# I/O-Efficient Contour Queries on Terrains 

Pankaj K. Agarwal* ${ }^{*} \quad$ Thomas Mølhave ${ }^{\dagger} \quad$ Bardia Sadri ${ }^{\ddagger}$


#### Abstract

A terrain $\boldsymbol{M}$ can be represented as a triangulation of the plane along with a height function associated with the vertices (and linearly interpolated within the edges and triangles) of $\boldsymbol{M}$. We investigate the problem of answering contour queries on $\boldsymbol{M}$ : Given a height $\ell$ and a triangle $f$ of $\boldsymbol{M}$ that intersects the level set of $\boldsymbol{M}$ at height $\ell$, report the list of the edges of the connected component of this level set that intersect $f$, sorted in clockwise or counterclockwise order. Contour queries are different from level-set queries in that only one contour (connected component of the level set) out of all those that may exist is expected to be reported. We present an I/Oefficient data structure of linear size that answers a contour query in $O\left(\log _{B} N+T / B\right) \mathrm{I} / \mathrm{Os}$, where $N$ is the number of triangles in the terrain and $T$ is the number of edges in the output contour. The data structure can be constructed using $O(\operatorname{Sort}(N))$ I/Os.


[^0]
## 1 Introduction

Over the last two decades there has been extensive work on modeling, analyzing, and visualizing terrain data in computational geometry, GIS, and spatial databases. A three-dimensional (digital elevation) model of a terrain is often represented as a triangulated $x y$-monotone surface, also known as a triangulated irregular network (TIN). This surface can be regarded as the graph of a piecewise-linear height or elevation function $h: \mathbb{R}^{2} \rightarrow \mathbb{R}$. A contour (or contour line) of a terrain $\boldsymbol{M}$ is a connected component of a level set of $h$. Each contour $K$ in a regular level set (see Section 2 for the definition) is a polygon in $\mathbb{R}^{2}$ (See Figure 1). Many terrain processing applications involve computing of one or many contours on a given terrain. In particular, a fundamental operation integral to such systems is answering contour queries: the user selects a point on the terrain (possibly using some provided rendering), and the system computes (and perhaps displays) the contour that goes through that particular point. These queries are related to level-set queries in which all contours at a chosen height, comprising an entire level set, are reported. Contour queries can be used, among other things, to compute the projected area of some feature of the terrain. In fact, computing the projected area of a hill or mountain was one of the earliest motivations for computing contours [14]. If $h$ represents other measured quantities over a planar domain, e.g. the barometric pressure, then contour queries can be used to compute, e.g. the domain or area of a high-pressure system. These queries also define many fundamental operations in terrain editing software, where a contour defines a "cut" in a terrain.

An alternative definition for contour queries replaces the query point with a pair $(f, \ell)$ where $\ell$ is a height and $f$ is a terrain triangle that spans height $\ell$. The target contour in this case is the (unique) contour at height $\ell$ that crosses $f$, denoted by $K_{M}(f, \ell) .{ }^{1}$ The contour-query problem thus can be formulated as follows: Preprocess a terrain $\boldsymbol{M}$ into a data structure so that for a height $\ell$ and a triangle $f$, the con-

[^1]

Figure 1: (a) the plane triangulation of a terrain together with a level set of that terrain consisting of two contours (shown in dotted lines). Only the bounded edges and triangles of the triangulation are displayed. (b) The three-dimensional interpretation of the same terrain.
tour $K_{M}(f, \ell)$ can be reported quickly in an outputsensitive manner - edges of the contour should be reported sorted in clockwise or counterclockwise order.

With recent advances in mapping technologies, such as laser based LIDAR, billions of points on a geographical terrain, at sub-meter resolution, can be acquired with very high accuracy ( $\sim 10-15 \mathrm{~cm}$ ) 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 [4]. 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 the number of I/Os it performs. See the recent surveys [5, 17] 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 $\operatorname{Sort}(N)$. It is common to assume that $M \geq B^{2}$. This is usually called the tall cache assumption.

Related work. A natural way of computing a contour $K$ of a terrain $M$ is simply to start at one triangle of $M$ that intersects $K$, trace out $K$ by walking through $M$ one triangle at a time, and stop on arriving at the starting point. 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 the best-known algorithm for blocking a planar triangulation $[1,15]$ only leads to an $O\left(T / \log _{2} B\right)$ I/O solution. See [7] for improved bounds in some special cases. The known I/O-efficient algorithms for answering level-set queries on terrains $[2,1,9]$ are not ideal for answering contour queries because the size of a level set to which a contour belongs can be much larger than the size of the contour. This is especially true for large global terrains that next-generation GIS software must face. A common approach when dealing with spatially restricted areas of large terrains is to create a subdivision by splitting the terrain into disjoint tiles of relatively small size, and then to extract contour information from a small subset of the tiles. However, finding the tiles containing a specific contour is hard and walking between tiles is in general no better than any other method of traversing a plane graph.

Agarwal et al. [2] proposed a linear-size data structure for storing a TIN $M$ of size $N$ so that a level-set query (but not a contour query) of output size $T$ can be answered in $O\left(\log _{B} N+T / B\right)$ I/Os. This data structure improves the previously best known I/O-efficient approaches [9, 1]: it reports the contours in circular order (unlike Chiang and Silva [9]), and it can be efficiently constructed in $O(\operatorname{Sort}(N))$ I/Os (unlike Agarwal et al. [1]). They prove the existence of, and present an efficient algorithm for computing, the so-called level ordering for any given $M$. A level ordering is a total order $\prec$ on the triangles of $M$ with the following two crucial properties:


Figure 2: Level-ordering on triangles intersecting two contours of the same level set. An arrow points to the successor of a triangle within the set.
(i) for any contour $K$ at any level, the $\prec$-sorted list of triangles that intersect $K$ enumerates these triangles in the cyclic order they appear along $K$ (in either clockwise or counterclockwise order), and
(ii) for any two contours $K_{1}$ and $K_{2}$ of the same level set, in the $\prec$-sorted list of triangles that intersect $K_{1} \cup K_{2}$, the triangles intersecting at least one of $K_{1}$ or $K_{2}$ appear contiguously (See Figure 2).

See Section 4 for precise definitions. In the $\prec-$ sorted sequence of all triangles that intersect a given level set, those triangles that intersect a particular contour appear in maximal contiguous subsequences, called fragments. Agarwal et al. [2] showed how the entire collection of contours, i.e., the whole level set, can be retrieved quickly from these fragments using (i) and (ii). Their data structure can be modified to answer a contour query using $O\left(\log _{B} N+J \log _{M / B} N / B+T / B\right)$ I/Os, where $T$ is the number of triangles that intersect the output contour and $J$ is the number of its fragments in the level ordering used by the algorithm. Unfortunately, $J$ may be as large as $T$ in the level ordering computed by Agarwal et al., so the query uses $O\left(\log _{B} N+T\right)$ I/Os in the worst case.

Our results. In this paper we propose an I/Oefficient algorithm for preprocessing $\boldsymbol{M}$, using $O(\operatorname{Sort}(N)) \mathrm{I} / \mathrm{Os}$, into a linear-size data structure so that a contour query can be answered in $O\left(\log _{B} N+\right.$ $T / B) \mathrm{I} / \mathrm{Os}$, where $N$ is the number of triangles in $\boldsymbol{M}$ and $T$ is the size of the query output. The structure relies on the level-ordering method of Agarwal et al. [2]. It tackles the fundamental problem of contour fragmentation by transforming $\boldsymbol{M}$ into another terrain $\boldsymbol{M}^{\prime}$ that has two key properties: (i) $\boldsymbol{M}^{\prime}$ "preserves" individual contours of $\boldsymbol{M}$. (ii) The image of each contour of $\boldsymbol{M}$ on $\boldsymbol{M}^{\prime}$ has $O(1)$ fragments in
the level ordering computed on $\boldsymbol{M}^{\prime}$. The key contributions of our paper are the following:
Stretchings We describe in Section 3 a general operation, called "stretching", that modifies a terrain $\boldsymbol{M}$ into $\boldsymbol{M}^{\prime}$ by changing the heights of some vertices but maintains its underlying triangulation. We prove that the "stetched" terrain $\boldsymbol{M}^{\prime}$ has the same (augmented) contour tree as that of $\boldsymbol{M}$. We preprocess $\boldsymbol{M}$ and $\boldsymbol{M}^{\prime}$ in $O(\operatorname{Sort}(N))$ I/Os into a linear-size data structure (see Section 3) that converts, in $O\left(\log _{B} N\right) \mathrm{I} / \mathrm{Os}$, any contour query $(f, \ell)$ on $\boldsymbol{M}$ into a contour query $\left(f, \ell^{\prime}\right)$ on $\boldsymbol{M}^{\prime}$ such that $K_{M}(f, \ell)$ and $K_{M^{\prime}}\left(f, \ell^{\prime}\right)$ intersect the same sequence of triangles of their commmon underlying triangulation.
Bounding Fragmentation In Section 4, we describe an algorithm that computes a specific stretching $\boldsymbol{M}^{\prime}$ of $\boldsymbol{M}$ and a level ordering of $\boldsymbol{M}^{\prime}$, using $O(\operatorname{Sort}(N))$ I/Os, such that each contour of $\boldsymbol{M}$ maps to a contour in $\boldsymbol{M}^{\prime}$ with $O(1)$ fragments in the level ordering of $\boldsymbol{M}^{\prime}$.

