ON THE FAITHFUL CONVERSION OF CONTOURS

Albert H J Christensen
7836 Muirfield Court,
Potomac, MD, 20854, 301 983 9004
E-mail: ahchris@erols.com

The arrival of computers more than four decades ago enabled cartographers and other scientists to represent single-value functions in a number of ways other than contouring, until then the prevailing way of portraying those functions. The new approaches to representing single-value functions, particularly the grid and the Triangular Irregular Network (TIN), were and still are more amenable to computer processing although being far less graphically powerful than contouring. The advantages of processing grids and TINs with a computer, together with the consideration of the World's one million plus topographic sheets, gave birth to the processes known as contour-to-grid and contour-to-TIN conversions. However, despite the affordability of large computers and the power of ancillary software brought in by the ensuing decades, the accurate conversion of contour sheets into grids or TINs is still beset by seemingly insurmountable problems. The cause of those problems, also found in many other computer applications, is the limitation of computers to understand a "whole" picture other than a small piece at a time. This limitation led to the establishment of the entirely new discipline known as Pattern Recognition, and in particular to the development of techniques commonly called Medial-Axis Transformations (MATs). This presentation endeavors to demonstrate how such a Pattern Recognition technique and the time-honored cartographic device of waterlining can be assembled into an efficacious Contour-to-TIN conversion procedure. The demonstration will be centered in the fundamental property of the results: that of verifying the ruled surfaces defined between adjacent contours.

The Title

In the context of this paper "conversion" means transforming a set of digital contours into another set of geometric entities which could be handled in a simpler and more straightforward manner than contours. Although a contour map, or a set of digital contours, is the most flexible and most graphically powerful terrain representation, the machine handling of contours is problematic. Therefore, since computers began to be used in the process and analysis of topographic data, contours have been converted into more easily manipulated structures. The structure of choice was then, and still is, the uniform grid of elevations. Although proposals of variable grids were made as the shortcomings of uniform grids become understood, see for instance [HDC, 1977], alternatives to the uniform grid never really took off. The major US producers of digital terrain data, NIMA and the USGS, still publish their models as uniform grids. However, the reason for this loyalty is not just simplicity in the process of conversion set up by those agencies and their contractors but simplicity in the exploitation of the converted data as well.

The qualifier in the title has to do with the model topographers; civil engineers, and cartographers relied on since isolines were invented. When presented with a topographic map, those users understood the contours as being intersections of the ruled surfaces that approximated the topographic surface. Consequently when they interpolated intermediate points and generated profiles they assumed that the terrain model behaved linearly in between adjacent contour lines. Why should have they assumed any other type of behavior? If there was an accident between two contour lines that deserved recording, the topographer would have depicted it in the map, if need be by tracing a supplementary contour or by drawing rocks. The principle behind this model is that of parsimony, so powerful that an otherwise very comprehensive text book used long ago by this author, disposed of the matter with a single line, stating that, "...of course, in between contours the variation is linear [Manfredi, 1950].

The conversion informed in this paper, for brevity's sake and for the reasons to come later, is designated as "MA conversion". The "faithfulness" in the title refers to the linear model. Furthermore, it aims to set in opposition the conformity of TINs with the linear model to the correlative nonconformity of the uniform grid.

That nonconformity of uniform grids with the linear model has been abundantly proved by reports on unwanted features, diversely known as fantasies, unnatural features, or artifacts, found in the grids. See for instance [USGS, 1990]. Specific among those features are the "ramps", "stars', "false depressions" or "dams", as well as low frequency noise. As a rule, those unwanted accidents do not show in correspondence with areas densely cover by well-shaped contours. However, in sparse areas or in areas with contours that do not exhibit a strong spatial coherence, the gridding algorithms tend to create fantasies.

The MA conversion was designed with the purpose of avoiding the creation of those fantasies, by rigorously honoring the linear model. This distinction is certainly not the only that can be established between a particular TIN and grid methods. On the contrary, the differences are so many and so significant that cannot possibly be considered in this paper.

