

I/O-Efficient Algorithms for Computing Contours on a Terrain *

Pankaj K. Agarwal
Duke University
Durham, NC

Lars Arge
University of Aarhus
Aarhus, Denmark

Thomas Mølhave
University of Aarhus
Aarhus, Denmark

Bardia Sadri
Duke University
Durham, NC

ABSTRACT

A terrain \mathbf{M} is the graph of a bivariate function. We assume that \mathbf{M} is represented as a triangulated surface with N vertices. A *contour* (or *isoline*) of \mathbf{M} is a connected component of a level set of \mathbf{M} . Generically, each contour is a closed polygonal curve; at “critical” levels these curves may touch each other or collapse to a point. We present I/O-efficient algorithms for the following two problems related to computing contours of \mathbf{M} :

- (i) Given a sequence $\ell_1 < \dots < \ell_s$ of real numbers, we present an I/O-optimal algorithm that reports all contours of \mathbf{M} at heights ℓ_1, \dots, ℓ_s using $O(\text{SORT}(N) + T/B)$ I/Os, where T is the total number edges in the output contours, B is the “block size,” and $\text{SORT}(N)$ is the number of I/Os needed to sort N elements. The algorithm uses $O(N/B)$ disk blocks. Each contour is generated individually with its composing segments sorted in clockwise or counterclockwise order. Moreover, our algorithm generates information on how the contours are nested.
- (ii) We can preprocess \mathbf{M} , using $O(\text{SORT}(N))$ I/Os, into a linear-size data structure so that all contours at a given height can be reported using $O(\log_B N + T/B)$ I/Os, where T is the output size. Each contour is gen-

*The first and fourth authors are supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEG-P200A070505, and by a grant from the U.S.–Israel Binational Science Foundation. The second and third authors are supported in part by the US Army Research Office through grant W911NF-04-01-0278, by an Ole Roemer Scholarship from the Danish National Science Research Council, a NABIIT grant from the Danish Strategic Research Council, and by the Danish National Research Foundation, and in part by MADALGO: Center for Massive Data Algorithmics, a Center of the Danish National Research Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG’08, June 9–11, 2008, College Park, Maryland, USA.
Copyright 2008 ACM 978-1-60558-071-5/08/04 ...\$5.00.

erated individually with its composing segments sorted in clockwise or counterclockwise order.

Categories and Subject Descriptors

F.2.2 [Nonnumerical Algorithms and Problems]: Geometrical problems and computations

General Terms

Algorithms, Theory

Keywords

terrains, contours, geographical information systems, I/O-efficient algorithms

1. INTRODUCTION

Motivated by a wide range of applications, there is extensive work in many research communities on modeling, analyzing, and visualizing terrain data. A three-dimensional (digital elevation) model of a terrain is often represented as a two-dimensional uniform grid, with a height associated with every grid point, or a triangulated xy -monotone surface; the latter is also known as triangulated irregular network (TIN). A *contour* (or *isoline*) of a terrain \mathbf{M} is a connected component of a level set of \mathbf{M} . *Contour maps* (aka *topographic maps*), consisting of contour lines at regular height intervals, are widely used to visualize a terrain and to compute certain topographic information of a map; this representation goes back to at least the eighteenth century [19]. In this paper we propose efficient algorithms for computing contour maps and computing contours at a given level.

With the recent advances in mapping technologies, such as laser based LIDAR technology, hundreds of millions of points on a terrain, at sub-meter resolution with very high accuracy (~ 10 – 15 cm), can be acquired in a short period of time. The terrain models generated from these data sets are too large to fit in main memory and thus reside on disks. Transfer of data between disk and main memory is often the bottleneck in the efficiency of algorithms for processing these massive terrain models (see e.g. [12]). We are therefore interested in developing efficient algorithms in the two-level I/O-model [3]. In this model, the machine consists of a main memory of size M and an infinite-size disk. A block of B consecutive elements can be transferred between main memory and disk in one *I/O operation* (or simply *I/O*). Computation can only take place on elements in main memory, and the complexity of an algorithm is measured in terms of

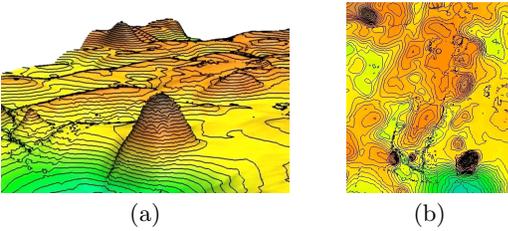


Figure 1: Examples of equidistant contour lines of a terrain. (a) Rendered on a perspective view of the terrain in 3d. (b) Projected onto the 2d plane.

the number of I/Os it performs. Over the last two decades, I/O-efficient algorithms and data structures have been developed for several fundamental problems, including sorting, graph problems, geometric problems, and terrain modeling and analysis problems. See the recent surveys [4, 23] for a comprehensive review of I/O-efficient algorithms. Here we mention that sorting N elements takes $\Theta\left(\frac{N}{B} \log_{M/B} \frac{N}{B}\right)$ I/Os, and we denote this quantity by $\text{SORT}(N)$.

Related work. A natural way of computing a contour K of a terrain \mathbf{M} is simply to start at one triangle of \mathbf{M} intersecting the contour and then tracing out K by walking through \mathbf{M} until we reach the starting point. If we have a starting point for each contour of a level set of \mathbf{M} , for a given level ℓ , we can compute all contours of that level set in linear time in the internal-memory model. The so-called *contour tree* [9] encodes a “seed” for each contour of K . Many efficient internal-memory algorithms are known for computing a contour tree; see e.g. [9]. Hence, one can efficiently construct a contour map of K . This approach of tracing a contour has been extended to higher dimensions as well, e.g., the well-known marching-cube algorithm for computing iso-surfaces [16].

An $O(\text{SORT}(N))$ algorithm in the I/O-model was recently proposed by Agarwal et al. [2] for constructing a contour tree of \mathbf{M} , so one can quickly compute a starting point for each contour. However, it is not clear how to trace a contour efficiently in the I/O-model, since a naive implementation requires $O(T)$ instead of $O(T/B)$ I/Os, to trace a contour of size T . Even using a provable optimal scheme for blocking a planar (bounded degree) graph, so that any path can be traversed I/O-efficiently [1, 17], one can only hope for an $O(T/\log_2 B)$ I/O solution. Nevertheless, I/O-efficient algorithms have been developed for computing contours on a terrain. Chiang and Silva [11] designed a linear-size data structure for storing a TIN terrain \mathbf{M} on disk such that all T edges in the contours at a query level ℓ can be reported in $O(\log_B N + T/B)$ I/Os, but their algorithm does not sort the edges along each contour. Agarwal et al. [1] designed a data structure with the same bounds so that each contour at level ℓ can be reported individually, with its edges sorted in either clockwise or counterclockwise order. However, while the space and query bounds of these structures are optimal, preprocessing them takes $O(N \log_B N)$ I/Os. This bound is more than a factor of B away from the desired $O(\text{SORT}(N))$ bound. Thus using this structure one can at best hope for an $O(N \log_B N + T/B)$ I/O algorithm to compute a contour map; here T is the total size of all the output contours. We refer the reader to the tutorial [18] and references therein for a review of practical algorithms for contour and iso-surface

extraction problems, and to [12, 15] for a sample of I/O-efficient algorithms for problems arising in terrain modeling and analysis.

Our results. Let \mathbf{M} be a terrain represented as a triangulated surface (TIN) with N vertices. For a contour K of \mathbf{M} , let $F(K)$ denote the set of triangles intersecting K . We prove (in Section 3) that there exists a total ordering \preceq on the triangles of \mathbf{M} that has the following two crucial properties:

- (C1) For any contour K , if we visit the triangles of $F(K)$ in \preceq order, we visit them along K in either clockwise or counter clockwise order.
- (C2) For any two contours K_1 and K_2 on the same level set of \mathbf{M} , $F(K_1)$ and $F(K_2)$ are not interleaved in \preceq ordering, i.e., suppose the first triangle of $F(K_1)$ in \preceq appears before that of $F(K_2)$, then either all of the triangles in $F(K_1)$ appear before $F(K_2)$ in \preceq , or all triangles of $F(K_2)$ appear between two consecutive triangles of $F(K_1)$ in \preceq .

We call such an ordering a *level-ordering* of the triangles of \mathbf{M} . We show that \preceq can be computed using $O(\text{SORT}(N))$ I/Os. Next, we present two algorithms that rely on this ordering.

Computing a contour map. Given as input a sorted list $\ell_1 < \dots < \ell_s$ of levels in \mathbb{R} , we present an algorithm (Section 4) that reports all contours of a \mathbf{M} at levels ℓ_1, \dots, ℓ_s using $O(\text{SORT}(N) + T/B)$ I/Os and $O(N/B)$ blocks of space, where T is the total number of edges in the output contours. Each contour is generated individually with its edges sorted in clockwise or counterclockwise order. Moreover, our algorithm reports how the contours are nested; see Section 4 for details.

Answering a contour query. We can preprocess \mathbf{M} , using $O(\text{SORT}(N))$ I/Os, into a linear-size data structure so that all contours at a given level can be reported using $O(\log_B N + T/B)$ I/Os, where T is the output size. Each contour is generated individually with its edges sorted in clockwise or counterclockwise order (Section 4.4).