## 2 Preliminaries

A triangulation of $\mathbb{R}^{2}$ is a planar subdivision of $\mathbb{R}^{2}$ in which each face is a (possibly unbounded) triangle. Let $M$ be a triangulation, and let $V(M), E(M)$, and $F(M)$ denote the sets of vertices, edges, and faces (triangles), respectively, of $M$. We add a vertex $v_{\infty}$ at infinity and connect all unbounded edges (rays) of $M$ to $v_{\infty}$. Triangles of $M$ that have $v_{\infty}$ as a vertex are unbounded and all other triangles are bounded. A height function consistent with $M$ is any continuous function $h: \mathbb{R}^{2} \rightarrow \mathbb{R}$ whose restriction to each triangle of $M$ is linear and approaches $-\infty$ at infinity, i.e., $h\left(v_{\infty}\right)=-\infty$. A terrain $\boldsymbol{M}=(M, h)$ is a triangulation $M$ of $\mathbb{R}^{2}$ endowed with a height function $h$ that is consistent with $M . M$ can be viewed as the graph of $h$, an $x y$-monotone piecewise-linear surface in $\mathbb{R}^{3}$. For simplicity we assume that vertices of terrains have distinct heights. We sometimes write an edge $u v$ as a pair $(u, v)$ to indicate that $h(u)<h(v)$, i.e., we implicitly regard $u v$ to be oriented from its lower endpoint $u$ to its higher endpoint $v$.

Critical vertices. For any vertex $v$ the $\operatorname{link} \operatorname{lk}(v)$ is the cycle formed by the edges that are not incident to $v$ but belong to triangles that are incident to $v$. The upper (resp. lower) link $\mathrm{lk}^{+}(v)$ (resp. $\mathrm{lk}^{-}(v)$ ) is the subgraph of $\mathrm{lk}(v)$ induced by vertices $u$ with $h(u)>$ $h(v)($ resp. $h(u)<h(v))$. If $\mathrm{lk}^{-}(v)\left(\right.$ resp. $\left.\mathrm{lk}^{+}(v)\right)$ is empty then $v$ is a (local) minimum (resp. maximum). An extremum is a minimum or a maximum (See Figure 3). A non-extremal vertex $v$ is called regular if


Figure 3: Determining the type of a vertex; lower link vertices are depicted by filled circles.
$\mathrm{lk}^{+}(v)\left(\right.$ or $\left.\mathrm{lk}^{-}(v)\right)$ is connected, otherwise it is called a saddle. Non-regular vertices (extrema and saddles) are called critical, and their heights are critical values of $h$. For simplicity, we assume that all saddles are simple, i.e., their upper (equivalently lower) links have exactly two connected components. (see [13] and Figure 3).

Level sets and contours. For $\ell \in \mathbb{R}$, the level set, sublevel set, and superlevel set of a terrain $\boldsymbol{M}=$ $(M, h)$ at height $\ell$ are subsets $\boldsymbol{M}_{\ell}, \boldsymbol{M}_{<\ell}$, and $\boldsymbol{M}_{>\ell}$ of $\mathbb{R}^{2}$ consisting of points $x$, respectively with $h(x)=\ell$, $h(x)<\ell$, and $h(x)>\ell . \boldsymbol{M}_{\leq \ell}$ and $\boldsymbol{M}_{\geq \ell}$ are defined analogously.

Each connected component of a level set $\boldsymbol{M}_{\ell}$ is called a contour of $\boldsymbol{M}$ at level $\ell$. A contour of $\boldsymbol{M}_{\ell}$ can be represented by a pair $(f, \ell) \in F(M) \times \mathbb{R}$, as the contour $K_{M}(f, \ell)$ at height $\ell$ that intersects the triangle $f$ of $M$. Each vertex $v \in V(M)$ is contained in exactly one contour in $\boldsymbol{M}_{h(v)}$, which we call the contour of $v$. Contours of regular vertices are simple polygons; contours of extrema are single points; and contours of saddles are polygons that self-intersect exactly once at their corresponding saddles. Being a Jordan curve, the contour $K$ of a regular vertex partitions $\mathbb{R}^{2} \backslash K$ into the inside, $K^{\text {in }}$, and outside, $K^{\text {out }}$, of $K$.


Figure 5: The 3D interpretation of $\bar{F}_{M}(K)$ for a contour $K$.

A contour is generic if it is not the contour of any vertex of the terrain. A generic level set $\boldsymbol{M}_{\ell}$ is a level set in which all contours are generic, i.e., $\ell \neq h(v)$ for any vertex $v$. The set of triangles that intersect $\boldsymbol{M}_{\ell}$ is denoted by $F\left(\boldsymbol{M}_{\ell}\right)$. For a generic
contour $K$ in $\boldsymbol{M}_{\ell}$, let $F_{\boldsymbol{M}}(K)$ be the set of triangles in $F(M)$ that intersect $K .{ }^{2}$ Each triangle in $F(K)$ meets $K$ at a single segment. The order in which $K$ intersects the triangles induces a cyclic ordering on $F(K)$. Let $\bar{F}(K)$ be the resulting cyclically ordered set (we take a cyclic ordering and its reverse as equal; See Figure 5). We refer to $\bar{F}(K)$ as the (combinatorial) representation of $K$. Given $\bar{F}(K)$, $K$ (as a plane polygon) can be generated in one scan of $\bar{F}(K)$. Therefore it suffices to generate the cyclic sequence $\bar{F}(K(f, \ell))$ to answer a contour query $(f, \ell)$.

Let $V^{-}(K)$ and $V^{+}(K)$ respectively be vertices from triangles in $F(K)$ with heights smaller and greater than $\ell$. The lower and upper supporting vertices of $K$ are respectively defined as

$$
\begin{aligned}
& v^{-}(K)=\arg \max _{v \in V^{-}(K)} h(v) \\
& v^{+}(K)=\arg \min _{v \in V^{+}(K)} h(v)
\end{aligned}
$$

Two contours $K$ and $K^{\prime}$ are equivalent if $\bar{F}(K)=$ $\bar{F}\left(K^{\prime}\right)$. It can be verified that $K$ and $K^{\prime}$ are equivalent if and only if they have the same lower and upper supporting vertices.

Up- and down-contours. Let $\varepsilon=\varepsilon(\boldsymbol{M})$ denote a sufficiently small positive value, in particular, smaller than height difference between any two vertices of $\boldsymbol{M}$. An up-contour of a vertex $v$ is any contour of $\boldsymbol{M}_{h(v)+\varepsilon}$ that intersects an edge incident on $v$. Similarly, a down-contour of $v$ is any contour of $\boldsymbol{M}_{h(v)-\varepsilon}$ that intersects an edge incident on $v$. If $v$ is maximum (resp. minimum), then its up- (resp. down-) contour is not defined. If $v$ is not a saddle, then its up- and down-contours are unique and they converge to the contour of $v$ as $\varepsilon \rightarrow 0$. A saddle is called positive if it has two up-contours and one down-contour and negative if it has two down-contours and one upcontour. All simple saddles are either negative or positive.

[^2]

Figure 4: Left: a terrain with one positive and one negative saddle. The saddle contours and their up and down contours are marked in their corresponding colors (or, dotted for blue and dashed for red). Middle: the sub level set of the lower saddle, indicating (together with the colors of the up- and down-contours that it is a red positive saddle. Right: sub level set of the upper saddle: a red negative one.

For a generic contour $K$, if $v^{-}(K)$ has a unique up-contour $K_{v}$, then $K$ and $K_{v}$ are equivalent. Otherwise, $v^{-}(K)$ is a positive saddle and has two upcontours, one of which is equivalent to $K$ and the other intersects no triangles in $F(K)$. The same can be said about down-contour(s) of $v^{+}(K)$.