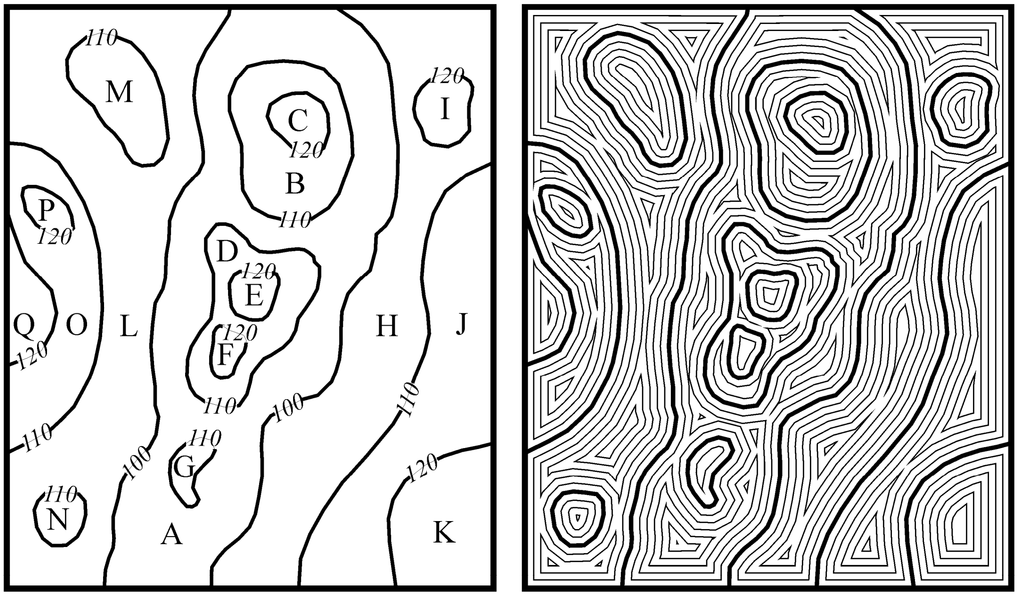

Figure 1a. Input contours and Units (A to Q)   Figure 1b. Waterlining of the Units

Consequently, this paper is limited to a description of the MA conversion; the practicalities of its application, and the enhancements required to eventually making it an efficient production tool.

The principle of the MA Conversion

Within the context of this paper, the best-known characterization of contour lines says that they are unmanageable by machines. That is quite right, but not very helpful. This paper proposes to round it off, by adding that they are unmanageable because they provide insufficient data. Now this characterization shows the way to a solution. To handle contours machines need supplementary data.

As regards this supplementary data, the original proposal --- so far as this author knows---, was made by himself back in 1987 [Christensen, 1987], after the ideas discussed with other colleague as early as 1975, at the National Research Institute for Mathematical Sciences in Pretoria. However, in 1987 personal computers were puny by current standards and the only implementation of the concept of extra data the author could possible attempt was a brute force approach, reported in that year [Christensen, 1987], and limited to diminutive data sets. However, the principle was sound, as this paper and the presentation at the ICA’2001 will hopefully prove.

The principle is quite simple. In order to force the TIN to honor linearity between two adjacent contours ***a*** and ***b***, a network of lines **MA** should be available that flow in between and at equal distances from ***a*** and ***b***. Furthermore, MA should be constructed to make sure that every edge in ***a*** and ***b*** could be linked to a vertex in **MA** without intersecting any other edge in ***a***, ***b***, or ***MA***. Then, to honor the linear model it would suffice to assign to the **MA** the mean of the elevations of **a** and ***b*** and to force every triangle to straddle the space between contour and MA. The special case of ***a*** and ***b*** having the same elevation will be discussed later.

The previous description of **MA** fits the concept of the structure known as the Medial-Axis of a shape, first introduced by Blum [1973]. Medial-axes can be determined by so-called medial-axis transformations (MAT). Most of reported MATs operate in raster mode. Fewer in vector mode. As it can be said of almost all spatial processes, the preference for the raster mode is explained by the simplicity of the raster approach over its vector counterpart. On the other hand, vector MAT are more flexible and, in some sense, more powerful. The method applied in the MA Conversion, similar to the one originally proposed in (Monatanari, 1969), operates in vector mode throughout. Both methods are based on the gradual and systematic shrinking of the shape’s perimeter. The lines produced after each shrinking, designated as “waterlines” in (Christensen, 1999) because of their 18^{th} Century European prototype, show singular points that outline the medial-axis of the shape (Figure 1b).

The Current Implementation of the MA Conversion

The method is organized in several steps, as follows.

Delineation of the neatline.

In the current implementation of the MA conversion, neatlines are determined using solely the ends of clipped contours, provided that the original clipping lines had been either straight or close approximations to circles of very large radii of curvature. As the method can be applied in spherical coordinates (provided the project area is small relative to the globe), the delineation of the neatline is a very simple operation. However, if the triangulation should be executed on the plane, to ensure the success of the neatline delineation, the projection should be Mercator or similar. This condition is not restrictive at all: as TINs are invariant with projection changes --- a big plus if one compares them to the uniform grids---, it will do no harm to change the contour coordinates from the native projection to the Mercator, and to change them back once the TIN is completed.

Construction of the complex polygons and the contour tree

Once the neatline is determined, the contours are grouped into three classes:

1. Contours clipped at the neatline;
2. Contours wholly inside the neatline;
3. Contours interrupted (significantly) inside the neatline.

Next the ends of adjacent Class 1 contours are joined by sections of neatline, after which the assemblage is arranged so its listing sequence is anti-clockwise. As the area computation routine included in the software library yields positive results for anticlockwise-listed polygons, the polygons fashioned with Class 1 contours are designated as “positive”. Next the Class 2 contours that are found inside positive polygons are listed clockwise, --- which turns them into “negative” polygons---, and appended to the positive polygons. A positive polygon plus its retinue of negative polygon is really a complex polygon, but for the purpose of this paper and for brevity’s sake such a polygon is called a “unit”. At later stages of the procedure, the medial-axis and the triangulation established inside a unit are also thought as parts of the unit. Anticlockwise-listed Class 2 contours, and all the smaller negative polygons they may contain, are also units. Class 3 contours must be completed before they can be incorporated into the other two classes. See Section on Future Enhancements.