2. PRELIMINARIES

Let $\mathbb{M} = (V, E, F)$ be a triangulation of \mathbb{R}^2 , with vertex, edge, and face (triangle) sets V , E , and F , respectively. We assume that V contains a vertex v_∞ , set at infinity, and that each edge $\{u, v_\infty\}$ is a ray emanating from u . The triangles in \mathbb{M} incident to v_∞ are unbounded. Let $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a continuous *height function* with the property that the restriction of h to each triangle of \mathbb{M} is a linear map. Given \mathbb{M} and h , the graph of h is a *terrain*. It is represented as an xy -monotone triangulated surface whose triangulation is induced by \mathbb{M} . That is, vertices, edges, and faces of \mathbf{M} are in one-to-one correspondence with those of \mathbb{M} and with a slight abuse of terminology we refer to V , E , and F , as vertices, edges, and triangles of both the terrain \mathbf{M} and the triangulation \mathbb{M} .

For convenience we assume that $h(u) \neq h(v)$ for all vertices $u \neq v$, and that $h(v_\infty) = -\infty$. Within each bounded triangle $f \in F$, h is uniquely determined as the linear interpolation of the height of the vertices of f . This is not

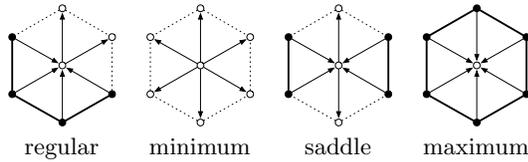


Figure 2: Link of a vertex; lower link is depicted by filled circles and bold edges. The type of a vertex is determined by its lower link.

the case for an unbounded face f since interpolation using $h(v_\infty) = -\infty$ is undefined; in which case to determine h on f an extra parameter, such as the height of a point in f , is needed.

For a given terrain \mathbf{M} and a level $\ell \in \mathbb{R}$, the ℓ -level set of \mathbf{M} , denoted by \mathbb{M}_ℓ , is defined as $h^{-1}(\ell) = \{x \in \mathbb{R}^2 \mid h(x) = \ell\}$. Equivalently, \mathbb{M}_ℓ is the vertical projection of $\mathbf{M} \cap z_\ell$ on the xy -plane, where z_ℓ is the horizontal plane $z = \ell$. The closed ℓ -sublevel and ℓ -superlevel sets of \mathbf{M} are defined respectively as $\mathbb{M}_{\leq \ell} = h^{-1}((-\infty, \ell])$ and $\mathbb{M}_{\geq \ell} = h^{-1}([\ell, +\infty))$, and the open ℓ -sublevel and ℓ -superlevel sets $\mathbb{M}_{< \ell}$ and $\mathbb{M}_{> \ell}$ are $\mathbb{M}_{\leq \ell} \setminus \mathbb{M}_\ell$ and $\mathbb{M}_{\geq \ell} \setminus \mathbb{M}_\ell$, respectively.

In much of what follows the height function of a terrain \mathbf{M} is only referred to for determining which endpoints of an edge of \mathbf{M} has greater height. We therefore “orient” each edge of \mathbf{M} toward its *higher* endpoint. Thus in the sequel, we treat \mathbf{M} as a directed triangulation and replace each undirected edge $\{u, v\} \in E$ with the ordered pair (u, v) , if $h(u) < h(v)$, or with (v, u) otherwise.

The dual graph $\mathbf{M}^* = (F^*, E^*, V^*)$ of the triangulation \mathbf{M} is defined as the planar graph that has a vertex $f^* \in F^*$ for each face $f \in F$, called the *dual* of f . There is an edge $e^* = (f_1^*, f_2^*) \in E^*$ if and only if the faces $f_1, f_2 \in F$ share an edge $e = (u, v)$ and f_1 is on the left and f_2 on the right side of the directed edge e . The graph \mathbf{M}^* is naturally embedded in the plane as follows: the vertex f^* is placed inside the face f and e^* is drawn as a curve that crosses e but no other edges of \mathbf{M} . A vertex $v \in V$ leads to a dual face v^* in \mathbf{M}^* that is bounded by the duals of the edges incident to v . The dual of \mathbf{M}^* is \mathbf{M} itself. For a given subset V_0 of V , we use the notation V_0^* to refer to the set of duals to the vertices in V_0 , i.e., $V_0^* = \{v^* : v \in V_0\}$. A similar notation is also used for sets of edges or faces.

Links and critical points. For a vertex v of \mathbf{M} , the *link* of v , denoted by $\text{Lk}(v)$, is the cycle in \mathbf{M} consisting of vertices adjacent to v , as joined by edges from triangles incident to v . The *lower link* of v , $\text{Lk}^-(v)$, is the subgraph of $\text{Lk}(v)$ induced by vertices lower (with smaller height) than v . The *upper link* of v , $\text{Lk}^+(v)$ is defined analogously; see Figure 2.

As we vary the height ℓ , the topology of $\mathbb{M}_{\leq \ell}$ changes at a discrete set $\{\ell_1, \dots, \ell_m\}$ of values, called the *critical levels* of h , where each ℓ_i is $h(v_i)$ for some vertex $v_i \in V$. The subset $\{v_1, \dots, v_m\}$ of V is the set of *critical vertices* of \mathbf{M} . A non-critical level of h is also called *regular*. Vertices with regular heights are *regular vertices*. By our assumption that the height of every vertex is distinct, there is only one critical vertex for each critical level.

There are three types of critical vertices: *minima*, *saddles*, and *maxima*. The type of a vertex v can be determined from the topology of $\text{Lk}^-(v)$: v is minimum, regular, saddle, or maximum if $\text{Lk}^-(v)$ is empty, a path, two or more paths, or a cycle, respectively. We assume that each saddle

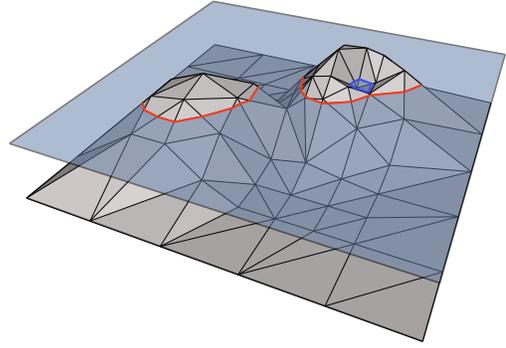


Figure 3: Red and blue contours in a level-set of a terrain.

is *simple* meaning that its lower link consists of only two paths. Multifold saddles can be split into simple saddles (e.g. [2]). Equivalently, a vertex can be classified based on the clockwise ordering of its incoming and outgoing edges: a minimum has no incoming edges, and a maximum has no outgoing edges. For other vertices v , we count the number of times incident edges switch between incoming to outgoing as we scan them around v in clockwise order. This number is always even. Two switches indicate a regular vertex, and four or more switches indicate a saddle.

A saddle vertex is further classified into two types. Let v be a saddle vertex, and let $\ell = h(v)$. The topology of $\mathbb{M}_{\leq \ell}$ differs from that of $\mathbb{M}_{< \ell}$ in one of two possible ways: either two connected components of $\mathbb{M}_{< \ell}$ join at v to become the same connected component in $\mathbb{M}_{\leq \ell}$, or the boundary of the same connected component of $\mathbb{M}_{< \ell}$ “pinches” at v introducing one more “hole” in $\mathbb{M}_{\leq \ell}$. Saddles of the former type are called *negative* and the latter type *positive* saddles. It is well-known that the number of minima (resp. maxima) is one more than the number of negative (resp. positive) saddles, and therefore

$$\#\text{saddles} = \#\text{minima} + \#\text{maxima} - 2. \quad (1)$$

This classification of saddles is related to persistent homology and a more general statement is proved in [13].

Contours. A *contour* of a terrain \mathbf{M} is a connected component of a level set of \mathbf{M} . Each contour K at a regular level is a simple cycle. K partitions $\mathbb{R}^2 \setminus K$ into two open sets — a bounded one called *inside* of K and denoted by K^i , and an unbounded one called *outside* of K and denoted by K^o . This is violated at critical levels at which a contour may shrink into a point (an extremum), or two contours may touch at a single point (a saddle). As the level ℓ changes from $-\infty$ to $+\infty$, contours in \mathbb{M}_ℓ change continuously except at critical levels where the contours containing the critical point undergo topological changes. Let K_1 and K_2 be two contours at levels ℓ_1 and ℓ_2 respectively with $\ell_1 < \ell_2$. We regard K_1 and K_2 as “the same” if K_1 continuously deforms into K_2 without any topological changes, as we sweep \mathbf{M} in the interval $[\ell_1, \ell_2]$.