Red and blue contours and vertices. A contour $K$ of $\boldsymbol{M}_{\ell}$ is called blue if $\boldsymbol{M}_{\leq \ell}$ is locally in the interior of $K$ and red otherwise. Equivalently, a contour $K$ is blue if for any (oriented) edge ( $u, v$ ) intersecting $K, u \in K^{\text {in }}$ and $v \in K^{\text {out }}$, and red otherwise. We define the color of a regular vertex to be the color of the contour that passes through it, which agrees with the color of its up- and down-contours. The color of $v_{\infty}$ is taken to be red, and the color of all other minima is blue; i.e., the color of their up-contour; the color of a maximum is taken to be red, i.e., the color of its down-contour. The color of a positive (resp. negative) saddle is taken to be the color of its unique down-contour (resp. upcontour). Under this coloring scheme, a red positive saddle will have a red down-contour and two red upcontours while a blue positive saddle will have a blue down-contour together with one red and one blue up-contours (Figure 6). Similarly, a blue negative saddle has all blue up- and down-contours while a red negative saddle has one red down-contour, one blue down-contour, and one red up-contour. The following lemma is equivalent to Lemma 3.1 of [3].

Lemma 2.1. [3] If a generic red (resp. blue) contour $K$ in $\boldsymbol{M}_{\ell}$ passes through a triangle $f$, then the vertices of $f$ with height greater than (resp. smaller than) $\ell$ are red (resp. blue).

Contour, join, and split trees. Finally, we define the contour tree, and its two relatives the join and
split trees of a terrain $[8,11],{ }^{3}$ which will play crucial roles in our algorithms. They respectively encode the topology of the level sets, sublevel sets, and superlevel sets of $\boldsymbol{M}$ at various heights. The (augmented) contour tree $\mathcal{C}_{\boldsymbol{M}}$ of $\boldsymbol{M}$ (or just $\mathcal{C}$ when $\boldsymbol{M}$ is clear from the context) is a tree on vertex set $V(M)$ in which $u v$ is an edge if $u$ is the upper and $v$ is the lower supporting vertex of some generic contour $K$ of $\boldsymbol{M}$. The edge $u v$ of $\mathcal{C}$ is then said to carry $K$. Equivalently, the contour tree of $\boldsymbol{M}$ is the space that is the quotient of $\boldsymbol{M}$ modulo taking points on the same contour as equivalent. Under the associated quotient $\operatorname{map} \pi_{\boldsymbol{M}}: \boldsymbol{M} \rightarrow \mathcal{C}_{\boldsymbol{M}}$, the image of each contour of $\boldsymbol{M}$ is a unique point in $\mathcal{C}_{\boldsymbol{M}}$. Contours of vertices of $\boldsymbol{M}$ are mapped into points that are regarded as corresponding vertices of $\mathcal{C}$ while generic contours are mapped into points along edges that carry them (See Figure 7. We color each vertex $v$ of $\mathcal{C}$ with the color of contour of $v$ on $\boldsymbol{M}$. Points along the edges of the contour tree are also given the color of the contours they represent. It can be verified that all contours mapped to the interior of the same edge of $\mathcal{C}$ have the same color. We choose $v_{\infty}$ as the root of the contour tree.

The augmented join tree $\mathcal{J}$ represents the merging of connected components of $\boldsymbol{M}_{<\ell}$ and the augmented split tree $\mathcal{S}$ represents the splittings of the connected components of $\boldsymbol{M}_{>\ell}$ when $\ell$ varies from $-\infty$ to $+\infty$. More precisely, $\mathcal{J}_{M}$ is a tree on vertex set $V(M)$ in which $u v$ is an edge if $v$ is the highest vertex in a connected component of $\boldsymbol{M}_{<h(u)}$ that contains a vertex of $\mathrm{lk}^{-}(u) . \quad \mathcal{S}_{M}$ is defined symmetrically by negating the height function $h$. We choose $v_{\infty}$ as the

[^3]

Figure 6: The four possible saddle types. Left to right: red positive, blue positive, blue negative, red negative. The up- and down-contours of the saddle points, but not the saddle contours themselves, are shown.


Figure 7: (a) A terrain. Level sets of all vertices shown (dashed for blue). Numbers indicate heights of vertices. (b) (left to right): colored contour tree, split tree, and join tree of the same terrain.
root of $\mathcal{S}$ and the highest maximum of $\boldsymbol{M}$ as the root of $\mathcal{J}$.

## 3 Stretching Terrains

In this section we introduce stretching as a general operation on a terrain $\boldsymbol{M}$, which "preserves" individual contours (as described below) while moving them amongst level sets. We then present an I/Oefficient algorithm for mapping a contour $K_{M}(f, \ell)$ to an equivalent contour on the stretched terrain.

The stretching operation. Let $K$ be a generic contour of a terrain $\boldsymbol{M}=(M, h)$, and let $\delta$ be an offset parameter; $\delta>0$ if $K$ is red and $\delta<0$ if $K$ is blue. The stretching of $\boldsymbol{M}$ at $K$ by $\delta$ is a terrain $\boldsymbol{M}^{\prime}=\left(M, h^{\prime}\right)$ where for any vertex $v$ :

$$
h^{\prime}(v)= \begin{cases}h(v)+\delta & \text { if } v \in K^{\text {in }} \\ h(v) & \text { if } v \in K^{\text {out }}\end{cases}
$$

A terrain $\boldsymbol{M}^{\prime}$ is called a stretching of $\boldsymbol{M}$ if there is a sequence $\boldsymbol{M}=\boldsymbol{M}_{1}, \ldots, \boldsymbol{M}_{k}=\boldsymbol{M}^{\prime}$ of terrains where each $\boldsymbol{M}_{i}$ is obtained from $\boldsymbol{M}_{i-1}$ by stretching it at one of its generic contours. Note that the underlying triangulations of $\boldsymbol{M}^{\prime}$ and $\boldsymbol{M}$ are the same. If $K_{1}$ and $K_{2}$ are equivalent generic contours of $\boldsymbol{M}$ with $K_{2} \subset K_{1}^{\text {in }}$, then there are no vertices
in $K_{1}^{\text {in }} \cap K_{2}^{\text {out }}$. Therefore, stretching $\boldsymbol{M}$ at $K_{1}$ or $K_{2}$ by $\delta$ produces the same terrain. The following lemma shows that certain properties of a terrain are preserved by stretching.

Lemma 3.1. Let $\boldsymbol{M}^{\prime}=\left(M, h^{\prime}\right)$ be a stretching of a terrain $\boldsymbol{M}=(M, h)$.
(1) For any adjacent vertices $u$ and $v$ of $M, h(u)>$ $h(v)$ if and only if $h^{\prime}(u)>h^{\prime}(v)$.
(2) The type (regular, minimum, maximum, positive saddle, or negative saddle) of any vertex $v \in$ $V(M)$ is the same in both $\boldsymbol{M}$ and $\boldsymbol{M}^{\prime}$.
(3) For any vertex $v$, if $K_{v}$ is the up-contour (resp. down-contour) of $v$ in $\boldsymbol{M}$ that intersects an edge vw and $K_{v}^{\prime}$ is the up-contour (resp. downcontour) of $v$ in $\boldsymbol{M}^{\prime}$ that also intersects $v w$, then $\bar{F}_{M}\left(K_{v}\right)=\bar{F}_{M^{\prime}}\left(K_{v}^{\prime}\right)$.

Proof. It suffices to prove the lemma assuming that $\boldsymbol{M}^{\prime}$ is the result of applying a single simple stretching at a contour $K$ of $\boldsymbol{M}$. The more general statement then follows inductively.
(1) Suppose $K$ is a red contour of $\boldsymbol{M}$ and therefore $\delta>0$. The endpoints of any edge $(u, v)$ that does not intersect $K$ lie at the same side of $K$
and therefore their heights are modified (if at all) by the same amount. For any (oriented) edge ( $u, v$ ) of $\boldsymbol{M}$ that intersects $K, u \in K^{\text {out }}$ and $v \in K^{\text {in }}$. Using the fact that $\delta>0$, we obtain $h^{\prime}(u)=h(u)<h(v)<$ $h(v)+\delta=h^{\prime}(v)$.

By part (1), the lower and upper links of every vertex remain the same after stretching. Therefore, regular vertices, extrema, and saddles of $\boldsymbol{M}$ remains the same in $\boldsymbol{M}^{\prime}$. Furthermore, since $K$ is taken to be generic, up- and down-contour(s) of any saddle are all contained either in $K^{\text {in }}$ or all in $K^{\text {out. }}$. Therefore the type of a saddle (positive or negative) is not affected by the stretching operation. The case where $K$ is blue is similar.

For part (3), assume that $K$ is red (the other case is symmetric). Let $K_{v}$ be an up-contour of $v$ and assume first that $K \subset K_{v}^{\text {in }}$. If $K_{v}$ is blue, then no edge $(a, b)$ of $\boldsymbol{M}$ that intersects $K_{v}$ can also intersect $K$ since this would mean that $b \in K_{v}^{\text {out }}$ and $b \in K^{\text {in }}$ while $K \subset K_{v}^{\text {in }}$. Therefore in this case all the triangles in $F_{M}\left(K_{v}\right)$ are contained in $K^{\text {out }}$ and remain unaffected by the stretching at $K$. Therefore, $K_{v}$ remains at the same height in $\boldsymbol{M}^{\prime}$ and continues to be an up-contour of $v$. Since it also continues to intersects the edge $v w$, it coincides with $K_{v}^{\prime}$ and in particular $\bar{F}_{M}\left(K_{v}\right)=\bar{F}_{M^{\prime}}\left(K_{v}^{\prime}\right)$.

