Arc"), with multiple points along it.

The properties of the Quad-Edge/Quad-Arc led us to examine our previous techniques. The Quad-Arc allows the management of any connected graph on the plane or orientable manifold. However, our earlier algorithms required closed polygons, with different labels, in order to identify the edges we wished to preserve. Thus unclosed linework would be lost, as the polygon label would be the same on both sides of the line. However, we could not find a satisfactory algorithm based on point distribution alone, without some labelling function.

In 1998 Amenta et al. published a seminal article where they showed that the connected set of boundary points (the "crust") could be extracted using the Voronoi diagram. Their objective was to eliminate all Delaunay edges that crossed the skeleton of the polygonal object under consideration, with the remaining edges forming the crust. They achieved this by first generating

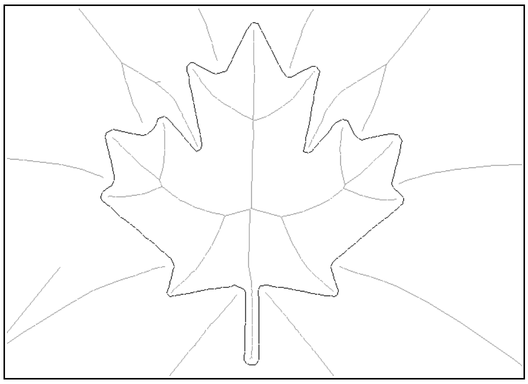

Fig. 2 Smoothing by skeleton retraction

the Voronoi diagram of the boundary points, and exporting the Voronoi vertices. In a second step, they constructed the Delaunay triangulation of both the original boundary points and the new Voronoi vertices. Any Delaunay edge connecting two boundary points must be part of the desired crust, as Delaunay edges that originally crossed from one side of the polygon to the other would now be cut by the newly-inserted "skeleton" points formed from the original Voronoi vertices. They showed that if the samples along the boundary were not more than (approximately) 0.25 times the distance to the skeleton, it was guaranteed that the complete crust (and no other edges) would be preserved.

However, our previous work had used the skeleton (a subset of the Voronoi edges) to define our topologically complete polygon boundaries, not the crust. Gold (1999) and Gold and Snoeyink (1999) showed that the two steps were unnecessary - a local test of associated Voronoi/Delaunay edge pairs sufficed. Our previously-used Quad-Edge structure contained this information directly. The work of Amenta et al. reduced to a simple test on each Delaunay edge of the original Voronoi/Delaunay diagram, as previously described. Thus it was possible to generate both the crust and the skeleton in one step, as a simple test on each Voronoi/Delaunay edge pair. Essentially, each edge pair (stored as the Quad-Edge structure of Guibas and Stolfi (1985)) would be assigned either to the crust or the skeleton.

In order to use our skeleton-based generalization algorithm, 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 Guibas and Stolfi (1985) 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. Gold (1998) 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, and islands may be connected to their enclosing polygon during extraction from the original Voronoi diagram.

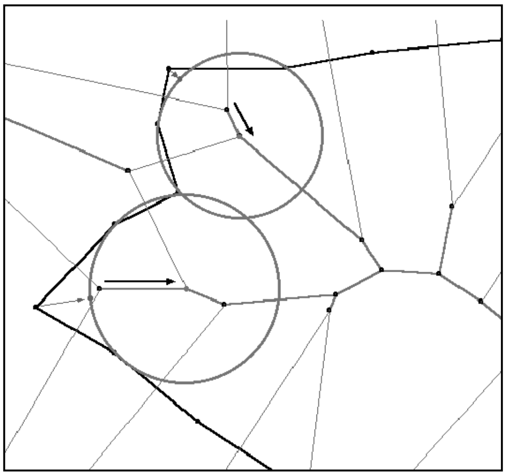

Fig. 3 Retraction process on a part of a polygon boundary

The second aspect that needs improvement is the existence of spurious "hairs" on the skeletons generated. This is a well-known artifact of skeleton generation, where any irregularities in the boundary generate unwanted skeleton branches. Ogniewicz attempted to reduce skeletons formed from raster boundary points to a simple form by retracting the leaf nodes of the skeleton until a specified minimum circumcircle was achieved, but with the development of the one-step crust and skeleton algorithm, this process may be greatly simplified.