Following [1], we call a contour K in \mathbb{M}_ℓ *blue* if, “locally”, $\mathbb{M}_{\leq \ell}$ lies in K^i , and *red* otherwise; see Figure 3. Every blue contour is born as a single point at a minimum. Conversely, a blue contour is born at every minimum except at v_∞ . Because of being placed at infinity, a red contour is born at v_∞ . Likewise, red contours “die” by shrinking into a single point at a maximum, and some red contour dies at every

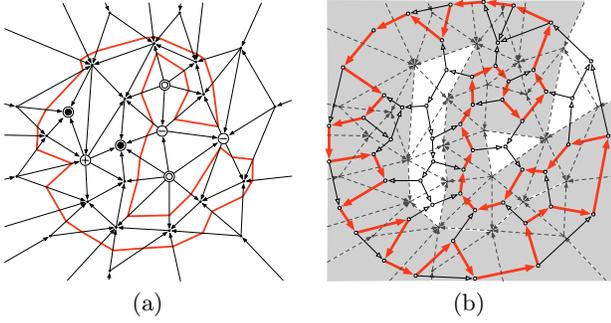


Figure 4: (a) Orientation of the edges in the plane triangulation \mathbb{M} of a terrain. Critical points and a contour K in a regular level set are shown. (b) the dual \mathbb{M}^* of \mathbb{M} . The representing cycle of K in \mathbb{M}^* is shown with bold edges. Triangles in $C(K, \mathbb{M})$ are shaded.

maximum. Two contours, with at least one of them being blue, merge into the same contour at a negative saddle. The resulting contour is red if one of the merging contours is red, and blue otherwise. A contour splits into two contours at a positive saddle. A red contour splits into two red contours while a blue contour splits into one red and one blue contour.

Two contours K_i and K_j of a level set \mathbb{M}_ℓ are called *neighbors* if no other contour K of \mathbb{M}_ℓ separates them, i.e., one of K_i and K_j is contained in K^i and the other in K^o . If K_i is neighbor to K_j and $K_i \subset K_j^i$, then K_i is called a *child* of K_j . If $K_i \subset K_j^o$ and $K_j \subset K_i^o$ then K_i is called a *siblings* of K_j . It can be verified that all children of a red (resp. blue) contour are blue (resp. red) contours while all siblings of a red (resp. blue) contour are red (resp. blue) contours.

We conclude this section by making a key observation, which is crucial for our main result. A contour of \mathbb{M} corresponds to a cycle in \mathbb{M}^* : let K be a contour in an arbitrary level set \mathbb{M}_ℓ , and let $F(K)$ (resp. $E(K)$) denote the set of faces (resp. edges) of \mathbb{M} that intersect K . If K is a red (resp. blue) contour, all the edges in $E(K)$ are oriented toward K^i (resp. K^o). Consequently, the vertices in $F^*(K)$ are linked by the edges in $E^*(K)$ into a cycle in \mathbb{M}^* . We refer to this cycle as *representing cycle of K* . We use $C(K, \mathbb{M})$ to denote the circular sequence of triangles dual to the representing cycle of K in \mathbb{M} . The sequence in $C(K, \mathbb{M})$ is oriented clockwise (resp. counterclockwise) if K is red (reps. blue).

3. LEVEL-ORDERING OF TRIANGLES

In this section we present our main result — the existence of a level-ordering of triangles of any terrain \mathbb{M} , i.e., an ordering that satisfies conditions (C1) and (C2). We begin by proving the existence of a level-ordering for a *simple terrains*, i.e., a terrain that does not have saddle vertices. Next we prove certain structural properties of a terrain and show that any arbitrary terrain can be transformed into a simple terrain by performing a surgery so that the contours of the original terrain are “preserved”. We then argue that a level-ordering on the transformed terrain corresponds to a level-ordering on the original one.

3.1 Simple terrain

Let \mathbb{M} be a simple terrain. The above discussion and (1) imply that \mathbb{M} has one (global) minimum, \tilde{v} , which coincides with v_∞ , and one (global) maximum, \hat{v} , and that every level

set consists of a single red contour that shrinks to a point at \hat{v} .

LEMMA 3.1. *Let $P \subset E$ be a directed (monotone) path in \mathbb{M} from \tilde{v} to \hat{v} . Then every cycle of \mathbb{M}^* contains exactly one edge from P^* . In particular, the graph $\mathbb{M}^* \setminus P^*$ obtained from deleting the edges in P^* from \mathbb{M}^* is acyclic.*

PROOF. We claim that \hat{v} is reachable in \mathbb{M} from every vertex $v \in V$. Recall that \hat{v} is the only local maximum in \mathbb{M} and that every other vertex has at least one outgoing edge. If one starts at v and follows an arbitrary outgoing edge at each step, the height of the new vertex is greater than the that of the previous one. This process can only stop at \hat{v} . By a similar argument, every vertex $v \in V$ is reachable from \tilde{v} .

Consider an arbitrary cycle C^* in \mathbb{M}^* . In the plane drawing of \mathbb{M}^* , C^* is a Jordan curve. Let $V_0 \subset V$ be the set of vertices that are contained in the inside of C^* (equivalently, $V_0^* \subset V^*$ is the set of faces of \mathbb{M}^* whose union is bounded by C^*). Let $C \subset E$ be the set of edges in \mathbb{M} dual to those in C^* . Since C^* is a cycle, the edges in C are either all oriented from V_0 to $V \setminus V_0$ or all from $V \setminus V_0$ to V_0 .

The former case cannot happen because $v_\infty \notin V_0$ and every vertex in V is reachable from v_∞ . If all edges of C are oriented from $V \setminus V_0$ to V_0 , then $\hat{v} \in V_0$ because otherwise \hat{v} cannot be reachable from the vertices of V_0 . Since $\hat{v} \in V_0$ and $v_\infty \in V \setminus V_0$, $|P \cap C| \geq 1$. There is no edge directed from V_0 to $V \setminus V_0$, so once P reaches a vertex of V_0 it cannot leave V_0 , implying that $|P \cap C| = 1$. Every cycle of \mathbb{M}^* is destroyed by the removal of the edges in P^* , implying that $\mathbb{M}^* \setminus P^*$ is acyclic. \square

Let P be the path from \tilde{v} to \hat{v} as defined in Lemma 3.1. The graph $\mathbb{M}^* \setminus P^*$ has all of the vertices of \mathbb{M}^* . Thus every face f of \mathbb{M} is represented by f^* in $\mathbb{M}^* \setminus P^*$. Let \prec be the a binary relation on F (triangles in \mathbb{M}) defined as $f_1 \prec f_2$ if $(f_1^*, f_2^*) \in E^* \setminus P^*$. Since by Lemma 3.1 $\mathbb{M}^* \setminus P^*$ is acyclic, \prec is a partial order on F . We call \prec the *adjacency partial order* induced by the acyclic graph $\mathbb{M}^* \setminus P^*$. A linear extension of \prec is any total order \preceq on F that is consistent with \prec , i.e. $f_1 \prec f_2$ implies $f_1 \preceq f_2$. Such a linear extension can be obtained by topological sorting of $\mathbb{M}^* \setminus P^*$. By definition, the existence of a directed path from f_i^* to f_j^* in $\mathbb{M}^* \setminus P^*$ implies that $f_i \preceq f_j$. This immediately results the following statement.

COROLLARY 3.2. *Let \mathbb{M} be a simple terrain, and let P be a directed (monotone) path from \tilde{v} to \hat{v} in \mathbb{M} . Let \preceq be a linear extension of the adjacency partial order induced by $\mathbb{M}^* \setminus P^*$. Then \preceq is a level-ordering of the triangles of \mathbb{M} .*

3.2 Red and blue cut-trees

Consider now a non-simple terrain with saddle vertices. Let S_\oplus and S_\ominus be the sets of positive and negative saddles of \mathbb{M} , respectively. We first introduce the notions of *ascending (red)* and *descending (blue) cut-trees* of \mathbb{M} as subgraphs of the triangulation \mathbb{M} , which we later use to turn \mathbb{M} into a simple terrain $\tilde{\mathbb{M}}$. Contours of each level set of \mathbb{M} will then be encoded in a corresponding level set of $\tilde{\mathbb{M}}$, each of which is a single contour.

A *descending* (resp. *ascending*) path on \mathbb{M} from a vertex $v \in V$ is a path v_0, v_1, \dots, v_r where $v_0 = v$ and $h(v_i) < h(v_{i-1})$ (resp. $h(v_i) > h(v_{i-1})$) for $i = 1, \dots, r$. For each

negative saddle v , let $P_1(v) = u_0, u_1, \dots, u_r$ and $P_2(v) = w_0, \dots, w_s$ be two descending paths from v such that u_r and w_s are both minima and u_1 and w_1 belong to different connected components of $\text{Lk}^-(v)$. Furthermore assume that for any two negative saddles u and w , if $P_i(u) = u_0, \dots, u_r$ and $P_j(w) = w_0, \dots, w_s$, for some $i, j \in \{1, 2\}$, and $u_k = w_l$ for some $1 \leq k \leq r$ and $1 \leq l \leq s$, then $u_{k+1} = w_{l+1}$, in other words, we assume descending paths from different vertices can join but then cannot diverge. Such a set of paths always exist: one can assign such paths to negative saddles in increasing order of their heights. At any negative saddle u , we follow a descending paths through each of the two connected components of $\text{Lk}^-(u)$ until it either reaches a minimum or joins a path already assigned to a lower negative saddle. Let $P(u) = P_1(u) \cup P_2(u)$ for any negative saddle u . Since $P_1(u) \setminus \{u\}$ and $P_2(u) \setminus \{u\}$ are contained in different connected components of $\mathbb{M}_{<h(u)}$, the underlying undirected graph of $P(u)$ is a simple path. For a positive saddle u , $P_1(u)$ and $P_2(u)$ are defined similarly using ascending paths that start at different connected components of $\text{Lk}^+(u)$ and end in maxima.