If $K_{v}$ is red, then $V^{-}\left(K_{v}\right) \subset K_{v}^{\text {out. Thus }}$ $h^{\prime}(u)=h(u)<h(v)$ for any $u \in V^{-}\left(K_{v}\right)$ while $h^{\prime}(u) \geq h(u)>h(v)$ for all $u \in V^{+}\left(K_{v}\right)$. This means that the triangles in $F_{M}\left(K_{v}\right)$ all intersect some upcontour of $v$ in $\boldsymbol{M}^{\prime}$. Since the triangles incident to the edge $v w$ are among these triangles, the mentioned up-contour has to be $K_{v}^{\prime}$.

The case where $K$ is contained in $K_{v}^{\text {out }}$ is handled similarly and the argument for down-contours is done symmetrically.

Corollary 3.1. If $\boldsymbol{M}^{\prime}$ is a stretching of $\boldsymbol{M}$, then $\mathcal{C}_{\boldsymbol{M}}\left(\right.$ resp. $\left.\mathcal{J}_{M}, \mathcal{S}_{\boldsymbol{M}}\right)$ and $\mathcal{C}_{M^{\prime}}\left(\right.$ resp. $\left.\mathcal{J}_{M^{\prime}}, \mathcal{S}_{M^{\prime}}\right)$ are combinatorially identical.

The stretching operation has a natural interpretation in terms of the contour trees. Let $K$ be a contour carried by an edge $u v$ of $\mathcal{C}$. Let $v$ be the parent of $u$ in $\mathcal{C}$ (with respect to the root $v_{\infty}$ of $\mathcal{C}$ ). Let $\mathcal{C}^{u}$ denote the subtree of $\mathcal{C}$ rooted at $u$. It can be verified that the vertices of $\mathcal{C}^{u}$ are exactly the vertices of $\boldsymbol{M}$ that are contained in $K^{\text {in }}$. In particular, the leaves of $\varrho^{u}$ are the extrema in $K^{\text {in }}$. Furthermore, if $K^{\prime}$ is a contour carried by an edge of $\mathcal{C}^{u}$, then $K^{\prime} \subset K^{\text {in }}$. Thus stretching $\boldsymbol{M}$ at $K$ by $\delta$ corresponds to adding $\delta$ to the heights of all vertices in $\mathcal{C}^{u}$. Since stretching $\boldsymbol{M}$ at equivalent contours, i.e., contours carried by the same edge of $\mathcal{C}$, results the same stretched terrain, one can regard stretching $\boldsymbol{M}$ at a contour $K$ by


Figure 8: The effect of stretching on contour tree: on the left a red edge $e_{1}$ and a blue edge $e_{2}$ of a contour tree $\mathcal{C}$ are respectively stretched by $\delta_{1}>0$ and $\delta_{2}<0$. The heights of vertices in the subtree rooted at the tail of $e_{1}$ are increased by $\delta_{1}$. Those vertices in the subtree rooted at tail of $e_{2}$ are displaced by an extra $\delta_{2}$ amount.
$\delta$ as associating an offset amount $\delta$ to the edge of $\mathcal{C}$ that carries $K$. The total change to $h(v)$ for a vertex $v$ after a series of stretches is then determined by the sum of the offsets over the edges on the path from $v$ to $v_{\infty}$ in $\mathcal{C}$ (See Figure 8).

Mapping contour queries. Part (3) of Lemma 3.1 along with the discussion on up- and down-contours implies that if one canonically represent a generic contour $K$ by, say, the appropriate up-contour of $v=$ $v^{-}(K)$, then $\bar{F}(K)$ is identical to $\bar{F}\left(K^{\prime}\right)$ where $K^{\prime}$ is the corresponding up-contour of $v$ in the stretching $\boldsymbol{M}^{\prime}$ of $\boldsymbol{M}$; the same is true for the down-contour of $v^{+}(K)$. In other words, a generic contour $K$ in $\boldsymbol{M}$ can be mapped to one in $\boldsymbol{M}^{\prime}$ that intersects the same set of triangles in the same order. Since a contour is represented by a pair $(f, \ell)$, we need to find the lower (or upper) supporting vertex of $K_{M}(f, \ell)$ to map it to $\boldsymbol{M}^{\prime}$. The following lemma shows how join and split trees can be used to do this.

Lemma 3.2. Let $\boldsymbol{M}$ be a terrain with augmented join and split trees $\mathcal{J}$ and $\mathcal{S}$. Let $K$ be a generic blue (resp. red) contour in $\boldsymbol{M}_{\ell}$ and let $u$ be a vertex in $V^{-}(K)$ (resp. $V^{+}(K)$ ). Then the lower (resp. upper) supporting vertex of $K$ is the highest blue (resp. lowest red) vertex $v$ on the path from $u$ to the root of $\mathcal{J}$ (resp. S) that satisfies $h(v)<\ell$ (resp. $h(v)>\ell)$.

Proof. We only prove the lemma for the case where $K$ is blue; the other case is similar. Let $v=v^{-}(K)$ be
the lower supporting vertex of $K$, and let $u \in V^{-}(K)$. Let $r_{\mathcal{J}}$ denote the root of $\mathcal{J}$, i.e. the highest maximum in $\boldsymbol{M}$. We show that $v$ is the highest blue vertex with height less than $\ell$ on the path between $u$ and $r_{\mathcal{J}}$ in $\mathcal{J}$.

Let $f$ be a triangle incident upon $u$ that intersects $K$ and let $K_{u}$ be the up-contours of $u$ that intersect $f$. By Lemma $2.1 u$ is blue and therefore so is $K_{u}$. For any $\alpha$, let $U_{\alpha}$ be the connected component of $\boldsymbol{M}_{\leq \alpha}$ that contains $u$. At $\alpha=h(u), u$ lies on the boundary of $U_{\alpha}$. By the definition of $\mathcal{J}$, as $\alpha$ grows larger, every time the boundary of $U_{\alpha}$ reaches a new vertex $w$, this vertex is turned into the parent of the highest vertex of $U_{\alpha}$ in $\mathcal{J}$. The color of $w$ is determined by the color of the contour that is the boundary component of $U_{\alpha}$ that reaches $w$. We argue that $v$ is the highest blue vertex on $P$ with $h(v)<\ell$ : being a connected component of a sublevel set with a blue contour on its boundary, $U_{h(v)+\varepsilon}$ has to be bounded by a single blue contour $K_{0}$ (which is the up-contour of $v$ equivalent to $K$ ), and zero or more red contours $K_{1}, \ldots, K_{r}$ where $K_{i} \subset K_{0}^{\text {in }}$ for all $i=1, \ldots, r$. When $\alpha$ sweeps the interval $[h(v)+\varepsilon, \ell]$, the blue boundary component of $U_{\alpha}$ continuously changes from $K_{0}$ to $K$ without passing through any vertices (since $v$ is the lower supporting vertex of $K$ ). This means that all other vertices encountered at the boundary of $U_{\alpha}$ (and therefore added to $P$ between $v$ and $r_{\jmath}$ ) are met by one of the red boundary components of $U_{\alpha}$ and are therefore all red.

Let $(f, \ell)$ be the input to a contour query. Let $v_{0}$ (resp. $v_{1}$ ) be the lowest (resp. highest) vertex of $f$. $K=K(f, \ell)$ intersects the edge $v_{0} v_{1}$. By Lemma 2.1, if $K$ is blue then $v_{0}$ is blue, and if $K$ is red then $v_{1}$ is red. By Lemma 3.2, if $K$ is blue, then $v^{-}(K)$ is the highest blue vertex in $\mathcal{J}$ between $v_{0}$ and the root with height less than $\ell$, and if $K$ is red, then $v^{+}(K)$ is the lowest red vertex in $\mathcal{S}$ between $v_{1}$ and the root with height greater than $\ell$. If the colors of $v_{0}$ and $v_{1}$ agree, $K$ has to be of the same color. Otherwise, by Lemma 2.1, $v_{0}$ must be blue, and $v_{1}$ red, and the color of contours that intersect $v_{0} v_{1}$ flips once from blue to red at a unique height which we denote by $\eta\left(v_{0} v_{1}\right)$. The following lemma is relatively straightforward:

Lemma 3.3. $\eta\left(v_{0} v_{1}\right)$ is the height of the saddle that is the lowest common ancestor of $v_{0}$ and $v_{1}$ in the contour tree of $\boldsymbol{M}$.