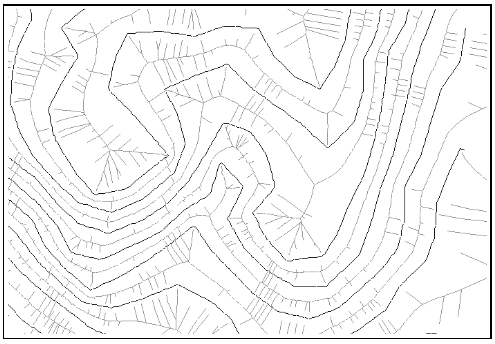

Fig. 4 Contour map and skeleton

Alt and Schwartzkopf (1995), as well as Blum (1967) showed that leaf nodes of a skeleton correspond to locations of minimum curvature on the boundary. For a sampled boundary curve this means that three adjacent points are cocircular, with their centre at the skeleton leaf. If we wish to simplify the skeleton we should retract leaf nodes to their parent node location. This means that we now have four cocircular points instead of three. The retraction is performed by taking the central point of the three defining the leaf node, and moving it towards the parent node of the skeleton until it meets the parent node circumcircle. This smooths outward-pointing salients in the boundary of the object. The same should be done from the other side of the boundary, retracting those salients also. This may displace some of the points involved in the first smoothing step, but as the process is convergent a small number of iterations suffices to produce a smoothed curve having the same number of points as the original, but with a simplified skeleton, as shown in Figures 1 and 2. The retraction process is shown in figure 3.

This process is interesting for several reasons. Firstly, we have retracted the leaf nodes of the skeleton, greatly reducing its complexity. This then gives us a resulting skeleton closer to Blum's idea of the Medial Axis Transform, but without the artifacts due to perturbation of the samples along the "ideal" boundary. Blum's motivation was to provide stable descriptors of shape for "biological" objects, and the skeletons we have derived are well suited to that. However, retraction of skeleton leaves may be appropriate in the middle of a curve, but inappropriate at, for example, the corner of a building. Such boundary points would need to be flagged prior to the smoothing operation. Note that this smoothing process has no metric tolerance associated with it - so there is no problem with scale. Convergence is based on cocircularity. (This is particularly useful in the case of contour

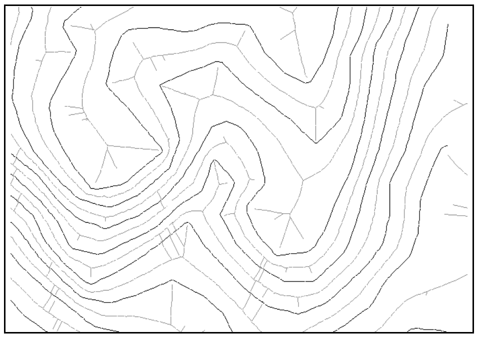

Fig. 5 Retraction smoothing

maps, as slope is perpendicular to the contour segments, so drainage is always towards skeleton nodes.)

Secondly, the skeleton may act as a linkage between adjacent objects - for example between adjacent curves of a contour map. The retraction process can not cause one curve to cross another (as can happen with simple curve smoothing), because the two crusts must have a separating skeleton. We thus get simplified curves that fit together. This is completely general, and provides a new model for map generalization. It is particularly relevant for contour map generalization or smoothing, because the usual one-dimensional curve smoothing procedures can not guarantee that adjacent contours do not overlap.