We define the *descending (blue) cut-tree* $\check{\mathbb{T}} = (\check{V}, \check{E})$ of \mathbb{M} to be $\cup_{u \in S_\ominus} P(u)$, and the *ascending (red) cut-tree* of $\hat{\mathbb{T}} = (\hat{V}, \hat{E})$ to be $\cup_{u \in S_\oplus} P(u)$. It is, of course, not clear that $\check{\mathbb{T}}$ and $\hat{\mathbb{T}}$ are trees but this and some of their other properties are proven below.

Remark. The definitions of red and blue cut-trees are closely related to notions of *split* and *join trees* in the context of *contour tree* [21]. In fact, if join and split trees are subgraphs of the terrain triangulation, they can replace the red and blue cut-trees. However, this is not always the case.

LEMMA 3.3. *The underlying undirected graph of a blue (resp. red) cut-tree $\check{\mathbb{T}}$ (resp. $\hat{\mathbb{T}}$) has no cycles.*

PROOF. We prove the claims for blue cut-tree. The argument for red cut-tree can be made symmetrically. Let u_1, \dots, u_r be the list of all negative saddles of \mathbb{M} in increasing order of height. Let $\check{\mathbb{T}}_0$ be the empty graph and for each $i = 1, \dots, r$, let $\check{\mathbb{T}}_i = \cup_{j=1}^i P(u_j)$; $\check{\mathbb{T}}_{i-1}$ is a subgraph of $\check{\mathbb{T}}_i$, and $\check{\mathbb{T}}_r = \check{\mathbb{T}}$. We prove by induction on i that each $\check{\mathbb{T}}_i$ is a forest. The empty graph $\check{\mathbb{T}}_0$ is trivially a forest. Assume now that $\check{\mathbb{T}}_i$ is a forest. By construction, adding $P(u_{i+1})$ connects two distinct connected components of $\check{\mathbb{T}}_i$, each contained in one of the two connected components of $\mathbb{M}_{<h(u_{i+1})}$ that join at u_{i+1} .

Moreover, once $P(u_{i+1})$ (or $P_2(u_{i+1})$) meets a vertex of $\check{\mathbb{T}}_i$, it follows a path of $\check{\mathbb{T}}_i$, therefore it does not create a cycle within a component of $\check{\mathbb{T}}_i$. \square

For a set $U \subseteq \mathbb{R}^2$, let $\check{\mathbb{T}}(U)$ (resp. $\hat{\mathbb{T}}(U)$) be the union of the paths $P(u)$ for all negative (resp. positive) saddles $u \in U$. In particular, $\check{\mathbb{T}} = \check{\mathbb{T}}(\mathbb{R}^2)$ and $\hat{\mathbb{T}} = \hat{\mathbb{T}}(\mathbb{R}^2)$.

LEMMA 3.4. *For a blue (resp. red) contour K , the underlying undirected graph $\check{\mathbb{T}}(K^i)$ (resp. $\hat{\mathbb{T}}(K^o)$) connects all of the minima in K^i (resp. K^o). A symmetric statement can be made for $\hat{\mathbb{T}}$ and maxima by switching “red” and “blue”.*

PROOF. We prove the lemma for $\hat{\mathbb{T}}$ and blue contours. The other cases are similar. Let K be a blue contour in \mathbb{M}_λ for some $\lambda \in \mathbb{R}$. We show that for each $\ell \in \mathbb{R}$, the minima in each connected component of $U_\ell = \mathbb{M}_{<\ell} \cap K^i$ are connected by $\hat{\mathbb{T}}(U_\ell)$. The statement of the lemma then



Figure 5: Cutting a triangulation along a tree.

follows by taking ℓ to be larger than the height of all vertices in K^i .

To prove the lemma we sweep ℓ from $-\infty$ toward $+\infty$ and verify the claim for U_ℓ . Every time ℓ reaches the height of a minimum in K^i , a new connected component is added to U_ℓ . The lemma holds for this new component since it originally has only a single minimum which is vacuously connected by $\check{\mathbb{T}}(U_\ell)$ to every other minimum in that component. The validity of the claim as ℓ continues to raise can only be altered when ℓ reaches the height of a negative saddle u in K^i at which two connected components U_1 and U_2 of $U_{<\ell}$, for $\ell = h(u)$, join at u . At this time the path $P(u)$ is added to $\hat{\mathbb{T}}(U_\ell)$. The crucial observation here is that because K is a blue contour, no descending path started at a vertex $u \in K^i$ can reach K^o . Thus the endpoints of $P(u)$ have to be minima in K^i . In other words $P(u)$, which reaches a minimum in U_1 and another in U_2 , connects $\hat{\mathbb{T}}(U_1)$ and $\hat{\mathbb{T}}(U_2)$ as desired. \square

COROLLARY 3.5. *The underlying undirected graphs of $\check{\mathbb{T}}$ and $\hat{\mathbb{T}}$ are trees. Moreover, all minima are vertices of $\check{\mathbb{T}}$ and all maxima are vertices of $\hat{\mathbb{T}}$.*

We conclude this discussion by mentioning a property of $\check{\mathbb{T}}$ and $\hat{\mathbb{T}}$ that follows from the construction.

LEMMA 3.6. *Let u be a vertex of $\hat{\mathbb{T}}$ (resp. $\check{\mathbb{T}}$). If u is a positive (resp. negative) saddle, then u has two outgoing (resp. incoming) edges in $\hat{\mathbb{T}}$ (resp. $\check{\mathbb{T}}$) — one to each connected component of the upper (resp. lower) link of u in \mathbb{M} . If u is a regular vertex or a negative saddle, then u has one outgoing (resp. incoming) edge. Finally, if u is a maximum (resp. minimum), then it has no outgoing (resp. incoming) edges.*

3.3 Surgery on terrain

Let $\hat{\mathbb{T}}$ be a red cut-tree for \mathbb{M} . Consider the following combinatorial operation on \mathbb{M} . First we duplicate every edge e of $\hat{\mathbb{T}}$, thus creating a face f_e that is bounded by the two copies of e . We then perform an Eulerian tour on the subgraph of \mathbb{M} induced by the copies of the edges of $\hat{\mathbb{T}}$ in which at each vertex the next edge of the tour is the first unvisited edge of the subgraph in clockwise order, relative to the previous edge of the tour. We then combine all of the faces f_e into a single face \hat{f} that is bounded by the Eulerian tour by making as many copy of each vertex as its degree in $\hat{\mathbb{T}}$ (or equivalently the number of times the tour has passed through it) and connecting non-tree edges incident on u to appropriate copies of u ; see Figure 5. Geometrically, the above modification of the terrain triangulation can be interpreted as “puncturing” the plane, thus creating a single hole in it, by cutting it along the edges of $\hat{\mathbb{T}}$. One can think of the created hole as the new face \hat{f} in this modified triangulation that is bounded by the $2|\hat{E}|$ edges in the Euler tour.

We then subdivide \hat{f} by placing a new vertex \hat{v} inside it and connecting \hat{v} via incoming edges (u, \hat{v}) to every vertex

u on the boundary of \hat{f} . The result is a triangulation $\mathbb{M}_0 = (V_0, E_0, F_0)$; see Figures 6 (a) and (b). The newly added triangles are all incident to \hat{v} , and we refer to them as \hat{v} -triangles. The edge e opposite to \hat{v} in a \hat{v} -triangle f (which is a copy of a \hat{T} edge) is called the *base* of f and f is said to be *based* at e . One can modify the plane drawing of \mathbb{M} into a (singular) plane drawing of \mathbb{M}_0 , that has faces of zero area and edges that bend overlap, by jamming all the new faces and edges in the (zero-area) hole that results from cutting the plane through \hat{T} .

\mathbb{M}_0 can be regarded as the triangulation of a terrain \mathbf{M}_0 : Fáry's theorem [14] can be used to straight-line embed \mathbb{M}_0 while preserving all its faces and the height function of \mathbf{M} induces a height function on triangles of \mathbb{M}_0 that are also in \mathbb{M} . The height of \hat{v} is then set higher than all vertices of \mathbf{M} and is used to linearly interpolate a height function on \hat{v} -triangles.

LEMMA 3.7. \mathbb{M}_0 has no positive saddles and exactly one maximum, namely \hat{v} . The minima of \mathbb{M}_0 are precisely those of \mathbf{M} . Each negative saddle of \mathbb{M}_0 is a copy of a negative saddle of \mathbb{M} , and only one copy of each negative saddle of \mathbb{M} is a negative saddle of \mathbb{M}_0 .