One can compute $\eta(e)$ of all terrain edges $e$ with nonmatching endpoint colors in $O(\operatorname{Sort}(N))$ I/Os using the so-called time-forward processing technique [10].

We can compute one of $v^{-}(K)$ or $v^{+}(K)$ as follows: We preprocess the blue (resp. red) vertices of $\mathcal{J}$ (resp. $\mathcal{S}$ ) in a persistent B-tree $[6,16] B_{\mathcal{J}}$ (resp. $B_{\mathcal{S}}$ ), using their heights as the key. Algorithm

Compute $B_{\mathcal{J}}$ (Figure 9 left) computes $B_{\mathcal{J}}$. A similar algorithm is used to compute $B_{8}$. The recorded step number $\tau(v)$ is later used to access the instance $B_{\mathcal{J}}[\tau(v)]$ of $B_{\mathcal{J}}$ in which $v$ was first inserted. $B_{\mathcal{J}}[\tau(v)]$ stores the sequence of blue vertices on the path from $v$ to the root, sorted by their heights. Algorithm Translatequery (Figure 9 right) computes $v^{-}(K)$ if $K$ is blue and $v^{+}(K)$ if $K$ is red. It then returns a pair $\left(f, \ell^{\prime}\right)$ where $\ell^{\prime}$ is the height in $M^{\prime}$ of the upcontour $v^{-}(K)$ in the former case, or that of the down-contour of $v^{+}(K)$ in the latter case. We have thus shown the following:

Theorem 3.1. Let $\boldsymbol{M}=(M, h)$ be a terrain of size $N$ and let $M^{\prime}=\left(M, h^{\prime}\right)$ be a stretching of $M$. $M$ and $M^{\prime}$ can be preprocessed using $O(\operatorname{Sort}(N))$ $I / O s$ into a linear size data structure that converts, in $O\left(\log _{B} N\right) I / O s$, any contour query $(f, \ell)$ on $\boldsymbol{M}$ into a contour query $\left(f, \ell^{\prime}\right)$ on $\boldsymbol{M}^{\prime}$ such that $\bar{F}_{M}(K(f, \ell))=\bar{F}_{M^{\prime}}\left(K\left(f, \ell^{\prime}\right)\right)$.

## 4 Bounding Fragmentations by Stretching

We now describe an algorithm for computing a stretching of $M$ that enables us to answer contour queries quickly.

Level-ordering and fragmentations. Let $\rho$ : $F(M) \rightarrow \mathbb{N}$ be an injective rank function on the triangles of $\boldsymbol{M}$. A generic contour $K$ agrees with $\rho$ if the circular list $\bar{F}(K)$ can be written as $\left\langle f_{1}, \ldots, f_{k}, f_{1}\right\rangle$, such that $\rho\left(f_{1}\right)<\cdots<\rho\left(f_{k}\right)$. The function $\rho$ is a level-ordering for $M$ if the following two conditions hold:

C1. Any generic contour $K$ agrees with $\rho$.
C 2 . Let $K, K^{\prime}$ two generic contours of a level set $\ell$, let $f_{1}, f_{2} \in F(K)$, and let $f_{1}^{\prime}, f_{2}^{\prime} \in F\left(K^{\prime}\right)$. If $\rho\left(f_{1}\right)<$ $\rho\left(f_{1}^{\prime}\right)<\rho\left(f_{2}\right)$ then $\rho\left(f_{1}\right)<\rho\left(f_{2}^{\prime}\right)<\rho\left(f_{2}\right)$.

Given a level-ordering $\rho$ on $\boldsymbol{M}$, let $F_{\rho}\left(\boldsymbol{M}_{\ell}\right)$ denote the list of triangles that intersect $\boldsymbol{M}_{\ell}$, ordered by $\rho$. A fragment of a generic contour $K$ in $\boldsymbol{M}_{\ell}$ is a maximal subsequence $f_{1}, \ldots, f_{r}$ of $F_{\rho}\left(\boldsymbol{M}_{\ell}\right)$ consisting of triangles in $F(K)$. $\boldsymbol{M}$ may have several level-orderings, but we focus on the level-ordering computed by Agarwal et al. [2]. As mentioned in Introduction, the only obstacle in achieving efficient contour queries using the algorithm of Agarwal et al. [2] is the lack of any control over the number of fragments of contours under the computed level-ordering. Theorem 3.1 shows that if some stretching $\boldsymbol{M}^{\prime}$ of the input terrain $\boldsymbol{M}$ avoids this obstacle, at least for the contours of $\boldsymbol{M}^{\prime}$ that are mapped from $\boldsymbol{M}$, then it can be used in place of $\boldsymbol{M}$. This section proves the existence of such a stretching and gives an algorithm to compute it in

```
Algorithm: Compute \(B_{\mathcal{~}}\)
    Compute an Euler tour \(\Pi\) of \(\mathcal{J}_{M}\) starting at its root
    for a vertex \(v\) visited at step \(i\) of \(\Pi\) do
        if \(v\) is blue
            if this is the first visit to \(v\)
                    insert \(v\) into \(B_{\mathcal{J}}\); assign time-stamp \(i\)
                    to \(B_{\mathcal{J}}\); and record \(\tau(v)=i\)
            if this is the last visit to \(v\)
                    delete \(v\) from \(B_{\mathcal{J}}\)
```

```
Algorithm: TranslateQuery \((f, \ell)\)
```

Algorithm: TranslateQuery $(f, \ell)$
$v_{0} \leftarrow$ lowest vertex of $f$.
$v_{0} \leftarrow$ lowest vertex of $f$.
$v_{1} \leftarrow$ highest vertex of $f$.
$v_{1} \leftarrow$ highest vertex of $f$.
if $v_{1}$ is blue or $\eta\left(v_{0} v_{1}\right)<\ell$
if $v_{1}$ is blue or $\eta\left(v_{0} v_{1}\right)<\ell$
$u \in B_{\mathcal{J}}\left[\tau\left(v_{0}\right)\right]$ highest vertex with $h(u)<\ell$
$u \in B_{\mathcal{J}}\left[\tau\left(v_{0}\right)\right]$ highest vertex with $h(u)<\ell$
return $\left(f, h^{\prime}(u)+\varepsilon\right)$
return $\left(f, h^{\prime}(u)+\varepsilon\right)$
else
else
$u \in B_{\mathcal{S}}\left[\tau\left(v_{1}\right)\right]$ lowest vertex with $h(u)>\ell$
$u \in B_{\mathcal{S}}\left[\tau\left(v_{1}\right)\right]$ lowest vertex with $h(u)>\ell$
return $\left(f, h^{\prime}(u)-\varepsilon\right)$

```
        return \(\left(f, h^{\prime}(u)-\varepsilon\right)\)
```

Figure 9: Algorithms for terrain preprocessing (left) and conversion of contour queries (right).
$O(\operatorname{Sort}(N))$ I/Os. We begin by introducing the notion of ascending and descending cut-trees, which are used by Agarwal et al.'s level-ordering algorithm.

Ascending and descending cut-trees. For a terrain edge $e=(u, v)$, let $\lambda^{+}(e)$ be the first outgoing edge of $v$ in clockwise order after $e$, and let $\lambda^{-}(e)$ be the first incoming edge of $u$ in clockwise order before $e$. A rightmost ascending (leftmost descending) path $P$ from $v_{1}$ to $v_{k}$ is a path $v_{1} \ldots v_{k}$ where for each $1<i<k,\left(v_{i}, v_{i+1}\right)=\lambda^{+}\left(\left(v_{i-1}, v_{i}\right)\right)$ (resp. $\left(v_{i+1}, v_{i}\right)=\lambda^{-}\left(\left(v_{i}, v_{i-1}\right)\right)$. An ascending (resp. descending) path $P=v_{1} \ldots v_{k}$ is maximal if $v_{k}$ is a maximum (resp. minimum). Let $V_{\ominus}, V_{\oplus}$ respectively be the sets of negative and positive saddles of $\boldsymbol{M}$. For each positive saddle $s$, let $P_{1}(s)$ and $P_{2}(s)$ be two maximal rightmost ascending paths from $s$ such that the vertex following $s$ on each of $P_{1}$ and $P_{2}$ is the counterclockwise first neighbor of $s$ in a distinct upper link component of $s$. Similarly, for each negative saddle $s$, let $P_{1}(s)$ and $P_{2}(s)$ be maximal leftmost descending paths from $s$ such that the vertex following $s$ in each of $P_{1}$ and $P_{2}$ is the clockwise first neighbor of $s$ in a distinct lower link component of $s$. For each saddle $s$, the subgraph $P(s)=P_{1}(s) \cup P_{2}(s)$ of $M$ is uniquely determined. We call the subgraphs $\hat{\boldsymbol{T}}=\cup_{s \in V_{\oplus}} P(s)$ and $\check{\boldsymbol{T}}=\cup_{s \in V_{\ominus}} P(s)$ of $M$ the ascending and the descending cut-trees, respectively, of $\boldsymbol{M}$. Agarwal et al. [2] proved that $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$ are indeed both trees and include all maxima and minima, respectively, of $\boldsymbol{M}$ among their vertices. They also showed that they can be computed in $O(\operatorname{Sort}(N))$ I/Os.