Fig. 6 Detail of Figure 4

In the case of contour maps, the result consisted of a connected set of contours if the curves were sufficiently well sampled. In between each curve is a skeleton, or medial axis, consisting of the Voronoi edges separating curves or portions of curves. Figure 4 shows the crust and skeleton for a simple contour map, and Figure 5 shows the smoothed version of Figure 4. Figures 6 and 7 show enlargements of portions of those figures.

The remaining secondary branches of the skeleton are valuable as they indicate the probable intermediate form of the surface between contours - such as minor ridges and valleys. Derived where there are re-entrants in a single contour, they approximate the human interpretation of the surface form, and can serve to improve runoff modelling. Problems may exist where the skeleton breaks through the crust, due to the sampling conditions of Amenta et al. (1998) not being maintained, especially at sharp corners. Additional filtering procedures are being developed to improve this situation. To get the best generalization, the algorithm must make several iterations, until no further points are displaced. In many cases a single iteration is sufficient, but up to five have been needed on occasion.

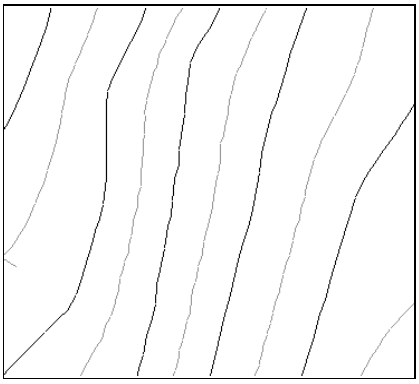

Fig. 7 Detail of Figure 5

Polygon generalization also gives excellent results, as shown in Figures 8, 9, and 10, where the original and the generalized curve are shown separately, and then superimposed. It can be seen that the form of the polygon has been preserved, and the skeleton simplified - but both are recognizably similar to the ungeneralized version. This conforms to the original concept of the Medial Axis Transform as a shape descriptor, as originally defined by Blum in 1967.

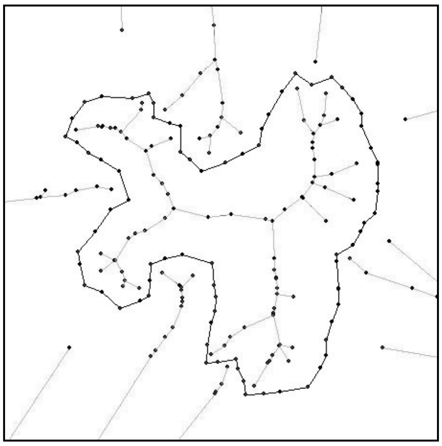

Fig. 8 Original polygon boundary

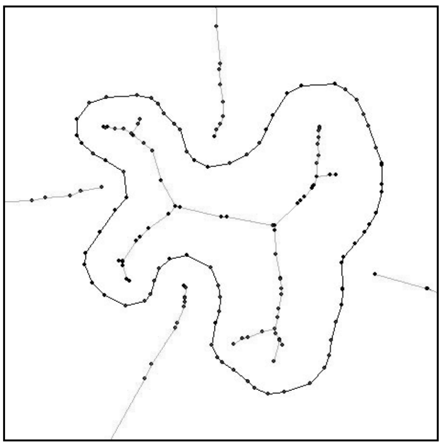

Fig. 9 Smoothed polygon boundary

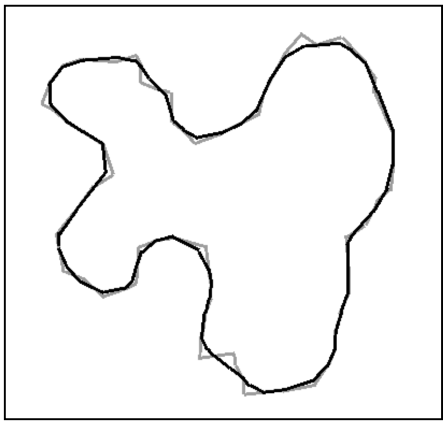

Fig. 10 The original and the smoothed curve superimposed

The current algorithm does not reduce the number of points - it merely simplifies the form of the curves. However, point reduction is made much simpler with the generalized curves. Research is ongoing on this topic. While excellent for individual polygons, generalization of connected polygon sets can rupture the crust if the node at the junction of three arcs is displaced as required by the algorithm. A method has been developed to flag such nodes as immobile.