PROOF. For a vertex $u \notin \hat{T}$, $\text{Lk}(u)$ in \mathbb{M} and \mathbb{M}_0 is the same modulo taking copies of \hat{T} vertices as identical. In particular, minima of \mathbb{M} stay minima in \mathbb{M}_0 . Thus it suffices to consider \hat{v} and copies of \hat{T} vertices. Clearly, \hat{v} is a maximum. Let u be a vertex of \hat{T} , and let u' be a copy of u in \mathbb{M}_0 . Let e'_1 and e'_2 be copies of \hat{T} edges that enter and leave u' , respectively, in the Eulerian tour of \hat{T} . Both of these edges remain incident on u' in \mathbb{M}_0 . Let v'_1 and v'_2 be other endpoints of e'_1 and e'_2 , respectively in \mathbb{M}_0 . Let v_i and e_i , $i = 1, 2$, be the vertex and edge in \mathbb{M} corresponding to v'_i and e'_i , respectively. $\text{Lk}(u')$ consists of a path $\pi(u')$ from v'_1 to v'_2 followed by \hat{v} . Moreover, e'_1 and e'_2 are the only edges incident on u' that are copies of \hat{T} edges, and $\pi(u')$ is also a path in $\text{Lk}(u)$ in \mathbb{M} , modulo taking copies of \hat{T} vertices as identical.

First, u' cannot be a maximum because u' is adjacent to \hat{v} . It cannot be a minimum either because then $\pi(u') \subseteq \text{Lk}^+(u)$ and e_1 and e_2 are outgoing edges from u in \hat{T} connected to some component of $\text{Lk}^+(u)$, which contradicts Lemma 3.6. Next, if $\text{Lk}^+(u')$ is not connected, then its component U not containing \hat{v} does not contain v'_1 and v'_2 either and thus u lies in the interior of the path $\pi(u')$. Then U is also a connected component of $\text{Lk}^+(u)$ in \mathbb{M} . Unless u is a negative saddle, by Lemma 3.6, there is an outgoing edge in \hat{T} from u to a vertex in U , contradicting the fact that e'_1 and e'_2 are the only edges adjacent to u' that are copies of \hat{T} edges. Hence, unless u is a negative saddle, $\text{Lk}^+(u')$ is connected and u' is a regular vertex in \mathbb{M}_0 .

Finally, suppose u is a negative saddle, with two components U_1 and U_2 in $\text{Lk}^+(u)$. By Lemma 3.6, u has exactly one outgoing edge e in \hat{T} . Without loss of generality assume that e is connected to U_1 . Then U_2 will appear as a connected component of the upper link of exactly one copy u' of u , namely if $U \subseteq \pi(u')$, and u' will be a negative saddle in \mathbb{M}_0 . The upper link of all other copies of u will be connected — consisting of \hat{v} and possibly a portion of U_1 . Consequently, one copy of every negative saddle in \mathbb{M} becomes a negative saddle in \mathbb{M}_0 and other copies become regular vertices. This completes the proof of the lemma. \square

Next we perform a similar surgery on \mathbb{M}_0 only using a red cut-tree \hat{T} of \mathbb{M}_0 . As above, the idea is to slice the plane at \hat{T} and insert a new vertex \check{v} in the resulting face and connect \check{v} to every copy u of a vertex in \hat{T} by an outgoing edge (\check{v}, u) . We call the resulting triangulation $\tilde{\mathbb{M}} = (\tilde{V}, \tilde{E}, \tilde{F})$. A slight technicality arises in this case as a result of the fact that v_∞ is a minimum of \mathbb{M}_0 which by Corollary 3.5 is a vertex of \hat{T} . We omit the details from this version and conclude using the same argument as in Lemma 3.7, the following:

LEMMA 3.8. $\tilde{\mathbb{M}}$ does not have saddle vertices.

LEMMA 3.9. If (f_1^*, f_2^*) is an edge of \mathbb{M}^* , then there is a path from f_1^* to f_2^* in $\tilde{\mathbb{M}}^*$.

PROOF. If f_1 and f_2 are adjacent in $\tilde{\mathbb{M}}$ then (f_1^*, f_2^*) is an edge in $\tilde{\mathbb{M}}^*$. Thus we only need to consider the case in which and edge e shared by f_1 and f_2 in \mathbb{M} is an edge of \hat{T} or \tilde{T} (or both). Suppose e is an edge of \hat{T} . In constructing \mathbb{M}_0 , e is duplicated to create two edges e_1 and e_2 , respectively, incident to f_1 and f_2 . Let ϕ_1 and ϕ_2 respectively be the \hat{v} -triangles based at e_1 and e_2 . By construction, f_1 is to the left and ϕ_1 to the right of e_1 and therefore (f_1^*, ϕ_1^*) is an edge in \mathbb{M}^* . Similarly (ϕ_2^*, f_2^*) are edges in \mathbb{M}^* . Consider the subgraph of \mathbb{M}^* induced by \hat{v} -triangles. Since all the edges incident to \hat{v} are incoming, their duals make a cycle in \mathbb{M}^* which includes ϕ_1^* and ϕ_2^* . Since there is a path from ϕ_1^* to ϕ_2^* on this cycle and there are edges from f_1^* to θ_1^* and from θ_2^* to f_2^* in \mathbb{M}^* , we get a path from f_1^* to f_2^* . It is easy to observe that the same argument extends to neighboring $\tilde{\mathbb{M}}$ triangles that are separated by the edges of \tilde{T} or both \tilde{T} and \hat{T} . \square

3.4 Encoding of contours in simplified terrain

Although we argued above that $\tilde{\mathbb{M}}$ can be realized as the triangulation of a terrain $\tilde{\mathbf{M}}$, in what follows we use a degenerate realization of it that is different from what a straight-line embedding of $\tilde{\mathbb{M}}$ results and substantially simplifies the arguments of the rest of this section. Given a terrain \mathbf{M} , we surgically modify the surface it represents in \mathbb{R}^3 as follows. The graph of the restriction of the height function h of \mathbf{M} to the red cut-tree \hat{T} of \mathbf{M} is a tree \hat{T} embedded on \mathbf{M} (and therefore in \mathbb{R}^3). We cut \mathbf{M} along \hat{T} and represent each edge (u, \hat{v}) of $\tilde{\mathbf{M}}$ (where u is a copy of a \hat{T} vertex) as a vertical ray in \mathbb{R}^3 (parallel to the z -axis) that emanates from the vertex of \hat{T} corresponding to u in the positive direction of the z -axis. In other words, we take \hat{v} to be at infinity in the positive z -direction. A \hat{v} -triangle based at an edge (u, w) (where (u, w) is a copy of a \hat{T} edge) is then represented by a vertical wall erected on top of the edge corresponding to e in \mathbf{M} that extends to $+\infty$ and is bounded at its vertical sides by the vertical rays that represents (u, \hat{v}) and (w, \hat{v}) ; see Figure 7. We then cut the resulting surface along a blue cut-tree of it and carry out a similar construction for the edges and triangles incident to \check{v} , only this time the rays and walls extend toward $-\infty$. We call the final surface $\tilde{\mathbf{M}}$. Note that although $\tilde{\mathbf{M}}$ is topologically a triangulated surface whose triangulation has the same combinatorial structure as $\tilde{\mathbb{M}}$ (modulo taking all vertical edges and triangles that go to $+\infty$ or $-\infty$ as being incident to \hat{v} or \check{v} respectively), after cutting $\tilde{\mathbf{M}}$ along the cut-trees all the copies of each vertex of the tree still reside at the same point in \mathbb{R}^3 . Thus the vertical rays and walls that represent edges or triangles based at

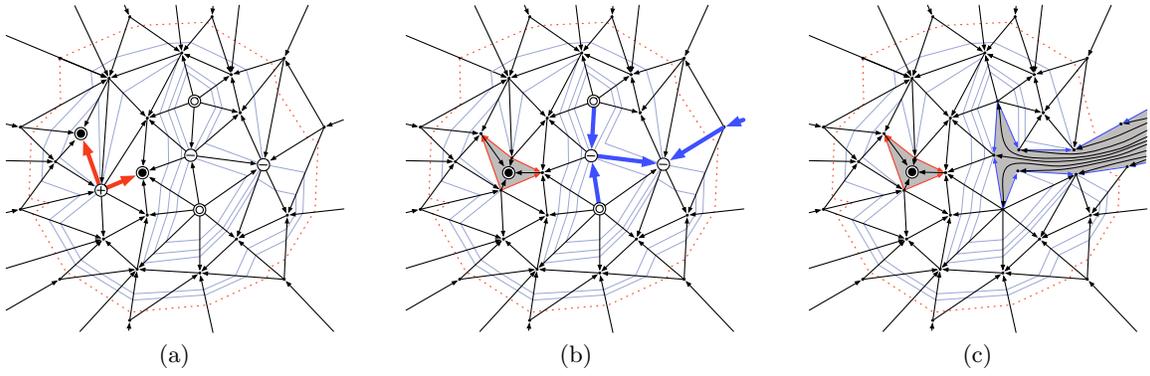


Figure 6: (a) A red (ascending) cut-tree marked \hat{T} marked on the terrain \mathbb{M} of Figure 4. (b) Construction of the graph \mathbb{M}_0 : the terrain is cut along \hat{T} and a new maximum \hat{v} is inserted in the opened face. On the right, a blue cut-tree of \mathbb{M}_0 is marked. (c) Construction of the graph $\tilde{\mathbb{M}}$: the terrain is cut open on the red cut-tree and a new maximum is inserted.

different copies of the same cut-tree vertices or edges overlap.