Once this assemblage is complete, the units can be recognized as being a tree of areas, with the root being the area enclosed by the neatline. The MA conversion never resorts explicitly to that tree. The only evidence of its existence is the sequence in which the units are listed in core or in auxiliary storage. In Figure 1a, units are indicated by capital letters.

The Waterlining of the Units

From this previous step on, the computer procedures that implement the MA conversion operate on a unit at a time, a remarkable benefit for both processing time and storage. Waterlining proceeds along the way described in [Christensen , 1999].

The operation must be preceded by the selection of a number of parameters. This step requires some insight into the operation of the procedure, particularly in what concerns linear tolerances as well as the constants that determinate the intervals between subsequent waterlines. The first interval is the distance between a contour and its first waterline. The subsequent intervals can be constant or variable, with the variability being linear or quadratic. As a rule, the parameter should be chosen so that there would enough room anywhere in the unit to trace at least four waterlines. This requirement may be, up to a point, satisfied by a self-calibrating mechanism incorporated into the waterlining procedure.

The construction of the medial-axis

The reader is referred to [Christensen, 1999] for a discussion on the pros and cons of vector vis-à-vis raster approaches to the Medial-Axis Transformation.

For each unit the waterlining-based MAT produces a MA. The MA is structured as a network, fully described in two data sets: a directory of lines and a linked list of arcs vertices. The directory of lines, of the usual sort in spatial application, can be easily inverted into a directory of nodes. The list of vertices deserves perhaps a little more attention. In that list each vertex is associated with a set of integers as well as with a real number. The integers point to the contour edges from which the vertex originated. Those integers are passed form one waterline to the next, to be finally assigned to the vertices in the MA. By construction, they point to the nearest contour vertices in the unit at both sides of the MA. The real number is a measure of the distance from the vertex to the contour at either side. Pointers and distances are elements basic to the step of the conversion described next.