5. TIN ENRICHMENT

The resulting diagram of the crust and skeleton may then be processed to extract the terrain model. The enriched triangulation consists of the triangulation of the original data points forming the crust, plus the Voronoi vertices that formed the skeleton. The next step is to estimate the elevation at each skeleton vertex. In most cases this is straightforward: the skeleton is at the mid point between two contours, and hence its elevation is half way between. This can be detected because each skeleton point is a Voronoi vertex, the centre of a circumcircle touching three of the original contour data points. If these vertices do not have the same elevation, the skeleton is part of an intermediate contour, with an elevation half way between the originals. If they do have the same elevation we have the case of minor re-entrants or promontories, where there are branches of the skeleton separating the two parts of the contour curve, connecting the main skeleton to the head of the ridge or valley - or else we have a closed contour forming a peak or pit.

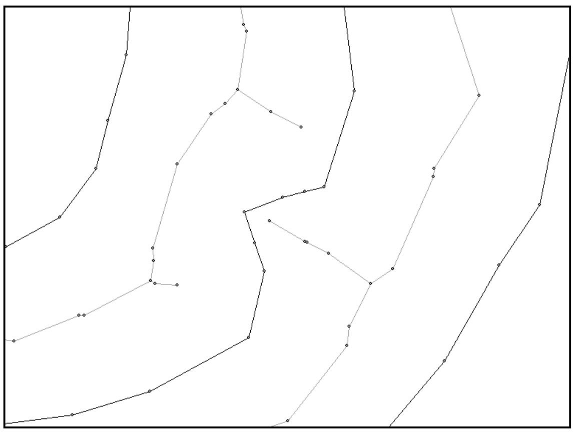

Fig. 11 Detected ridges

Various workers have worked on the problem of generating these intermediate points for an enriched triangulation. In particular, Aumann et al. (1991) used raster techniques to estimate the skeleton of these re-entrants. In our case these are obtained automatically from the Voronoi/Delaunay construction, with the crust/skeleton classification. Figure 11 shows some detected ridges. Each of these points is a Voronoi vertex, and it's circumcircle touches both sides of the contour re-entrant. As minor valleys become narrower towards their heads, the radius of any skeleton point on the re-entrant, by comparison with the skeleton circle at the junction of the branch with the main medial skeleton, gives the elevation of skeleton point as a proportion of the elevations of the medial contour and the one forming the re-entrant. Figure 12 shows the method, and some of the calculated results. In the case of peaks and pits, where only contour points of the same elevation are available, and no medial contour, slope information is necessary.