The surface $\tilde{\mathbb{M}}$ is not xy -monotone and therefore violates our definition of a terrain. All points on rays and vertical triangles based at copies of \hat{T} vertices and edges vertically project into \hat{T} and a similar statement holds for \check{T} . It is nonetheless possible to define the ℓ -level sets $\tilde{\mathbb{M}}_\ell$ of $\tilde{\mathbb{M}}$ in the same way as they are defined for xy -monotone surfaces as the vertical projection into xy -lane of $\tilde{\mathbb{M}} \cap z_\ell$. On the other hand, all the triangles of \mathbb{M} are also in $\tilde{\mathbb{M}}$ and correspond to the same triangles of \mathbb{M} and $\tilde{\mathbb{M}}$ in \mathbb{R}^3 . This implies that $\mathbb{M}_\ell \subseteq \tilde{\mathbb{M}}_\ell$ for all $\ell \in \mathbb{R}$ and $\tilde{\mathbb{M}}_\ell \setminus \mathbb{M}_\ell \subset \hat{T} \cup \check{T}$. In other words $\tilde{\mathbb{M}}_\ell$ consists of the contours of \mathbb{M}_ℓ connected to each other by pieces of red and blue trees. Let $\hat{T}_\ell = \hat{T} \cap \tilde{\mathbb{M}}_\ell$ and $\check{T}_\ell = \check{T} \cap \tilde{\mathbb{M}}_\ell$.

LEMMA 3.10. *Let K_0 be a blue (resp. red) contour in a level set \mathbb{M}_ℓ and let K_1, \dots, K_r be its children. Then K_0, K_1, \dots, K_r are connected to each other in $\tilde{\mathbb{M}}_\ell$ only through paths in \hat{T}_ℓ (resp. \check{T}_ℓ). Moreover, if G_0 is defined as a graph that results from the restriction of \hat{T} (resp. \check{T}) to $K_0^i \setminus (K_1^i \cup \dots \cup K_r^i)$ by contracting each contour K_i into a vertex, $i = 0, \dots, r$, then G_0 is a tree. The statement of the lemma is also valid for the outermost contours that do not have a parent.*

PROOF. We prove the lemma for the case where K_0 is blue. The proof for the case where it is red is symmetric. For the first claim, observe that the subset S of the plane defined as $K_0^i \setminus (K_1^i \cup \dots \cup K_r^i)$ is contained in the sublevel set $\mathbb{M}_{<\ell}$. Therefore vertical triangles that intersect S must connect edges with at least one endpoint in $\tilde{\mathbb{M}}_{<\ell}$ to \hat{v} and are therefore \hat{v} -triangles.

For the second part we prove a cycle in G_0 , implies a cycle in the underlying graph of \hat{T} and this contradicts Corollary 3.5. Consider any contour K_j with $j = 1, \dots, r$ and let e_1 and e_2 be two edges of \hat{T} that intersect K_j . Since K_j is red, all edges crossing it have their higher endpoint in K_j^i . Since e_1 is an edge of \hat{T} , its higher endpoint is followed by an ascending path that ends in a maximum. Since no ascending path can leave K_j^i , \hat{T} reaches a maximum v_1 in K_j^i through e_1 . Similarly, \hat{T} reaches a maximum v_2 in K_j^i through e_2 . On the other hand Lemma 3.4 implies that v_1 and v_2 are connected by a path in \hat{T} contained in K_j^i . In other words, any two branches of \hat{T} that enter K_j^i meet in K_j^i . A similar

argument shows that any two branches of \hat{T} that enter K_0^i meet in K_0^i . Thus, a cycle in G_0 implies a cycle in the underlying undirected graph of \hat{T} . \square

The nested structure of red and blue contours together with Lemma 3.10 result the following statement.

COROLLARY 3.11. *Let K_1, \dots, K_ℓ be all of the contours in \mathbb{M}_ℓ and let G be the graph that results from $\hat{T}_\ell \cup \check{T}_\ell$ by contracting every contour K_i into a vertex. Then G is a tree.*

Since $\tilde{\mathbb{M}}$ does not have saddles, $\tilde{\mathbb{M}}^* \setminus P^*$ for a \hat{v} - \hat{v} path P is acyclic, by Lemma 3.1. Let \prec be adjacency partial order on \tilde{F} induced by $\tilde{\mathbb{M}}^* \setminus P^*$. Since $F \subset \tilde{F}$, \prec is also a partial order on F .

LEMMA 3.12. *Let \preceq be linear extension of \prec on F . If K and K' are two contours of a level set \mathbb{M}_ℓ and $f_1, f_2 \in F(K)$ and $f'_1, f'_2 \in F(K')$ are such that $f_1 \prec f'_1 \preceq f_2$, then $f_1 \preceq f'_2 \preceq f_2$.*

PROOF. Since $\tilde{\mathbb{M}}$ is a simple terrain, $\tilde{\mathbb{M}}_\ell$ consists of a single contour. Let $C^* = C(K, \tilde{\mathbb{M}}^*)$ be the representing cycle of $\tilde{\mathbb{M}}_\ell$ in $\tilde{\mathbb{M}}^*$. If $f_1, f_2 \in F(K)$ for some contour K in \mathbb{M}_ℓ and the edge common to f_1 and f_2 does not belong to either of \hat{T} or \check{T} , then (f_1^*, f_2^*) is an edge in C^* . By Lemma 3.1, C^* has exactly one edge in P^* . Thus $C^* \setminus P^*$ is a path Q^* that is exactly one edge short of C^* .

If a \hat{v} -triangle based at one copy of a \hat{T} edge e intersects the plane z_ℓ in some segment, the \hat{v} -triangle based at the other copy of e does so as well in the same segment. Thus all the segments in \hat{T}_ℓ or \check{T}_ℓ correspond to two overlapping segments in the contour $\tilde{\mathbb{M}}_\ell$. In other words $\tilde{\mathbb{M}}_\ell$ is a closed curve that connects contiguous fragments of individual contours that are not crossed by the red or blue cut-trees with paths from \hat{T}_ℓ and \check{T}_ℓ in which $\tilde{\mathbb{M}}_\ell$ overlaps itself “twofold”; see Figure 8. Let G be the tree of Corollary 3.11. Each contour K_i appears as a vertex in G . Let $K = K_0, \dots, K_r = K'$ be the sequence of vertices (contracted contours) visited in the unique path in G that connects K to K' . By Lemma 3.10 for each $i = 1, \dots, r$, K_{i-1} is connected to K_i by a unique path in either \hat{T}_ℓ or \check{T}_ℓ . Let s be an arbitrary segment from either \hat{T}_ℓ or \check{T}_ℓ on the unique path connecting K_0 to K_1 .

For $f_1 \preceq f'_1 \preceq f_2$ to hold, Q^* must visit $f_1^*, f_1'^*$ and f_2^* in this order. Assume without loss of generality that

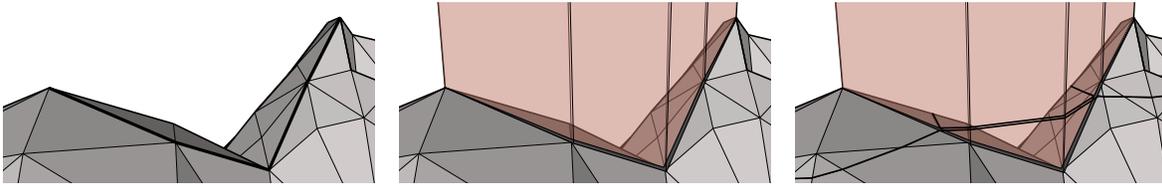


Figure 7: Left: The red cut-tree \hat{T} of a terrain. Middle: The terrain is cut along \hat{T} and \hat{v} -triangles are represented by vertical walls. Right: A contour of the simplified terrain overlapping itself on \hat{v} -triangles.

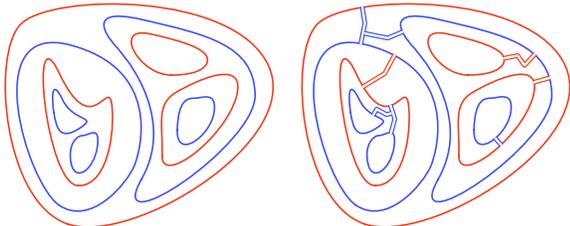


Figure 8: Contours of M_ℓ (left) versus those of \tilde{M}_ℓ (right).

$f'_1 \preceq f'_2$. In order for $f_1 \preceq f'_2 \preceq f_2$ not to hold, one must have $f_2 \preceq f'_2$ which means Q^* must visit f'_2 after f_2 . But this corresponds to going from K to K' , then back to K and then again to K' . This corresponds to \tilde{M}_ℓ passing through s three times, a contradiction. \square

Lemmas 3.9 and 3.12 respectively prove that the total order \preceq has properties (C1) and (C2) of a level-ordering.

THEOREM 3.13. *For any terrain M with triangulation \mathbb{M} , there is exists a level level-ordering of the triangles of M .*

4. CONTOUR ALGORITHMS

In this section we describe I/O-efficient algorithms for computing contour maps as well as an I/O-efficient data structure for answering contour queries.