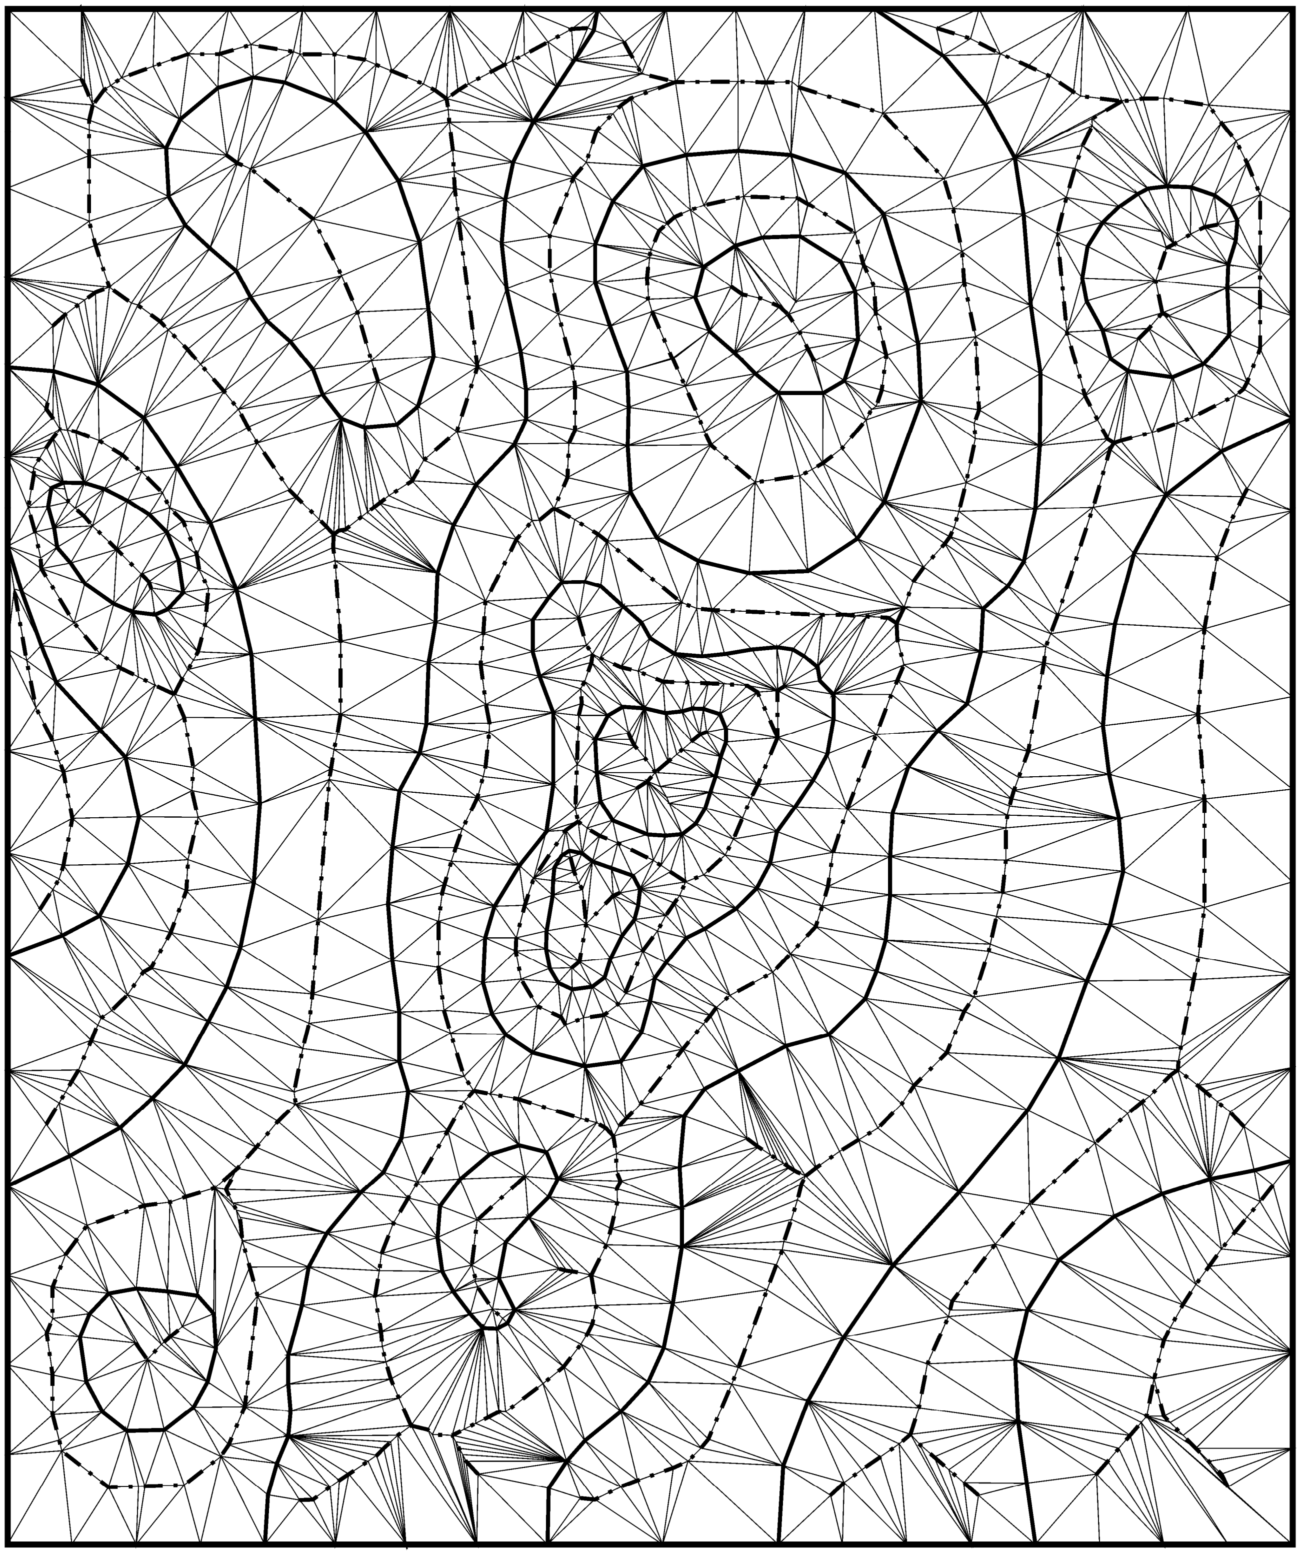

Figure 2. Contours (heavy, solid); medial-axes (dot-dash), and triangulation (light)

