A Watershed Algorithm for Triangulated Terrains

Michael McAllister
University of British Columbia
mcallist@cs.ubc.ca

Abstract

We study the problem of automatically identifying watershed boundaries from digital elevation data. Nelson et al. have a vector-based algorithm to solve this problem, but their approach cannot guarantee that each watershed is a single polygon because of the modeling assumptions. We propose a vector-based algorithm that is provably consistent with its modeling assumptions and that guarantees one polygon per watershed. Our algorithm finds the watershed boundaries for local minima in the terrain in $\Theta(n^3)$ worst-case time.

1 Introduction

The watershed of a river, the region of land that is drained by the river, is a natural management unit for land preservation or development. Any land-based activity within a watershed can affect the health of the river that defines the watershed. Consequently, good land management requires accurate watershed boundaries.

Historically, watershed boundaries for both rural and urban areas have been manually traced from topographic maps. More recently, algorithms have been developed to identify watersheds from different digital terrain representations. The most common approaches use raster grids of elevations to locate terrain features such as ridges and valleys [1, 2, 7, 13, 17, 18, 23] or to simulate rainfall on the terrain [3, 5, 10, 12, 15] Unfortunately, there are difficulties with extracting watershed boundaries from raster-based algorithms. For example, the boundary edges for the watersheds are not explicitly defined. The edges and their connections to one another must be derived from the discrete raster structure.

Vector-based algorithms generate crisper watershed boundaries at the expense of using non-local properties of the data. The most recent vector algorithms use terrain-adjusted Voronoi edges as watershed boundaries [20] or groups of area elements of a triangulated terrain (a TIN [11]) formed by tracing paths of steepest descent on the terrain [6, 14, 16]. Although these latter algorithm are common, there exist degenerate terrain conditions where the algorithms can return incorrect results.

In this paper, we present an algorithm for finding watersheds that is provably consistent with the same modeling assumptions used by Nelson et al. and de Berg et al.:

- a TIN correctly approximates the drainage characteristics of the terrain,
- water follows the path of steepest descent on a terrain, and
- every terrain point has unique path of steepest descent.

Unlike Palacios-Velez and Cuevas-Renaud [16] and Nelson, Jones, and Miller [14], our approach for identifying watershed boundaries focuses on the edges of the boundaries rather than on the area between the boundaries. We identify a set of edges and paths on a TIN that form a superset of the watershed boundary edges for the local minima in the terrain and topologically connect these edges and paths into a graph. We prove that the graph is planar and that each face of the graph is the watershed of one local minimum on the terrain. Finally, a simple graph traversal algorithm extracts the watershed boundaries where each watershed boundary is a single closed and connected curve.

2 Definitions

A Triangulated Irregular Network (TIN) is a piecewise-planar approximation to a terrain where each facet of the approximation is a triangle. Our goal is to identify the watershed boundaries for local minima in a terrain when a TIN approximates the terrain and when water follows a unique path of steepest descent on the TIN. For a point p, let trickle-path(p) be the unique path of steepest descent from point p on the terrain. Then the watershed of a point q is $\{p \mid q \in \text{trickle-path}(p)\}$, that is all points whose paths of steepest descent include q.

The water flow assumption implies that water has neither volume to overflow negligible local minima in the terrain nor inertial to restrict the types of paths that water can follow. However, the simplicity of the trickle-paths make the paths locally reversible on faces of the TIN as paths of steepest ascent. Our algorithm uses paths of steepest ascent, called *trace-ups*, to delimit local areas that can flow into a point. For simplicity, we stop trace-ups when they encounter a ridge or saddle point in the terrain.

The ridges, valleys, and pits of a terrain govern its drainage characteristics [22]. A TIN edge e is a valley if the normals to the faces incident on e both point towards e. A TIN edge e is a ridge if the normals to its incident faces both point away from e. A point p is a pit if it is a local minimum in the elevation of the terrain. We impose a boundary condition where the TIN is surrounded by a cycle of ridges, as if the terrain appeared in a crater.