4.1 Level-ordering of terrain triangles

We describe an I/O-efficient algorithm for computing the triangulation \tilde{M} of the simplified terrain \tilde{M} , and a monotone path P from \tilde{v} to \hat{v} in \tilde{M} . We can then compute a level-ordering of triangles in \tilde{M} , which also gives a level-ordering of triangles in M , by doing a topological sort on the vertices of $\tilde{M}^* \setminus P^*$ in $O(\text{SORT}(N))$ I/Os.

Computing the red cut-tree. The first in computing \tilde{M} is to compute a red (ascending) cut-tree \hat{T} of M . The I/O-efficient topological persistence algorithm of Agarwal et al. [2] can determine the type of every vertex of M in $O(\text{SORT}(N))$ I/Os. Moreover, for every vertex $v \in M$, it can also compute, within the same I/O bound, a vertex from each connected component of $\text{Lk}^+(v)$. Since each saddle of M is assumed to be simple, $\text{Lk}^+(v)$ has at most two connected components.

\hat{T} is computed using the *time-forward processing* technique [10] and using a priority queue Q : we scan the vertices of M in the increasing order of their heights. We store a subset of vertices in Q , namely the upper endpoints of the edges of \hat{T} whose lower endpoints have been scanned. The priority of a vertex v in Q is its height $h(v)$. Suppose we are scanning a vertex v of M and u is the lowest priority vertex in

Q . If $h(v) < h(u)$ and v is not a positive saddle, we move to a new vertex in M . Otherwise, i.e. if $h(u) = h(v)$ or v is a positive saddle, we choose a vertex w from each connected component of $\text{Lk}^+(v)$, which we have already computed in the preprocessing step. We add the edge (v, w) to \hat{T} and add w to Q . Since each operation on Q can be performed in $O\left(\frac{1}{B} \log_{M/B} N/B\right)$ I/Os, \hat{T} can be computed in $O(\text{SORT}(N))$ I/Os.

Adding the blue cut-tree. The second step in computing \tilde{M} is to compute a blue cut-tree \tilde{T} of M_0 . However, we can compute \tilde{T} directly on M if we ensure that \hat{T} and \tilde{T} do not cross each other, even though they can share edges. This property can be ensured by choosing the ascending and descending edges, in \hat{T} and \tilde{T} , respectively, out of each vertex v , more carefully. Specifically, we use the following rule:

1. On an ascending path, the edge following (u, v) is (v, w) where (v, w) is the first outgoing edge out of v in clockwise order from (u, v) , and
2. On a descending path, the edge following (v, u) is (w, v) where (w, v) is the first incoming edge of v in counter-clockwise order from (v, u) ,

It can be verified that \hat{T} and \tilde{T} do not cross. One can therefore compute \tilde{T} precisely in the same way as \hat{T} directly on M .

Computing a monotone \tilde{v} - \hat{v} path P . While computing \tilde{T} we also compute a descending path starting at the lowest *positive* saddle v_1 of M as though v_1 were another negative saddle. This path P , which ends at a \tilde{T} vertex v_0 , together with (\tilde{v}, v_0) and (v_1, \hat{v}) serves as a monotone path in M connecting \tilde{v} to \hat{v} .

Generating $\tilde{M}^* \setminus P^*$. The topological sorting algorithm by Arge et al. [8] takes as input a planar directed acyclic graph, represented as a list of vertices upon with the list of edges incident upon them in circular order. Given M , \hat{T} , \tilde{T} , and P , we need to compute such a representation of $\tilde{M}^* \setminus P^*$. Since each face in \tilde{M} is a triangle, the degree of each vertex in \tilde{M}^* is three. It is easy to compute the circular order of edges incident upon a vertex of \tilde{M}^* whose dual triangle is neither a \hat{v} - or \tilde{v} -triangle, nor adjacent to a copy of a \hat{T} or \tilde{T} edge. The main task is then to compute the \hat{v} - and \tilde{v} -triangles. This can be accomplished by computing the Eulerian tours of \hat{T} and \tilde{T} , which takes $O(\text{SORT}(N))$ I/Os [8]. Putting everything together, we obtain the main result of this paper.

THEOREM 4.1. *Given a terrain M with triangulation M , a level-ordering of the triangles of M can be computed in $O(\text{SORT}(N))$ I/Os, where N is the number of vertices of M .*

4.2 Contour maps of simple terrains

Let $L = \{\ell_1, \dots, \ell_s\}$ be a set of input levels with $\ell_1 < \dots < \ell_s$. Given a simple terrain \mathbb{M} , the goal is to compute the contour map of \mathbb{M} for levels in L . Since \mathbb{M} is simple, each \mathbb{M}_{ℓ_i} is a single contour. Generating the segments of \mathbb{M}_{ℓ_i} in clockwise or counterclockwise order is equivalent to listing the triangles of \mathbb{M} , the contour \mathbb{M}_{ℓ_i} intersects in that order, i.e. reporting $C(\mathbb{M}_{\ell_i}, \mathbb{M})$.

Our algorithm uses a *buffer tree* \mathcal{B} to store the terrain triangles of \mathbb{M} that intersect a level set. The buffer tree [5] is a variant of a B-tree, which propagates updates from the root to the leaves in a lazy manner, using buffers attached to the internal nodes of the tree. As a result, a sequence of N updates (inserts and deletes) can be performed in amortized $O(\text{SORT}(N))$ I/Os. Moreover, one can perform a *flush* operation on a buffer tree that results in the writing of all its stored elements on the disk in sorted order. Flushing a tree with T elements takes $O(T/B)$ I/Os. After the flushing, only those triangles that intersect the sweeping plane remain in \mathcal{B} .

It is more intuitive to describe the algorithm as a plane sweep of \mathbb{M} in \mathbb{R}^3 . At the first step, the algorithm computes a level-ordering of the terrain triangles using Corollary 3.2. Then starting at $\ell = -\infty$, the algorithm sweeps a horizontal plane z_ℓ at height ℓ in the positive z -direction. A *target level* ℓ_* is initially set to ℓ_1 . At any time the algorithm maintains a list of triangles in \mathbb{M} that intersect the sweeping plane in a buffer tree \mathcal{B} ordered by \prec . Whenever the sweep plane encounters the bottom-most vertex of a triangle f of \mathbb{M} , we insert f into \mathcal{B} ; f is deleted again from \mathcal{B} when the plane reaches the top-most vertex of f . When the sweep plane z_ℓ reaches the height of the target level ℓ_* , it flushes the buffer tree. The generated list of triangles (vertices of \mathbb{M}^*) are precisely the set of triangles in \mathbb{M} that intersect the horizontal plane $z = \ell_*$, ordered by \prec . Corollary 3.2 implies that the output is exactly $C(\mathbb{M}_{\ell_*}, \mathbb{M})$. The algorithm then raises the target level ℓ_* to the next level in L and continues.

Level-ordering the terrain triangles takes $O(\text{SORT}(N))$ I/Os (Theorem 4.1). Preprocessing for the sweep algorithm consists of sorting the vertices in their increasing order of heights which can also be done in $O(\text{SORT}(N))$ I/Os. During the sweep each update on the buffer tree takes $O(\frac{1}{B} \log_{M/B} |N|)$ amortized I/Os [5]. Thus all the $O(N)$ updates can be performed in $O(\text{SORT}(N))$ I/Os in total. Each flushing operation takes $O(1 + T'/B)$ I/Os, where T' is the number of triangles in \mathcal{B} . If a triangle is in \mathcal{B} but has been deleted, it is not in \mathcal{B} after the flushing operation, so a “spurious” triangle is flushed only once.

Hence, the total number of I/Os is $O(\text{SORT}(N) + T/B)$, where T is the output size. Finally, in addition to storage used for the terrain the algorithm uses $O(N/B)$ blocks to store the buffer tree and thus uses $O(N/B)$ blocks in total.

4.3 Generalization to arbitrary terrains

Given a general terrain \mathbb{M} with saddles, one can still compute by Theorem 4.1 a level-ordering of the triangles of \mathbb{M} in $O(\text{SORT}(N))$ I/Os. If one runs the algorithm of the previous section on $\tilde{\mathbb{M}}$ the output generated for each input level ℓ_i is $C(\tilde{\mathbb{M}}_{\ell_i}, \tilde{\mathbb{M}})$. Running the algorithm on \mathbb{M} is equivalent to running it on $\tilde{\mathbb{M}}$ but ignoring all \hat{v} and \check{v} -triangles. Consequently, the produced output for level ℓ_i is the same sequence of triangles only with \hat{v} and \check{v} -triangles omitted. By Theorem 3.13 this is a subsequence $R = \langle f_1, \dots, f_k \rangle$ of

the $C(\tilde{\mathbb{M}}_{\ell_i}, \tilde{\mathbb{M}})$ in which $C(K, \mathbb{M})$ of each contour K in \mathbb{M}_{ℓ_i} appears as a subsequence R_K . Thus all one needs to do is to extract the subsequence R_K and write it separately in the same order as it appears in R . Property (C2) of a level-ordering allows this to be done in $O(k/B)$ I/Os: if in R some elements of R_K are later followed by elements of $R_{K'}$, then the appearance of another element of R_K , indicates that no more elements from $R_{K'}$ remain.