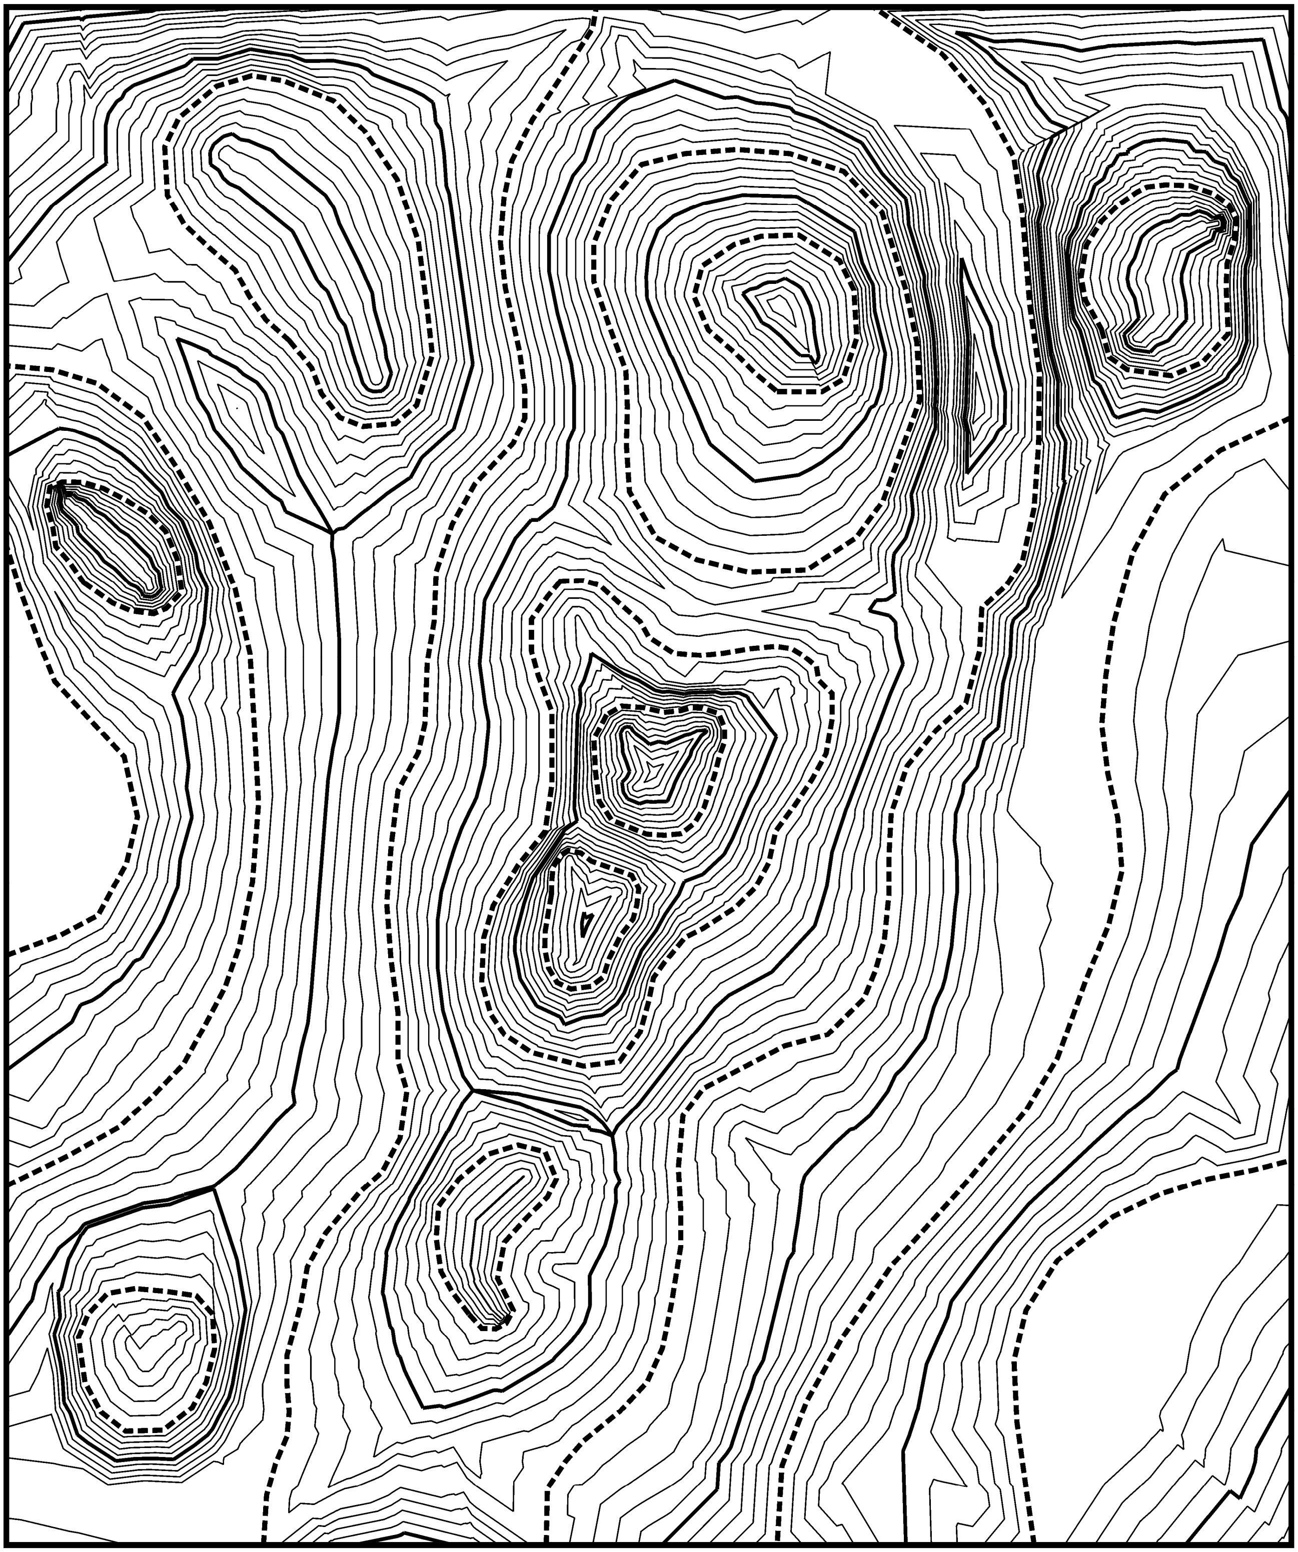

Figure 3. Derived contours, with input contours in heavy dashed lines

The synchronization of the MA and contours

To be able to triangulate the unit without exhaustive searches and checks, it is necessary first to restructure the MA network into a list that describes a closed polygon. Once that step is over, the software changes the ordering of the MA list to match the list describing the perimeters of the unit. Here, “to match” the two lists means that a point **P** that moves down the MA list one vertex at the time, and a point **Q** that does the some down the perimeter list, define a distance **P**;**Q** that never increases monotonously. If one imagines the triangulation of the unit to be a temporal process, one could very well call such reordering a synchronization. The transformation of the MA network into a MA list entails this duplication of all MA arcs, which is no waste if one considers the other alternative, that of synchronizing a complex boundary, with its positive and negative loops, with a network whose arcs have an internal listing that depend on the shape of the unit. Finally, the origins of two lists are shifted, to ensure that the distance between is a minimum. This step is repeated for every unit in the project area.