Lemma 3.1 implies that a stretching $\boldsymbol{M}^{\prime}$ of $\boldsymbol{M}$ has exactly the same ascending and descending cuttrees as $\boldsymbol{M}$. In [2], $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$ are used to compute a specific level-ordering of a terrain $\boldsymbol{M}$ in $O(\operatorname{Sort}(N))$ I/Os. This, in conjunction with the level-ordering construction of [2] implies the following:

Corollary 4.1. Let $\boldsymbol{M}^{\prime}$ be a stretching of $\boldsymbol{M}$. A level-ordering on $\boldsymbol{M}$ can be computed in $O(\operatorname{Sort}(N))$


Figure 10: The effect of stretching on contour fragmentation

## $I / O s$ that is also a level-ordering on $\boldsymbol{M}^{\prime}$.

The specific construction of the level-ordering $\rho$ in [2] entails a characterization of the fragments of contours in relation to $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$. This is summarized in the following lemma. See [2] for details and proof.

Lemma 4.1. Let $K$ and $K^{\prime}$ be contours in $\boldsymbol{M}_{\ell}$, and let $\rho$ be the computed level-ordering of $\boldsymbol{M}$. Then a fragment of $K$ is immediately followed by a fragment of $K^{\prime}$ in $F_{\rho}\left(\boldsymbol{M}_{\ell}\right)$ if and only if $K$ and $K^{\prime}$ are connected by a polygonal path contained either in $\hat{\boldsymbol{T}} \cap \boldsymbol{M}_{\leq \ell}$ or in $\dot{\boldsymbol{T}} \cap \boldsymbol{M}_{\geq \ell}$.
4.1 The intuitive idea. The intuition as to why stretching can help us with fragmentations is as follows. Consider the cartoon terrain of Figure 10. On the left a terrain $\boldsymbol{M}$ with two hills is shown. The saddle between the two hills contributes a path to the ascending cut-tree which ends at the peaks of the two hills. The two contours $K$ and $K^{\prime}$ on the shown level set are connected by the ascending cut-tree through their corresponding sublevel set. By Lemma 4.1, this means that some fragment of each of these contours is followed by a fragment of the other in the levelordered list of triangles that intersect their level. Let us assume that $K$ is split into two fragments with $K^{\prime}$ in between them. On the right, $\boldsymbol{M}$ is stretched at one of the up-contours of the saddle so that the


Figure 11: Rooting of the ascending cut-tree in a terrain with only positive saddles.
piece of the terrain bounded by that up-contour (the portion that contains $K$ ) is lifted above the other peak. Let $\boldsymbol{M}^{\prime}$ be this stretched terrain. $K$ is no longer fragmented by $K^{\prime}$ in $\boldsymbol{M}^{\prime}$. Although a contour $K^{\prime \prime}$ in the shaded area of $\boldsymbol{M}^{\prime}$ is fragmented by $K^{\prime}$, no contour of $\boldsymbol{M}$ is mapped to $K^{\prime \prime}$, so the fragmentation of $K^{\prime \prime}$ is no concern of us.

Next, consider the case when $\boldsymbol{M}$ does not have any negative saddles., i.e., $\boldsymbol{M}$ has only one minimum $v_{\infty}$; see Figure 11. The descending cut-tree of this terrain is therefore trivial. Each saddle of the ascending cut-tree can be regarded as the midpoint of a hyper-edge of the cut-tree that connects the two maxima at which the ascending paths from that saddle end. This step compresses the ascending-cut tree to another tree on the maxima of $\boldsymbol{M}$. Let us pick an arbitrary maximum as the root and orient the hyper-edges toward it. Of the two ascending paths out of each saddle, one ends at the head and the other at the tail of the oriented hyper-edge. We call ascending paths correspondingly as tail or head ascending paths. It was proved in [2] that the restriction of the ascending cut-tree to the inside of each (red) contour is connected. This in conjunction with the fact that each maximum can be the tail of at most one hyper-edge (though it could be the head of many hyper-edges), implies that each contour of the terrain intersects at most one tail ascending path. Now, we can use our observation from Figure 10 as follows: At each saddle, the up-contour crossing the head ascending path is stretched high enough so that in the final terrain every points inside this up-contour ends up higher than every point in the inside of the other up-contour of the saddle.

In the resulting terrain, any contour can only connect through the sublevel set to another contour of its level via the ascending cut-tree only by following


Figure 12: Stretching at a red negative saddle
the at most one tail ascending path that crosses it. By Lemma 4.1, this amounts to at most one fragmentation per contour. Terrains with only negative saddles can be handled in a similar way. However, the problem becomes more complicated when both positive and negative saddles are present because in this case blue positive and red negative saddles appear; see Figure 12. The figure on the right demonstrates the idea behind dealing with a red negative saddle. Notice how the outer down contour of the saddle is stretched upward so that the entire pit of the volcano rises above all the points outside of the stretched contour. This eliminates the possibility of the red down-contour of the saddle to be fragmented by its blue down-contour (the one inside the volcano). Notice that in this case the saddle itself also moves up.
4.2 The stretching algorithm We now present our algorithm for stretching $\boldsymbol{M}$. The stretched terrain $\boldsymbol{M}^{\prime}$ is derived by stretching exactly one upor -down-contour of every saddle of $\boldsymbol{M}$. We thus characterize the particular up- or down-contour of each saddle $s$ of $\boldsymbol{M}$ that is stretched together with the offset parameter $\delta(s)$.

Our algorithm uses compressed versions $\hat{T}$ and $\check{T}$ of $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$ respectively. Specifically, we define $\hat{T}=(\hat{V}, \hat{E})$ and $\check{T}=(\dot{V}, \check{E})$ where $\hat{V}$ is the set of maxima and $\check{V}$ is the set of minima of $\boldsymbol{M}$. For $u, v \in \hat{V}, u v \in \hat{E}$ if there is a positive saddle $s$ for which $P(s)$ ends in $u$ and $v$, in which case we label the edge $u v$ with $s$. $\check{T}$ is defined symmetrically. Furthermore, we pick root vertices for $\hat{T}$ and $\check{T}$ and orient their edges toward their roots. For the root $\check{r}$ of $\check{T}$ we pick $v_{\infty}$ and the root $\hat{r}$ of $\hat{T}$ is chosen arbitrary among the maxima reachable from $v_{\infty}$ in the contour tree $\mathcal{C}=\mathcal{C}_{M}$ through a path of red edges. ${ }^{4}$ See Figure $13 ; \hat{T}$ and $\check{T}$ shown in dotted black.

[^4]

Figure 13: Illustration of $\alpha(s), \beta(s)$, head $(s)$, and $\operatorname{tail}(s)$ for two choices of a saddle $s$ on a contour tree. In each case, the paths from $s$ to head $(s)$ and $\operatorname{tail}(s)$ are highlighted. The edges of $\hat{T}$ and $\check{T}$ are drawn in dashed lines. Vertices of $\hat{T}$ are the maxima and those of $\check{T}$ are the minima.

For an edge $u v$ of $\hat{T}$ (resp. $\check{T}$ ) labeled with positive (resp. negative) saddle $s$, if $v$ is the parent of $u$ with respect to $\hat{r}$ (resp. $\check{r}$ ), then we set $\operatorname{tail}(s)=u$ and $\operatorname{head}(s)=v$. There is a one-toone correspondence between the positive saddles and the edges of $\hat{E}$, and there is a similar correspondence between negative saddles and $\check{E}$. For any saddle $s$, we distinguish two of the three neighbors of $s$ in $\mathcal{C}$ as $\alpha(s)$ and $\beta(s): \alpha(s)$ is the neighbor on the path in $\mathcal{C}$ from $s$ to head $(s)$ and $\beta(s)$ is the neighbor on the path from $s$ to tail(s) (See Figure 13) For any saddle $s$, let $n(s)$ denote the number of leaves in the subtree $\mathcal{C}^{\beta(s)}$ of $\mathcal{C}$ rooted at $\beta(s)$.