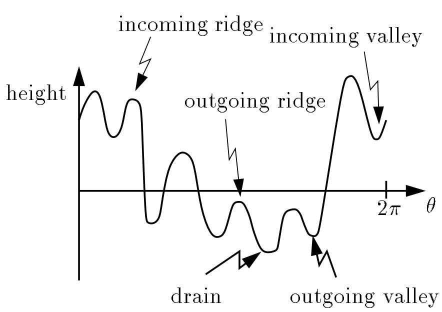

Figure 1: A sample height profile.

The valleys correspond to rivers in the terrain and the ridges correspond to watershed boundaries in the terrain. As noted by Palacios-Velez et al. and Nelson et al., not all watershed boundaries are edges of the TIN. Some watershed boundary edges can cross faces of the TIN. We use a height profile function to define a more general version of valleys and ridges around a point and to identify the cross-face edges.

Let the *height profile function* $h_{\epsilon,p}(\theta) : [0, 2\pi) \rightarrow \mathbb{R}$ at a point p to be a function of the angle θ that returns the elevation at which the cylinder $(x - p.x)^2 + (y - p.y)^2 = \epsilon^2$, $z \in \mathbb{R}$ intersects the terrain in at angle θ from p. We only consider the profile around one point whose elevation is assumed to be 0 and consider properties of $h_{\epsilon,p}(\theta)$ that are independent of ϵ (for small ϵ), namely its maxima, minima, and sign.

The local maxima and local minima of the height profile for a point p correspond to ridges and valleys at p. A local maximum identifies a direction where the flow of water splits around p: a ridge. A local minimum identifies a direction where the flow of water collects at p: a valley. We also distinguish between extrema whose elevation is higher than p as *incoming* features and extrema whose elevation is lower than p as *outgoing* (figures 1).

When the path of steepest ascent out of a point p is along a TIN face then that path follows a local maximum of the height profile function for p. Similarly, when the path of steepest descent out of a point p is along a TIN face then that path follows a local minimum of the height profile function for p. We use many trace-up paths from a single point p even though there is a unique path of steepest ascent; the multiple paths for p are the paths of steepest ascent that begin at an ϵ perturbation of p in the direction of each local maximum of the height profile function.

If the height profile value is negative at the global minimum then the direction of the global minimum is called the *drain* for the point. We assume that each point has at most one one drain.

3 Watershed Graph

We expect the watershed for a point p to be an irregular polygon that contains p, whose interior is a connected set, and whose boundary does not intersect itself. While natural watersheds meet these idealized conditions, the conditions are not consequences of our modeling assumptions.

We characterize terrains by their level of drainage complexity as nice terrains, normal terrains, or nasty terrains. Their watershed behaviours parallel the best, average, and worst-case analyses of algorithms. In *nice terrains*, the local maxima and minima for the height profile functions of every TIN point only occur at ridge and valley edges of the TIN. In *normal terrains*, the local maxima and minima for height profile functions can occur on faces of the TIN, but the trace-up paths along incoming ridges must end in the interior of ridges. Most TINs are normal terrains. In *nasty terrains*, there are no restrictions on the height profile function. Figure 2 shows three terrains that differ only in the elevation of a few TIN vertices but

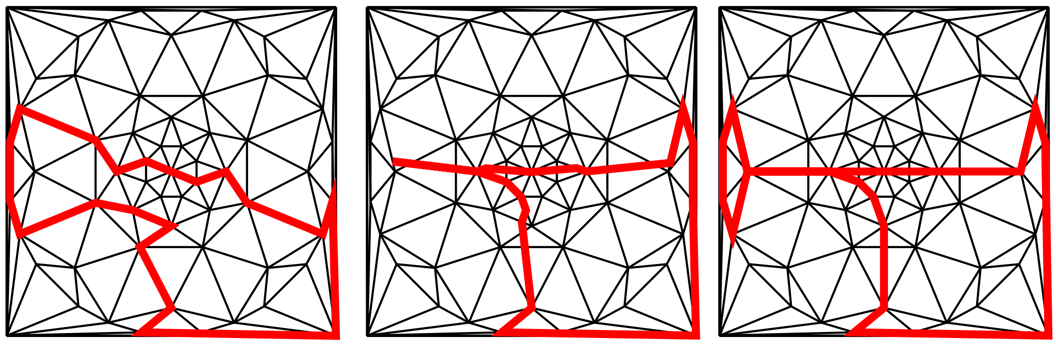

nice terrain   normal terrain   nasty terrain

Figure 2: Possible watersheds on terrains.

have very different watershed characteristics.