The Triangulation of a Unit

With the MA and the perimeter of a unit thus synchronized, the procedure triangulates the unit by applying the Delaunay test to the two potential triangles that share an already established base. The third vertex must be selected from either the MA or the contour. The new base is the edge of the new triangle that is supported on both the MA and contour. There are exceptions in this procedure, however, where the MA configuration is defective or where the contours are too angular. If an exceptional configuration is found, the software makes the best triangle that is supported only by one of the lists, usually the MA, and flags the triangle as pathological.

Why such occurrences are deemed pathological? To find the answer one must return to the principle, the linear model. If the MA has constant elevation, which is always the case with MA arcs in between contours of different elevations, a triangle with its three vertices on the MA would be horizontal. And that would violate the linearity principle. However, pathological triangles do not happen to be supported by MA vertices of the same elevation. As figure 5 shows, pathological triangles appear in the angle formed by two arcs of the MA, and one of those arcs is always given variable elevations, as it will be explained later. By construction, it is always possible to fashion a triangle that has at least one vertex on the MA. Therefore the instances are precluded of triangles that are indeed horizontal because they have their three vertices on a contour. Such horizontal triangles may show as narrow flat areas along contours in constrained Delaunay triangulations of contour maps. Their occurrence is behind one of the criticisms leveled against those solutions. See [Zhu, 2000] for an analysis and a proposed solution to the problem.

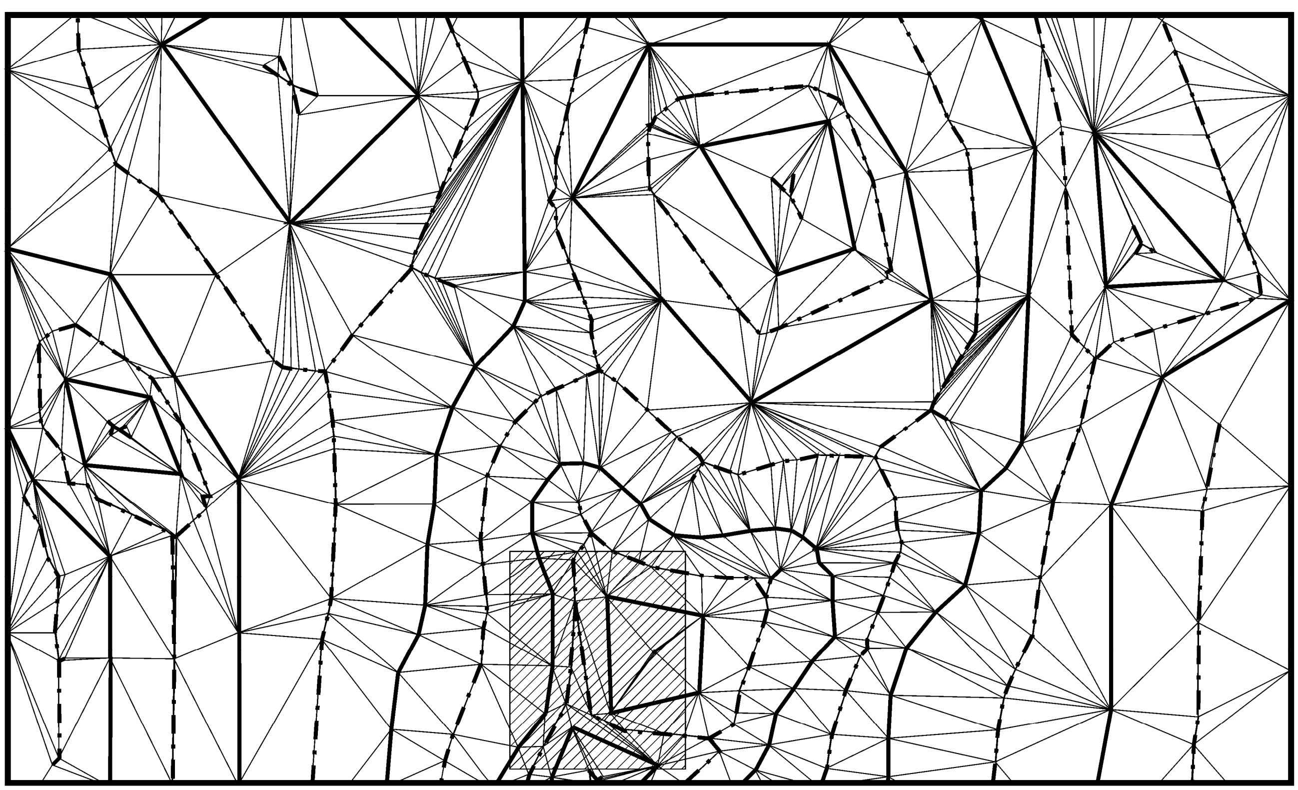

Figure 4. Triangulation of very angular contours. The cross-hatched area is enlarged in Figure 5.