We are now ready to characterize the stretch associated to each saddle $s$ : the stretch associated with $s$ is applied to the up- or down-contour of $s$ carried by the edge $s \alpha(s)$ of $\mathcal{C}$. The amount of the stretch is given by

$$
\delta(s)= \begin{cases}H \cdot n(s) & \text { if } s \text { is red } \\ -H \cdot n(s) & \text { if } s \text { is blue }\end{cases}
$$

where $H=\max _{u, v \neq v_{\infty}}|h(u)-h(v)|$.
For any vertex $v$, let $\Pi(v)$ denote the set of vertices in the path from $v$ to the root $v_{\infty}$ in $\mathcal{C}$. In our algorithm, the stretches that affect the height of a vertex $v$ are precisely those associated with saddles
$s$ for which $s \alpha(s)$ is an edge in $\Pi(v)$. We use the shorthand $\Sigma(v)$ for the set of all such saddles on $\Pi(v)$ and write $\Sigma_{B}(v)$ and $\Sigma_{R}(v)$ to respectively represent the sets of blue and red saddles in $\Sigma(v)$. For any saddle $s$, the edge $s \alpha(s)$ has the same color as $s$ itself. Thus defining

$$
\begin{aligned}
\Delta(v) & =\sum_{s \in \Sigma(v)} \delta(s) \\
.1) & =H \cdot\left(\sum_{s \in \Sigma_{R}(v)} n(s)-\sum_{s \in \Sigma_{B}(v)} n(s)\right)
\end{aligned}
$$

one has $h^{\prime}(v)=h(v)+\Delta(v)$.
This completes the description of our stretching algorithm. We describe in Section 5 how the stretched terrain can be computed using $O(\operatorname{Sort}(N))$ I/Os.
4.3 Bounding fragmentations in $\boldsymbol{M}^{\prime}$ We now argue that if a contour $K_{M}(f, \ell)$ on $\boldsymbol{M}$ maps to $K_{M^{\prime}}\left(f, \ell^{\prime}\right)$ in $\boldsymbol{M}^{\prime}$, then $K_{M^{\prime}}\left(f, \ell^{\prime}\right)$ has $\mathrm{O}(1)$ fragments in $F_{\rho}\left(\boldsymbol{M}_{\ell^{\prime}}^{\prime}\right)$. To prove this, we need to mention some structural properties of contour trees and relationship between cut trees and contour trees.

Properties of contour trees. As earlier, let $\mathcal{C}=$ $\mathcal{C}_{\boldsymbol{M}}$ be the contour tree of $\boldsymbol{M}$, rooted at $v_{\infty}$. We also assume that vertices and edges of $\mathcal{C}$ are colored red or blue as specified in Section 2. Each vertex $v$ of $\mathcal{C}$ has one, two, or three neighbors in $\mathcal{C}$. A neighbor $u$ of $v$ is an up-neighbor if $h(u)>h(v)$ and a downneighbor if $h(u)<h(v)$. An edge of $\mathcal{C}$ connecting $v$ to an up-neighbor or down-neighbor of $v$ is referred to as an up-going or down-going edge of $v$, respectively. An up-going path from $u$ to $v$ in $\mathcal{C}$ is a path along the vertices of which the height function is increasing. A down-going path is defined similarly. If we orient the edge of $\mathcal{C}$ toward $v_{\infty}$ (the root of $\mathcal{C}$ ) all red edges will be oriented from the higher endpoint to the lower and all blue edges will be oriented from the lower endpoint to the higher (See Figure 14).

The leaves of $\mathcal{C}$ (vertices with only one neighbor) are precisely the extrema of $\boldsymbol{M}$. A maximum has one down-neighbor and a minimum has one up-neighbor. Each regular vertex has one up- and one downneighbor.

Any saddle $s$ of $\boldsymbol{M}$ has three neighbors in $\mathcal{C}$. If $s$ is positive it has two up-neighbors and one downneighbor. If $s$ is a red positive saddle, all the edges incident to $s$ are red and if it is a blue positive saddle, then one of its up-going edges is red and its other two incident edges are blue. The case of negative saddles are symmetric with 'red' exchanged with 'blue' and 'up' with 'down'.


Figure 14: Colored contour tree rooted at $v_{\infty}$; The red components are shaded. The subtree $\mathcal{C}^{v}$ rooted at the indicated vertex $v$ is circled by a dashed line.

The subtree of $\mathcal{C}$ rooted at a vertex $v$ is denoted by $\mathcal{C}^{v}$. Consider the forest $\mathcal{C}_{R} \subseteq \mathcal{C}$ induced by the red edges of $\mathcal{C}$. We call each connected component of $\mathcal{C}_{R}$ a red component of $\mathcal{C}$. We define the forest $\mathcal{C}_{B}$ induced by the blue edges of $\mathcal{C}$ and the corresponding blue components similarly. We call red or blue components the color components of $\mathcal{C}$. Observe that if vertex $v$ has a red down-going edge incident to it, then all its up-going edges will be red as well. Symmetrically, if a vertex has a blue up-going edge, all its down-going edges will be blue. This means that red components of $\mathcal{C}$ are up-ward closed, i.e., all vertices reachable from a red vertex through up-going paths are also red. Similarly, blue components are down-ward closed.

By definition any two red components of $\mathcal{C}$ are vertex disjoint and the same holds for any two blue components. A vertex shared between a red and a blue component must be either a red negative saddle or a blue positive saddle. Since $\mathcal{C}$ is acyclic, a red and a blue component of $\mathcal{C}$ can have at most one vertex in common, in which case we call them neighboring components in $\mathcal{C}$. Let $v$ be a vertex common between two color components $C$ and $C^{\prime}$ of opposite colors. Then either $C$ or $C^{\prime}$ (but not both) is a subgraph of $\mathcal{C}^{v}$. If $C$ is a subgraph of $\mathcal{C}^{v}$, we say $C$ is a child component of $C^{\prime}$ and $C^{\prime}$ is the parent component of $C$. Note that in this case $v$ is the closest vertex of $C$ to the root $v_{\infty}$ in $\mathcal{C}$ and is therefore called the root of the color component $C$.

Recall that $v_{\infty}$ is a red vertex. Therefore, it will
be contained in a red component which we call the root component of $\mathcal{C}$. The children of a red component $R$ are blue components connected to $R$ through red negative saddles and the children of a blue component $B$ are red components connected to $B$ through blue positive saddles. Thus the root of a blue component is a red negative saddle and the root of a red component is a blue positive saddle. Besides the root, all vertices of a blue component are blue and all vertices in a red component are red.

Cut trees and contour trees. We consider the images of $\hat{\boldsymbol{T}}$ or $\check{\boldsymbol{T}}$ under the quotient map $\pi_{M}$. Observe that monotone paths on $\boldsymbol{M}$ are mapped into $\mathcal{C}$ injectively. Let $s$ be a positive saddle. The vertices $\operatorname{head}(s)$ and tail $(s)$ of $\hat{T}$ are the maxima of $\boldsymbol{M}$ in which $P(s)$ ends. Let $P_{1}(s)$ and $P_{2}(s)$ respectively be the ascending paths from $s$ to head $(s)$ and $\operatorname{tail}(s)$ that together compose $P(s)$. Since $P_{1}(s)$ and $P_{2}(s)$ are ascending, their images $\mathcal{P}^{+}(s)=\pi_{M}\left(P_{1}(s)\right)$ and $\mathcal{P}^{-}(s)=\pi_{M}\left(P_{2}(s)\right)$ are up-going paths in $\mathcal{C}$, i.e., the heights of vertices along these paths in $\mathcal{C}$ monotonically increase. Similarly, for a negative saddle $s$, we get down-going paths $\mathcal{P}^{+}(s)$ and $\mathcal{P}^{-}(s)$ in $\mathcal{C}$ from $s$ to head $(s)$ and tail $(s)$ respectively. For any saddle $s$, let $\mathcal{P}(s)=\mathcal{P}^{+}(s) \cup \mathcal{P}^{-}(s)$. Thus the image of the entire $\hat{\boldsymbol{T}}$ under the map $\pi_{M}$ is a subtree $\hat{\mathcal{T}}=\bigcup_{s \in V_{\oplus}} \mathcal{P}(s)$ of $\mathcal{C} . \check{\mathcal{T}}$ is defined similarly.

Stretching of saddles. Next, we bound relative heights of saddles in the stretched terrain $\boldsymbol{M}^{\prime}$, which in turn will bound the number of fragments in a contour in $\boldsymbol{M}^{\prime}$. We need the following two lemmas to prove Lemma 4.4, which states the main property of the saddles in $\boldsymbol{M}^{\prime}$.

Lemma 4.2. For any blue positive saddle $s, \alpha(s)$ and $\beta(s)$ are the up-neighbors of $s$ respectively reachable through the blue and red up-going edges at s. For any red ngative saddle $s, \alpha(s)$ and $\beta(s)$ are the downneighbors of $s$ respectively reachable through the red and blue down-going edges at $s$.