We define a watershed graph for each of these terrain types that abstracts the watershed boundaries on the terrain. The watershed graphs always have a planar embedding and each face of the graph corresponds to a single watershed in the TIN.

Section 3.1 describes the watershed graph for pits in terms of nice terrains. Sections 3.2 and 3.3 adapt the watershed graph for normal and nasty terrains.

3.1 Nice Terrain

In *nice terrains*, every local maximum and local minimum of the height profile function occurs along a ridge edge or a valley edge of the TIN. The set of ridge edges in the TIN form an embedded planar graph where each face either contains a pit or drains through one corner of the face. This graph is the basis for the watershed graph on nice terrains.

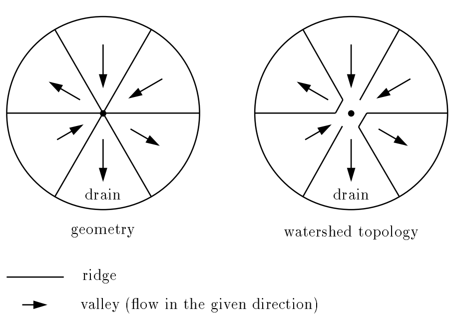

Figure 3: Watershed graph on a nice terrain.

Watersheds of nice terrains have two desirable characteristics: the interior of the watershed is a connected set and the boundary is not degenerate. These characteristics are typical of the results for watershed extraction algorithms in the literature and match our intuitive notion of a watershed.

We abstract the drainage characteristics as a *watershed graph*. Let S be the set of TIN vertices that are either endpoints of ridges or points with more than one outgoing valley. Let R be the set of ridges in the TIN and let I be copies of the elements of S; if point $p \in S$ has m incoming valleys then I has $m+1$ copies of p—one for each incoming valley and one for the drain. The vertices of the watershed graph are $R \cup I$. The edges of the watershed graph occur between an element of R and an element of I and represent the connection in the TIN between a ridge and its endpoints. If $r \in R$ is a ridge of the TIN with endpoint p then there is an edge in the watershed graph between the vertex of r and the vertex that corresponds to the first element of I that is clockwise from ridge r around p (figure 3).

For a planar embedding, the order of the edges around a vertex in the watershed graph is the same as the geometric order of the ridges that define the edges around the corresponding point in the TIN.

Lemma 1 *The watershed graph abstracts the local drainage characteristics on nice TINs.*

Proof: Let ν be a valley of the TIN around a point p. If ν is an incoming valley then $cw_p(\psi) = \nu$ and $cw_p(\tau)$ is the clockwise predecessor of ν in $U_p \cup \{u_{po}\}$; the face that contains ν is not closed at p and water along ν will reach p. If $\nu \neq u_{p0}$ is an outgoing valley bounded by ridges τ (clockwise of ν) and ψ (counterclockwise of ν) then $cw_p(\tau) = cw_p(\psi)$ since τ and ψ are only separated by ν and ν is not an incoming valley. The face of the watershed graph that contains ν is closed at p. If $\nu = u_{p0}$ and there is some incoming valley at p then $cw_p(\tau) = u_{pk}$, $cw_p(\psi) = u_{p0}$, and the face of the watershed graph that contains u_{p0} is open to anything that reaches the point p. Finally, if $\nu = u_{p0}$ and there is no incoming valley at p and p is a peak. All cw_p values at p are u_{p0} and all faces in the watershed graph are closed at p—no water runs around or over the peak. ∎

3.2 Normal Terrain

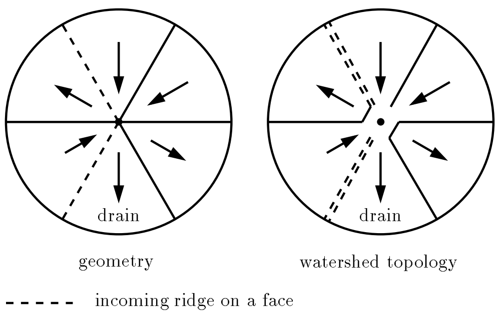

Figure 4: Watershed graph on a normal terrain.