The next section explains why a MA-supported pathological triangle may indeed look like a welcome member of a TIN.

The result of the triangulation of a unit is a simple list in which triangles are sets of six integer numbers. Three of them define the triangle vertices by pointing to the MA and the contour lists. The other three point to the adjacent triangles. Reciprocal pointers are recorded as attributes for every edge in the MA and in the contours.

The Assignment of Elevations to the Medial-Axis

This task, performed on a completed triangulation, is very simple in areas of a unit delimited by contour of different elevations. On the contrary, it is considerably complex if the contours have the same elevation, as it is necessary to decide whether the MA should be higher or lower than the contours. That decision is reached by analyzing the elevations of the adjacent units. As the process of the units proceed sequentially, it may happen that adjacent units may not have their MAs assigned yet. It is plain then that the assignment of elevations must be an iterative process. The elevations are estimated by first computing the slopes of the nearest triangles of adjacent units, and second, by raising or lowering the vertices of the MA to planes constructed with those slopes. This assignment of elevations is intricate and its complete description, for lengthy, is better left to a future paper. Let be said, though, that the elevations assigned to MA arcs in between contours of equal value cannot be constant. For each vertex they vary in proportion to the distance from the vertex to the end nodes of the arc.

As mentioned before, pathological triangles are usually nestled in between MA edges that are incident with a node. Additionally, it is always possible to verify that the angle formed by those two edges is rather narrow, and that at least one them belongs to an arc surrounded by contours of equal elevation. Hence the pathological triangles constructed on such arcs are not horizontal. But even if they were horizontal, they should be deemed anomalous, inasmuch as they should be considered as either at a crest or at the bottom of a valley. The concern for the pathological triangle is more of organization than shape. A triangulation, as any other data set, is handled better if there are no exceptions. And in a triangulation where triangles supported by both contour and MA is the rule, pathological triangles are the exceptions.

The Figures

The triangulation procedure has been tested with several data sets, the largest being a section of GEBCO-501, a bathymetry map with over 150 units and 50000 vertices. For this paper a much smaller data set was put together, that could be plotted distinctly in an 8.5 by 11 page. The contours were extracted from different areas of a USGS quad and placed so that they would configure several features in a single page that would be particularly critical to a contour-to-grid conversion.

The input contours assembled into units are shown in Figure 1a. Figure 1b is a waterlined map of those units. These waterlines were created for display only. The real ones, with which the test was run, are too close to be shown distinctly at the scale of the figure. Contours, medial-axes and triangulations are shown in Figure 2. The project area is 1000 units wide. In Table A the reader will find a selection of the columns of the statistics file compiled during the triangulation. A quick and dirty contouring program produced the contour map in Figure 3, at 1/10 of the original contour intervals. The irregularity of contours along the neatline, and particular those at the corners, is due to lack of data. In a real case, the contours within the map sheet should have been appended to those without, up to a certain distance of the neatline. This procedure would not only correctly shape the triangulation and the derived contours but would also ensure the matching of triangulations across neatlines.

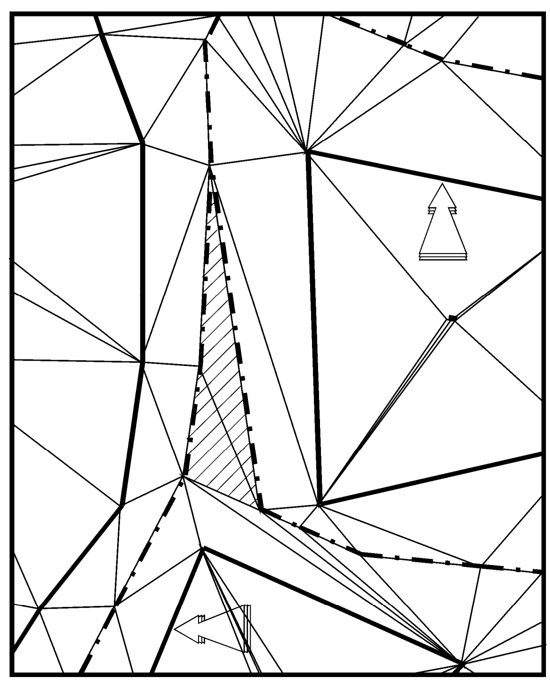

Figure 5. Two pathological triangles. Arrows point to angular input contours. Medial-axis as dot-dash lines

Figure 4 was prepared to demonstrate the influence that the degree of smoothness of the input contours has on the soundness of the triangulation. As it can be observed in this figure, the triangulations resulting from such angular perimeters are more irregular that those in Figure 2 and show more pathological triangles as well. Two of those pathological triangles are clearly indicated in Figure 5, an enlargement of the crosshatched area in Figure 4. To create the less desirable TIN in Figure 4, the contours in Figure 1a were thinned, so that some of them were reduced to quadrilateral and triangles. The arrows in Figure 5 point at two of them, both quadrilaterals. Such simple shapes are not impossible in real-life applications. They could be observed in overlays scanned at low resolution and thinned with large tolerances. If found in the input, the user would do well in clipping the sharp corners of the simple shapes or otherwise use a quadratic or a cubic interpolator on them before running the triangulation procedure.