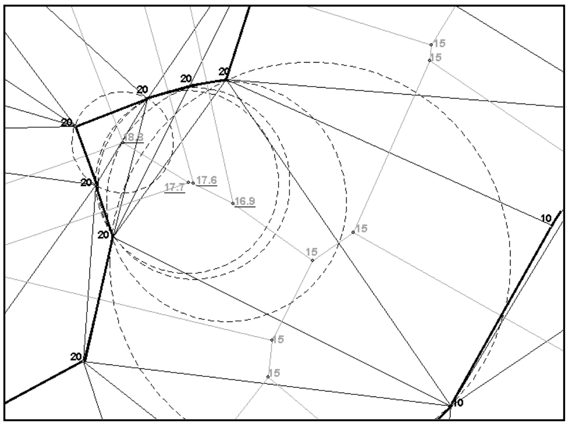

Fig. 12 Interpolation method using circle ratios

An advantage of the crust/skeleton model for contours is that slopes are obtainable directly. As noted previously, the line of maximum slope is perpendicular to any contour segment. In addition, Voronoi edges are perpendicular bisectors of their Delaunay edges. Thus, for all contour crust edges, the dual Voronoi edge gives the line of maximum slope. If the elevations of the Voronoi vertices have been calculated in the previous step, as just described, slopes may be ascribed to each crust edge, and then, by averaging, to each sample point on the contour. For peaks and pits a final step is needed to assign elevations to the interior Voronoi vertices, based on these slope values as they converge on the central skeleton.

6. CONCLUSION

On the basis of the crust/skeleton model just described, it is possible to build the components of a reasonable TIN model, using information about the supplementary skeleton points and the slope at contour segments. This is consistent with human interpretation of the original contours. Slope values at contour segments are also valuable for the generation of slope maps or other derived products. If a grid is desired, the area-stealing interpolation method may be used on the original data, the slopes at data points, plus the intermediate Voronoi vertices if required. If a triangulation is the preferred model, then the Voronoi vertices may be used to densify the original triangulation of the contour data. In each case the crust/skeleton model takes advantage of the human judgement implicit in the original contour construction, and mimics human interpretation of ridges, valleys and slopes, while preserving the basic topological consistency of the original data.

7. REFERENCES

Alt, H. and O. Schwarzkopf, 1995. The Voronoi diagram of curved objects. Proc. 11th Annual ACM Symposium on Computational Geometry, pp. 89-97.

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.

Aumann, G., H. Ebner, and L. Tang, 1991, Automatic derivation of skeleton lines from digitized contours, ISPRS Journal of Photogrammetry and Remote Sensing, 46 259-268, Elsevier Science Publishers B.V., Amsterdam.

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.

Carrara, A., G. Bitelli, and R. Carla, 1997. Comparaison of techniques for generating digital terrain models from contour lines, Int. J. Geographical Information Science, Vol.11, no.5, pp. 451-473

Garcia, A.B., C.G.Nicieza, J.B.O. Meréé, and A.M. Diaz, 1990. A contour line based triangulating algorithm. Proc. 5th Int. Symp. on spatial data handling. pp 273-284

Gold, C.M., 1989. Chapter 3 - Surface interpolation, spatial adjacency and G.I.S. In: Three Dimensional Applications in Geographic Information Systems (J. Raper, ed.), Taylor and Francis, Ltd., London, pp. 21-35.

Gold, C.M., 1997. Simple topology generation from scanned maps. Proceedings, Auto-Carto 13, ACM/ASPRS, Seattle, April, pp. 337-346.

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.

Gold, C.M., 1999, Crust and anticrust: a onestep boundary and skeleton extraction algorithm. In Proceedings of the ACM Conference on Computational Geometry, Miami, Florida, pp. 189-196.

Gold, C.M. and Snoeyink, J., 1999, A onestep crust and skeleton extraction algorithm. Algorithmica, (in press).

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.

Gold, C.M., D. Thibault and Z. Liu, 1999. Map Generalization by Skeleton Retraction. ICA Workshop on Map Generalization, Ottawa, August 1999.

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.

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.

Ogniewicz, R. and M. Ilg, 1990. Proc. 4th International Symposium on Spatial Data Handling, v. 1, pp. 15-24.

Okabe, A., B. Boots and K. Sugihara, 1992 Spatial Tessellations - Concepts and Applications of Voronoi Diagrams. John Wiley and Sons, Chichester.

Robinson, G.J., 1994 The accuracy of digital elevation models derived from digitised contour data, Photogrammetric Record, 14(83): 805-814,

Weibel, R., M. Heller, 1990. Digital terrain Modelling, In: Maguire, D.J., M.F. Goodchild, and D.W. Rhind, (eds.) Geographical Information Systems: Principles and applications, London; Longman.