Normal terrains relax the assumption of nice terrains: the local minima or maxima of the height profile can occur along faces of the TIN. In its place, we assume that all trace-up paths from any TIN vertex end at the interior of ridges.

The ridges on faces play a different role from ridge edges in determining the water flows around a point p. An incoming ridge on a face serves two roles: as a ridge that separates the water flows to either side of p and as a valley that collects the water along the ridge into p. We model the ridge with two parallel ridges that are separated by a valley in the TIN. An outgoing ridge on a face at p has no effect on the watershed boundary under the current assumptions.

The valleys on faces at p play the same role as valley edges of the TIN in dictating the topology of the watershed graph.

The incoming ridges along faces in the TIN change the characteristics of the watersheds. The interiors of the watersheds are still a single connected set, but the water along incoming ridges produce dangling edges in the watershed boundary.

There are two changes to the definition of the watershed graph from Section 3.1:

- incoming ridges along faces must have the effect of two parallel ridges, and
- incoming ridges along faces must be treated as an incoming valley.

Instead of creating a single edge in the watershed graph to represent an incoming ridge at a point p, we create two edges e_1 and e_2 for the ridge, and create a new graph vertex to represent the incoming valley between the two edges. Edges e_1 and e_2 are assigned different lower endpoints in the watershed group, but both meet the same ridge at their upper ends so their upper ends are connected (figure 4).

Lemma 2 *The watershed graph captures the local drainage characteristics of ridges of normal terrains.*

Proof: When an incoming ridge is modeled by two ridges that bound an incoming valley, the two ridges correspond to matching steepest ascent paths $t_{p,2i+1}$ and $t_{p,2i+2}$ in the watershed graph. Since the paths have different lower endpoints, the incoming valley bounded by the paths flows into the point p. ∎

3.3 Nasty Terrain

Nasty terrains have no assumptions on the height profile functions or on where trace-up paths stop. In particular, they allow trace-up paths to end at TIN vertices; the condition corresponds to a degeneracy between the TIN and the water flow model. Even though this does not normally occur in terrains, the water flow and terrain model assumptions do not forbid this behaviour so we must gauge its effects on the terrain's watersheds.

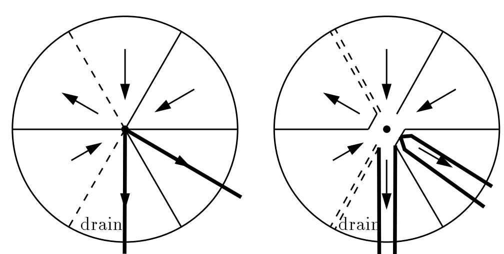

Figure 5: Topology of figure 4 with incoming paths of steepest ascent.

Trace-up paths only affect the drainage characteristics at incoming ridges along TIN faces. A trace-up path π that stops at a ridge creates a dangling edge in the watershed boundary and is handled in Section 3.2. A trace-up path π that stops at a TIN vertex creates a funnel along π. If the water flow assumption sends the water at p along π then π drains the land above p along a single line on the TIN. Path π can connect disjoint areas of the terrain in one watershed: the watersheds for nasty terrains can have disconnected interiors as well as degenerate boundaries.

The funnels, which appear as two parallel edges in the watershed graph, affect the topology of graph edges. In normal terrains, trace-up paths always meet at a common ridge so the corresponding graph edges meet at a common upper endpoint. In nasty terrains, trace-up paths can stop at a TIN vertex p that maps to several vertices in the watershed graph. The topology of the watershed graph in this latter case depends on the direction from which the paths approach p, both of which appear as dashed edges in figure 5. If the paths approach p from a non-drain direction then their corresponding graph edges meet at a common upper endpoint. If the paths approach p from its drain direction then the corresponding graph edges do not meet at their upper endpoints; the topology of the graph allows the drain of p to flow between the two graph edges.