We scan the sequence R in order and push the scanned triangles into a stack S_F . Every time the last element of a contour is pushed into the stack, the triangles of that contour make a suffix of the list of elements stored in S_F . At such a point, we pop all the elements corresponding to the completed contour and write them to disk. To recognize when a contour is completed and how many elements on the top of stack belong to it, we keep a second stack S_E of edges. For any triangle $f \in F(\mathbb{M}_{\ell_i})$, two of the edges of f intersect \mathbb{M}_{ℓ_i} . With respect to the orientation of these edges, f is to the right of one of them and to the left of the other one which we respectively call the *left* and *right* edges of f at level ℓ_i . If $e^* = (f_j^*, f_{j+1}^*)$ is an edge of the representing cycle of a contour in \mathbb{M}_{ℓ_i} , then e is the right edge of f_j and left edge f_{j+1} at level ℓ_i . We therefore check when scanning a triangle f_j whether its left edge is the same as the right edge of the triangle on top of S_F and insert f_j into S_F if this is the case. Otherwise, we compare the left edge of f_j with the edge on top of S_E . If they are not the same, we are visiting a new contour and we insert the left edge of f_j into S_E and f_j into S_F . Otherwise, f_j is the last triangle of its contour. Therefore we write it to the disk and successively pop and write to disk enough triangles from S_F until the left edge of a popped triangle is the same as the right edge of f_j . We also pop this edge from S_E . In this algorithm each scanned triangle is pushed to the stack once and popped another time. In a standard I/O-efficient stack implementation this costs $O(k/B)$ I/Os.

THEOREM 4.2. *Given any terrain \mathbb{M} with N vertices and a list $L = \{\ell_1, \dots, \ell_s\}$ of real levels with $\ell_1 < \dots < \ell_s$, one can compute using $O(\text{SORT}(N) + T/B)$ I/Os the contour map of \mathbb{M} for levels in L , where T is the total number of produced segments.*

Remark 4.3. In addition to reporting each contour individually, a number of applications call for computing how various contours are nested in each other. Let \mathbb{K} be the set of contours computed by the algorithm. We define a tree \mathcal{N} on \mathbb{K} as follows: Recall that a *contour tree* of \mathbb{M} is a tree embedded in \mathbb{R}^3 , each point of the tree is identified with one contour of \mathbb{M} . By marking the points in the contour tree that correspond to the contours in \mathbb{K} as the vertices and contracting some of the edges of the contour tree, we can construct \mathcal{N} . Using the I/O-efficient algorithm by Agarwal et al. [2] for constructing contour trees, we can construct \mathcal{N} in $O(\text{SORT}(N))$ I/Os. We omit the details from this version.

4.4 Answering contour queries

The sweep algorithm described in the previous section can easily be modified to construct a linear space data structure that given a query level ℓ can report the contours in the level set \mathbb{M}_ℓ I/O-efficiently. Unlike the previously known structure for this problem [1], our structure can be constructed in

$O(\text{SORT}(N))$ I/Os. To obtain the structure we simply replace the buffer tree \mathcal{B} with a partially persistent B -tree [6, 22]. To build the structure, we sweep \mathbf{M} by a horizontal plane in the same way as we did in the algorithm of Section 4.2, inserting the triangles when the sweep plane reaches their bottom-most vertex without checking for them to intersect any target levels and deleting them when the sweep plane passes their top-most vertex. There will also be no need to flush the tree.

Since $O(N)$ updates can be performed on a persistent B -tree in $\text{SORT}(N)$ I/Os [20, 7], the sweeping of the terrain require $O(\text{SORT}(N))$ I/Os. Since a persistent B -tree allows us to query any previous version of the structure and in particular produce the list of the elements stored in the tree in $O(\log_B N + T/B)$ I/Os if T is the number of reported elements, we can now obtain \mathbb{M}_ℓ in the same bound, simply by querying the structure for the triangles it contained when the sweep-plane was at height ℓ and then utilize Theorem 3.13 and the contour extraction algorithm discussed above to extract individual contours of \mathbb{M}_ℓ .

THEOREM 4.4. *Given a terrain \mathbf{M} with N vertices, one can construct in $O(\text{SORT}(N))$ I/Os a linear size data structure, such that given a query level ℓ , one can report contours of \mathbb{M}_ℓ in $O(\log_B(N) + T/B)$ I/Os where T is the size of the query output. Each contour is reported individually, and the edges of each contour are sorted in clockwise order.*

5. CONCLUSIONS

We defined level-ordering of terrain triangles and proved that every terrain has a level ordering that can be computed I/O-efficiently. Based on this, we provided algorithms that compute contours of a given terrain within similar I/O bounds. An immediate question is whether this approach can be generalized to general triangulated surfaces and arbitrary piecewise functions defined on them. For such surfaces, one can in particular consider height functions in various direction: is it possible to preprocess a given triangulated surface so that for any given direction, the contours of the height function for that direction can be computed I/O-efficiently?

6. REFERENCES

- [1] P. K. Agarwal, L. Arge, T. M. Murali, K. R. Varadarajan, and J. S. Vitter, I/O-efficient algorithms for contour-line extraction and planar graph blocking, *Proc. 9th ACM-SIAM Sympos. Discrete Algorithms*, 1998, pp. 117–126.
- [2] P. K. Agarwal, L. Arge, and K. Yi, I/O-efficient batched union-find and its applications to terrain analysis, *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, 2006, pp. 167–176.
- [3] A. Aggarwal and J. S. Vitter, The input/output complexity of sorting and related problems, *Commun. ACM*, 31 (1988), 1116–1127.
- [4] L. Arge, External memory data structures, in: *Handbook of Massive Data Sets* (J. Abello, P. M. Pardalos, and M. G. C. Resende, eds.), Kluwer Academic Publishers, 2002, pp. 313–358.
- [5] L. Arge, The buffer tree: A technique for designing batched external data structures, *Algorithmica*, 37 (2003), 1–24.
- [6] L. Arge, A. Danner, and S.-H. Teh, I/O-efficient point location using persistent B -trees, *Proc. Workshop on Algorithm Engineering and Experimentation*, 2003.
- [7] L. Arge, K. H. Hinrichs, J. Vahrenhold, and J. S. Vitter, Efficient bulk operations on dynamic R -trees, *Algorithmica*, 33 (2002), 104–128.
- [8] L. Arge, L. Toma, and N. Zeh, I/O-efficient topological sorting of planar dags, *Proc. 15th Annu. ACM Sympos. Parallel Algorithms and Architectures*, 2003, pp. 85–93.
- [9] H. Carr, J. Snoeyink, and U. Axen, Computing contour trees in all dimensions, *Computational Geometry: Theory and Applications*, 24 (2003), 75–94.
- [10] Y.-J. Chiang, M. T. Goodrich, E. F. Grove, R. Tamassia, D. E. Vengroff, and J. S. Vitter, External-memory graph algorithms, *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, 1995, pp. 139–149.
- [11] Y.-J. Chiang and C. T. Silva, I/O optimal isosurface extraction, *Proc. IEEE Visualization*, 1997, pp. 293–300.
- [12] A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge, and H. Mitasova, TERRASTREAM: From elevation data to watershed hierarchies, *Proc. ACM Sympos. on Advances in Geographic Information Systems*, 2007, 212–219.
- [13] H. Edelsbrunner, D. Letscher, and A. Zomorodian, Topological persistence and simplification, *Discrete Comput. Geom.*, 28 (2002), pp. 511–533.
- [14] I. Fáry, On straight lines representation of planar graphs, *Acta Sci. Math. Szeged*, 11 (1948), 229–233.
- [15] M. Isenburg, Y. Liu, J. Shewchuk, and J. Snoeyink, Streaming computation of Delaunay triangulations, *Proc. of SIGGRAPH*, 2006, pp. 1049–1056.
- [16] W. Lorensen and H. Cline, Marching cubes: a high resolution 3d surface construction algorithm, *Comput. Graph.*, 21 (1987), 163–170.
- [17] M. H. Nodine, M. T. Goodrich, and J. S. Vitter, Blocking for external graph searching, *Algorithmica*, 16 (1996), 181–214.
- [18] C. Silva, Y. Chiang, J. El-Sana, and P. Lindstrom, Out-of-core algorithms for scientific visualization and computer graphics, *Visualization'02*, 2002. Course Notes for Tutorial 4.
- [19] R. A. Skelton, Cartography, *History of Technology*, 6 (1958), 612–614.
- [20] J. van den Bercken, B. Seeger, and P. Widmayer, A generic approach to bulk loading multidimensional index structures, *Proc. International Conference on Very Large Databases*, 1997, pp. 406–415.
- [21] M. van Kreveld, R. van Oostrum, C. Bajaj, V. Pascucci, and D. Schikore, Contour trees and small seed sets for isosurface traversal, *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, 1997, pp. 212–219.
- [22] P. J. Varman and R. M. Verma, An efficient multiversion access structure, *IEEE Transactions on Knowledge and Data Engineering*, 9 (1997), 391–409.
- [23] J. S. Vitter, External memory algorithms and data structures: Dealing with MASSIVE data, *ACM Computing Surveys*, 33 (2001), 209–271.