Future Enhancements

The appendage of contours from adjacent sheets, mentioned in the previous section, is one of the enhancements planned for the MA Conversion. In the figures that illustrate this paper the reader will certainly notice the need for other enhancements. The next list contains those considered to be more pressing.

1. To produce smoother TINs, the criterion for selecting triangles should be modified. Whenever possible the improved algorithm would select triangles with bases on a contour rather than on the MA.
2. With the same purpose in mind, an effective thinning algorithm should be applied to the MA. The resulting triangulation would not only be smoother but will have fewer triangles as well.
3. MA branches (MA arcs with a dangling end) are required in some contour configuration in order to preclude the creation of horizontal triangles along the contours. As most of the branches are really unneeded, the current MA conversion tests them and clips or entirely removes them accordingly. The results shown in this paper prove that the test should be more stringent. If such enhancement is practicable as well as safe, the resulting TINs would be much improved, smoother and with fewer members.
4. As noted earlier, Class 3 contours, with endpoints at significant distances from the neatline, are mostly the evidence of feathering. They ought to be completed before they can be used in the conversion. The approach contemplated for this step is a preprocessing, that also requires triangulations. First a MA triangulation is established between any those pairs of Class 1 contours that contain Class 3 contours. Second, the resulting TINs are used to thread lines at the proper elevation from the corresponding ends of pairs of Class 3 contours. The completed Class 3 contours are then added to the Classes 1 and 2 and the triangulation of the project area is undertaken. As an aside, let it be said that this procedure could also shorten the rather outmoded stereoplotter task of manually dragging contours in steep areas.

On the Time Performance

The only steps of the MA conversion that require exhaustive searches are the first two, the shaping of units and their waterlining. Of those two only the waterlining step deserve attention as regards the time performance. Waterlining is a very onerous process in time and computer resources. The GEBCO sheet mentioned earlier needed about 20 hours of waterlining in a very reliable but small nine years old IBM RS6000. Today, machines of that line have improved greatly. With a new model of approximately the same price the process would take no more than 1/10 of the time. Even so, two hours is so much out of proportion with popular cartographic and GIS applications that the triangulation was designed to be based on any well-built MA, regardless of the MAT involved. Aside from waterlining, the process is fast. The process of the small data set used to illustrate this paper, waterlining included, took less than five minutes in the same old IBM. For information on the data size, please refer to Table A.

	Quantities		**Areas of**	
UNITS	**Vertices**	**Triangles**	**Polygons**	**Triangles**
1	320	487	280959.611	280959.611
2	74	109	54996.441	54996.441
3	19	32	11656.151	11656.151
4	115	201	46506.599	46506.599
5	31	52	8375.525	8375.525
6	17	37	5583.070	5583.070
7	15	32	7517.004	7517.004
8	182	267	199575.952	199575.952
9	17	37	14055.632	14055.632
10	81	121	130287.920	130287.920
11	47	64	54932.274	54932.274
12	239	351	242533.206	242533.206
13	15	38	35196.993	35196.993
14	11	24	9415.320	9415.320
15	112	160	72693.210	72693.210
16	17	34	7133.385	7133.385
17	21	34	18331.070	18331.070
Totals	**1387**	**2080**	**1199749.363**	**1199749.363**

Table A. Figures drawn from the Statistics File

References

Blum, H. (1973). Biological Shape and Visual Science. *Journal of Theoretical Biology*. 38, 205.

Christensen, A. H. J. (1987) Fitting a Triangulation to Contour Lines. *Proceedings of AutoCarto 8.* Baltimore, 3/30 to 4/3. pp. 57-67.

Christensen, A. H. J. (1999) The Revival of a Victorian Art: Waterlining with a Computer. *The Cartographic Journal*, 36, 1. pp. 31-41.

HDC (The Hydrological Engineer Center, USA Corps of Engineers) (1977). *Proceedings of a Seminar on Variable Grid Resolution – Issues and Requirements*. Davis, CA. August 18-19.

Manfredi, L. (1950). *Geometría Descriptiva*. O.B. Andrés, Editor, Buenos Aires. 2 volumes, 514 pages plus atlas.

Montanari, U. (1969). Continuous Skeletons from Digitized Images. *Journal of the ACM,* 16, 4. pp. 534-49.

USGS (1990). Procedure Manual for DEM Evaluation and Testing. U.S. Geological Survey, Reston, VA.

Zhu, H., and Schneider, K. (1999). Flat Feature Processes from Triangulated Irregular Networks for Hydrological Modeling. *Proceedings (CD) of GeoComputation Conference,* Mary Washington College, VA, July 25-29*.*