Lemma 3 *The watershed graph abstracts the local drainage patterns on nasty terrains.*

Proof: Given a path of steepest ascent up to a point q in S, we have a pair of graph edges: $t_{p(2i+1)}$ and $t_{p(2i+2)}$. Only one such pair of graph edges can approach q from each of its outgoing generalized valleys. If the pair approach from an outgoing generalized valley that is not the drain, then the cw_q function uses the same point of V_q as the endpoint of the edges, so the face is closed at q. If the pair approach from the drain of q then the cw_q function assigns u_{q0} as the endpoint of the most-counterclockwise edge and v_{qk} as the endpoint of the most-clockwise edge. The two edges close a face only if $k = 0$, which implies that there are no incoming generalized valleys at q and that q is a peak. ■

Since the edges of the watershed graph are derived from ridge edges of the TIN and from paths of steepest ascent on the TIN, we automatically have an embedding for the watershed graph.

Lemma 4 *The watershed graph for a nasty terrain is planar.*

So far, all the properties of watershed graphs are local around individual points. However, the watershed graph implies more global properties:

- Every face of the watershed graph contains exactly one pit.
- A point p belongs to the same face as the pit to which it drains.
- If the terrain is enclosed by an outer ridge box then the watershed graph is connected.

The first two properties combine to prove theorem 5.

Theorem 5 *The faces of the watershed graph for a TIN outline the watershed boundaries of the pits of the TIN.*

The connectedness property allows us to use a simple face-tracing algorithm to extract all the watershed boundaries from the watershed graph. These boundaries will be simple polygons.

3.4 Watersheds of Rivers

The watershed graph of Section 3.3 identifies the watershed boundaries for the pits in the terrain. When we want the watershed boundaries of rivers, we must augment the watershed graph with extra edges.

Let ρ be the river whose watershed is desired and let p be the most-downstream point of ρ. The watershed of ρ is the watershed of p. The trace-up paths from p are the boundaries for the land that can drain through p and the land that drains directly to a valley point lower than p; the paths belong to the watershed boundary for p. By adding p as a vertex of the watershed graph to the set S and by adding edges for the ridges at p to the watershed graph, we divide a face of the watershed graph into two parts: one part for the watershed of p and one part for the rest of the watershed into which p drains. Consequently, with minor changes to the watershed graph, we can find the watershed boundaries for rivers in the TIN.

4 Algorithm Complexity

The watershed graph of Section 3 contains the watershed boundaries for every pit in the terrain. In 1996, de Berg et al. [4] constructed a terrain whose river system had an $\Omega(n^3)$ complexity. The same construction demonstrates that the watershed graph for the terrain can have $\Omega(n^3)$ points along the watershed edges. However, the number of graph edges in the watershed graph can still be linear in the size of the TIN.

We use a straight-forward algorithm to construct watershed graphs. The algorithm identifies the vertices of the watershed graph (TIN ridges and saddles) in linear time, finds the $O(n)$trace-up paths from the saddles and ridge endpoints in $O(nk)$ time where k is the path complexity, sorts the ridges and trace-ups around each vertex in $O(n \log n)$ time, and re-constructs the graph topology with a linear scan of the sorted lists.

Lemma 6 *The straight-forward algorithm constructs the watershed graph of a TIN in* $O(nk + n \log n)$ *time where* k *is the size of the paths of steepest ascent on the TIN.*

Since a path of steepest ascent on a TIN can have $O(n^2)$ points, k is $O(n^2)$ and the algorithm's time complexity $\Theta(n^3)$.

5 Sample Watershed

Under the restrictions of the TIN terrain model and the water flow assumption, the sequence of terrain types in Section 3 describe increasingly-complex yet consistent drainage systems. Not only are the terrains consistent, figure 2 realizes each of the terrain types as TINs.

The terrains of figure 2 were constructed to exploit the weakness of the TIN. In a more general setting, degeneracies do not have as big of an impact on the watersheds. Figure 6 depicts manually-traced watersheds and rivers in a 35 km by 35 km mountain area north of Vancouver, Canada. An interesting feature of the watersheds is that some tributaries of rivers have their watersheds traced separately. Without user input or river network orders [9, 19, 21], an algorithm cannot be expected to isolate these watersheds separately.

The lines of figure 7 are the watershed boundaries as detected by our watershed graph. The internal structure inside each watershed does not appear in the figure. The watershed graph has 64 000 edges with the polyline for each edge having an average of 3 points and the longest edge having 39 points. If the internal structure of the watershed graph is eliminated, leaving only the watershed face boundaries, then the watershed graph has 18 200 edges with the polyline for each edge having an average of 2.5 points and the longest edge having 29 points. The small size of the graph edges is an encouraging result with respect to the time complexity of Lemma 6 for constructing the watershed graph. Our watershed graph contains many more faces than there are watersheds in figure 6. The extra faces come from pits in the TIN that are not present in the actual terrain.

The TIN under the watershed graph was created from a raster of elevations with a 40 metre error tolerance and has the river segments embedded in the TIN to help correct the TIN's drainage characteristics. The TIN for Vancouver uses 41 840 points of the original 694 650 raster points and 125 520 TIN edges. The original raster data has a 10 metre accuracy in the xy-plane and a 5 metre accuracy in the elevation and was converted into a TIN by an adaptation of an algorithm by Heckbert and Garland [8]. This TIN was not preprocessed to orient flat triangles and horizontal edges for the algorithm that constructs the watershed graph. Nevertheless, the automatically-detected boundaries follow the line-work of the manually-traced boundaries in most cases.

6 Conclusions

We examined the problem of identifying watershed boundaries from digital elevation data. As with previous approaches, we represent the terrain with a TIN terrain model and assume that water always follows the path of steepest descent on the terrain. We derive from the TIN a planar graph whose faces are the watersheds of the terrain. Standard face-tracing algorithms follow the faces of the graph to extract the watershed boundaries. Unlike previous approaches, our graph guarantees that each watershed is a single polygon and that the watersheds are consistent with the TIN topology and the assumption on the direction for water flow.

In the worst case, the size of our graph and the time complexity of the algorithm that constructs the graph are optimal. Our graph has $O(n)$ edges and at most $O(n^3)$ points along the edges, which match the lower bounds described by de Berg et al. [4] for river complexity.

Our approach for finding watershed boundaries has a main shortcoming: it does not handle flat areas such as plateaus and low-slope rivers. This assumption fails on flat faces and horizontal edges of the TIN.

Improvements in two aspects of the terrain and river system can lead to identifying better watershed

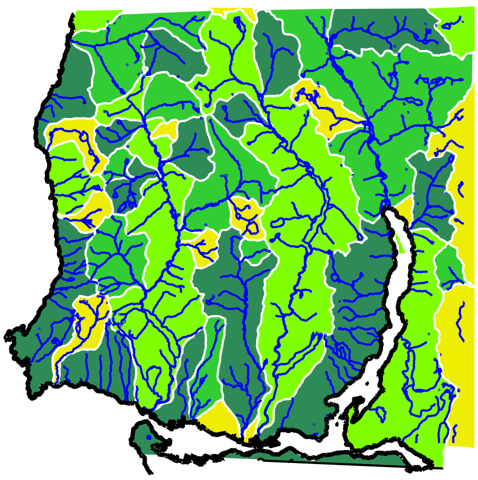

Figure 6: Rivers and digitized watersheds

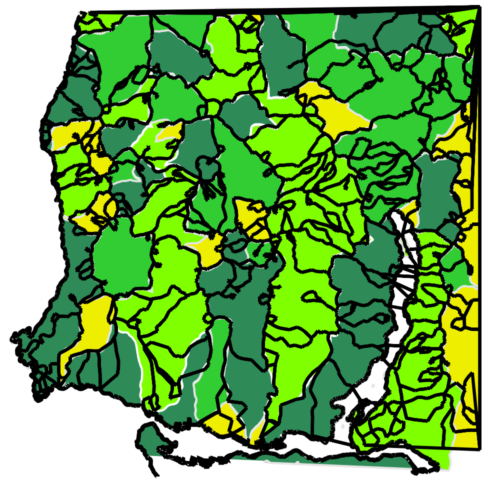

Figure 7: Derived (lines) and reference watersheds (shaded polygons)

boundaries. First, the terrain can be sculpted to better match the river system with the terrain valleys and to eliminate spurious pits in the terrain. Second, the hierarchical structure of the river network leads to a similar hierarchy in the watershed that we might exploit when tracing the boundaries of large rivers.

References

[1] L. Band. Topographic partition of watersheds with digital elevation models. *Water Resources Research*, 22(1):15–24, 1986.

[2] J. Chorowicz, C. Ichoku, S. Riazanoff, Y.-J. Kim, and B. Cervelle. A combined algorithm for automated drainge network extraction. *Water Resources Research*, 28(5):1293–1302, 1992.

[3] S. Collins. Terrain parameters directly from a digital terrain model. *Canadian Surveyor*, 29(5):507–518, 1975.

[4] M. de Berg, P. Bose, K. Dobrint, M. van Kreveld, M. Overmars, M. de Groot, T. Roos, J. Snoeyink, and S. Yu. The complexity of rivers in triangulated terrains. In *Proc. 8th Canad. Conf. Comput. Geom.*, pages 325–330, 1996.

[5] D. Douglas. Experiments to locate ridges and channels to create a new type of digital elevation model. *Cartographica*, 23(4):29–61, 1986.

[6] Environmental Modeling Research Laboratory, Bringham Young University. WMS. http://www.ecgl.byu.edu/wms.htm.

[7] P. Fua. Fast, accurate and consistent modeling of drainage and surrounding terrain. *International Journal of Computer Vision*, 26(1):1–20, 1998.

[8] P. S. Heckbert and M. Garland. Survey of surface approximation algorithms. Technical report, CS Dept., Carnegie Mellon U., 1995. CMU-CS-95-194, URL=http://www.cs.cmu.edu/ garland/scape.

[9] R. E. Horton. Erosional development of streams and their drainage basins—hydrophysical approach to quantitative morphology. *Geological Society of America Bulletin*, 56:275–370, 1945.

[10] S. Jensen. Automated derivation of hydrologic basin characteristcs from digital elevation model data. In *AutoCarto 7*, pages 301–310, 1985.

[11] M. P. Kumler. An intensive comparison of triangulated irregular network (TINs) and digital elevation models (DEMs). *Cartographica*, 31(2), 1994. monograph 45.

[12] D. Mark. Automated detection of drainage networks from digital elevation models. In *AutoCarto 6*, pages 288–298, 1984.

[13] L. Najman and M. Schmitt. Geodesic salency of watershed contours and hierarchical segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(12):1163–1173, 1996.

[14] E. J. Nelson, N. L. Jones, and A. W. Miller. Algorithm for precise drainage-basin delineation. *Journal of Hydraulic Engineering*, 120(3):298–312, Mar. 1994.

[15] J. F. O'Callaghan and D. M. Mark. The extraction of drainage networks from digital elevation data. *Comp. Vis. Graph. Image Proc.*, 28:323–344, 1984.

[16] O. L. Palacios-Velez and B. Cuevas-Renaud. Automated river-course, ridge and basin delineation from digital elevation data. *Journal of Hydrology*, 86:299–314, 1986.

[17] T. K. Peucker and D. H. Douglas. Detection of surface specific points by local parallel processing of discrete terrain elevation data. *Computer Graphics and Image Processing*, 4:375–387, 1975.

[18] J. Qian, R. Ehrich, and J. Campbell. DNESYS—an expert system for automatic extraction of drainage networks from digital elevation data. *IEEE Transactions on Geoscience and Remote Sensing*, 28(1):29–45, Jan. 1990.

[19] R. L. Shreve. Statistical law of stream numbers. *Journal of Geology*, 74:17–37, 1966.

[20] D. Skea, P. Friesen, J. Carr, A. Gallagher, I. Barrodale, E. Davies, F. Milanzzo, B. Corrie, and B. Noureddin. GeoData BC—TRIM watershed atlas. http://ssbux2.env.gov.bc.ca/srmb/twa_home.htm.

[21] A. N. Strahler. Quantitative analysis of watershed geomorphology. *Transactions of the American Geophysical Union*, 8(6):913–920, 1957.

[22] S. Yu, M. van Kreveld, and J. Snoeyink. Drainage queries in TINs: from local to global and back again. In *Proc. 7th Internat. Sympos. Spatial Data Handling*, pages 13A.1–13A.14, 1996.

[23] M. C. Zhang, J. B. Campbell, and R. M. Harlick. Automatic delineation of drainage basins within digital elevation data using the topographic primal sketch. *Mathematical Geology*, 22(2):189–209, 1990.