Proof. We prove the lemma for blue positive saddles; the other case is symmetric. Let $s$ be a blue positive saddle and let $u_{R}$ and $u_{B}$ be the up-neighbors of $s$ respectively reachable through the red and the blue up-going edges of $s$. Let $R$ and $B$ respectively be the red and blue components of $\mathcal{C}$ that contain the edges $s u_{R}$ and $s u_{B}$. Thus $s$ is the vertex shared between $R$ and $B$. Since $s$ is a blue positive saddle, $B$ is the parent and $R$ is the child component. As discussed above, one endpoint $v_{B}$ of the $\hat{T}$ edge labeled $s$ is in $\mathcal{C}^{u_{B}}$ and the other $v_{R}$ is in $\mathcal{C}^{u_{R}}$. We claim that $v_{R}=\operatorname{tail}(s)$ and $v_{B}=\operatorname{head}(s)$. To see this, we
observe that no ascending path initiated in $\mathcal{C}^{u_{R}}$ can reach $s$ since the only way for such a path to reach $s$ is through $u_{R}$ which is an up-neighbor of $s$. Thus the two endpoints of the $\hat{T}$ edge labeled $s^{\prime}$ for any $s^{\prime} \in \mathcal{C}^{u_{R}}$ are contained in $\mathcal{C}^{u_{R}}$. Since $\hat{r}$ (the root of $\hat{T}$ ) is in the root red component of $\mathcal{C}, v_{R}$ can reach $\hat{r}$ only through the parent component $B$ of $R$ and since $\hat{T}$ edge labeled $s$ is the only edge of $\hat{T}$ that has exactly one endpoint in $\mathcal{C}^{u_{R}}, v_{B}$ has to be head $(s)$. This means that $u_{B}=\alpha(s)$ and $u_{R}=\beta(s)$.

Lemma 4.3. Let $s$ and $s^{\prime}$ be two red (resp. blue) saddles so that $s, \alpha(s), s^{\prime}$, and $\alpha\left(s^{\prime}\right)$ are all on the same up-going (resp. down-going) path $\Pi_{0}$ made of only red (resp. blue) edges. Then $\mathcal{C}^{\beta(s)}$ and $\mathfrak{C}^{\beta\left(s^{\prime}\right)}$ are disjoint subtrees of $\mathcal{C}$.

Proof. We only prove the red version of the lemma. Let $R$ be the red component of $\mathcal{C}$ that contains $\Pi_{0}$ and let $r$ be the root of $R$. Being the root of a red component of $\mathcal{C}_{\boldsymbol{M}}, r$ is either $v_{\infty}$ or is a blue positive saddle. The path $\Pi_{0}$ can be extended to an up-going path $\Pi_{1}$ from $r$. For any vertex $v \in R, \Pi(v)$ must pass through $r$. Now, depending on $s$ being positive or negative, either both $\alpha(s)$ and $\beta(s)$ are up-neighbors of $s$ or they are both down-neighbors of $s$. Thus at most one of them can be present on $\Pi_{1}$. Since $\alpha(s)$ is on $\Pi_{1}, \beta(s)$ cannot be on $\Pi_{1}$. Similarly, $\beta\left(s^{\prime}\right)$ cannot be on $\Pi_{1}$. Since $\beta(s)$ is a neighbor of $s$ and $s$ is on a path $\Pi_{1}$ to $r, \Pi(\beta(s))$ has to go through $s$. Similarly, $\Pi\left(\beta\left(s^{\prime}\right)\right)$ goes through $s^{\prime}$. But this means that neither of $\beta(s)$ and $\beta\left(s^{\prime}\right)$ can be on the path from the other to $r$.

Lemma 4.4. For any positive (resp. negative) saddle $s$, and for any vertices $u \in \mathcal{P}^{-}(s)$ and $v \in \mathcal{P}^{+}(s)$ with $v \neq s, h^{\prime}(v)>h^{\prime}(u)\left(\right.$ resp. $\left.h^{\prime}(v)<h^{\prime}(u)\right)$.

Proof. We prove the lemma for positive saddles. The proof for negative saddles is symmetric. There are two cases depending on the color of $s$.
Blue saddle. We first consider the case in which $s$ is a blue positive saddle. Thus $s$ is the root of a red component $R$ of $\mathcal{C}$ whose parent $B$ is a blue component. We name the up-neighbors of $s$ in $B$ and $R$ as $u$ and $u^{\prime}$ respectively. By Lemma 4.2, $u=\alpha(s)$ and $u^{\prime}=\beta(s)$. Let $v=\operatorname{head}(s)$ and $v^{\prime}=\operatorname{tail}(s)$ as determined by $\hat{T}$ (See Figure 15). Since $R$ is upward closed, $v^{\prime}$ is a vertex of $R$ and the path $\Pi\left(v^{\prime}\right)$ passes through $s$ (the root of $R$ ) and $u$ which is the parent of $s$ in $\mathcal{C}$. Since $\mathcal{C}$ is the contour tree of both $\boldsymbol{M}$ and $\boldsymbol{M}^{\prime}$ and since $\mathcal{P}^{+}(s)$ and $\mathcal{P}^{-}(s)$ are up-going paths in $\mathcal{C}$, both $h$ and $h^{\prime}$ are increasing along both of these paths. It is therefore sufficient to prove that $h^{\prime}\left(v^{\prime}\right)<h^{\prime}(u)$. Since $u$ is


Figure 15: Illustrations for the proof of Lemma 4.4; the blue positive saddle case
a vertex of $\Pi\left(v^{\prime}\right), \Sigma(u) \subseteq \Sigma\left(v^{\prime}\right)$ and therefore by Eq (4.1) any saddle $s^{\prime} \in \Sigma(u)$ contributes the same amount to both $\Delta(u)$ and $\Delta\left(v^{\prime}\right)$. Therefore it is enough to show that $D=\sum_{s^{\prime} \in S_{0}} \delta\left(s^{\prime}\right)<0$ where $S_{0}=\Sigma\left(v^{\prime}\right) \backslash \Sigma(u)$. Since $D$ is an integer multiple of $H$, if negative, it must be smaller than or equal to $-H$. The statement of the lemma then follows from the fact that $h(u)>h\left(v^{\prime}\right)-H$.

Observe now that the only blue saddle in $S_{0}$ is $s$, and since $\delta$ is negative for blue saddles and positive for red ones, letting $S_{1}=S_{0} \backslash\{s\}$, we need to show that

$$
\begin{equation*}
\sum_{s^{\prime} \in S_{1}} n\left(s^{\prime}\right)<n(s) \tag{4.2}
\end{equation*}
$$

Recall that for any saddle $s, n(s)$ is the number of leaves in $\mathfrak{C}^{\beta(s)}$. Since $\beta(s)=u^{\prime}$ which is on the path from any vertex in $R$ to the root $r$ of $R$, for any $s^{\prime} \in S_{1}, \mathcal{C}^{\beta\left(s^{\prime}\right)}$ is a subtree of $\mathcal{C}^{\beta(s)}$. On the other hand, the saddles in $S_{1}$ all lie on an up-going red path, namely $\mathcal{P}^{-}(s)$, from $s$ to $v_{R}$. Therefore by Lemma 4.3 for any two saddles $s^{\prime}, s^{\prime \prime} \in S_{1}, \mathfrak{C}^{\beta\left(s^{\prime}\right)}$ and $\mathcal{C}^{\beta\left(s^{\prime \prime}\right)}$ are disjoint. This means that the total number of leaves in $\mathcal{C}^{\beta\left(s^{\prime}\right)}$ over all saddles $s^{\prime} \in S_{1}$ are at most equal to the number of leaves in $\mathcal{C}^{\beta(s)}$ and therefore $\sum_{s^{\prime} \in S_{1}} n\left(s^{\prime}\right) \leq n(s)$. The strictness of inequality in (4.2) follows from the observation that $v_{R} \neq \operatorname{tail}\left(s^{\prime}\right)$ for any $s^{\prime} \in S_{1}$, since $v_{R}=\operatorname{tail}(s)$, and therefore is contained in $\mathcal{C}^{\beta\left(s^{\prime}\right)}$ for no $s^{\prime} \in S_{1}$.
Red saddle. We now assume $s$ to be a red positive saddle. As in the previous case, let $u=\alpha(s), u^{\prime}=$ $\beta(s), v=\operatorname{head}(s)$ and $v^{\prime}=\operatorname{tail}(s)$ (See Figure 16).


Figure 16: Illustrations for the proof of Lemma 4.4; the red positive saddle case

Both $u$ and $u^{\prime}$ are up-neighbors of $s$ reached through up-going red edges of $s$. Again to prove the lemma it is enough to show that $h^{\prime}\left(v^{\prime}\right)<h^{\prime}(u)$. Notice that this time the parent of $s$ in $\mathcal{C}$ is its down-neighbor $w$. Thus the paths $\Pi\left(v^{\prime}\right)$ and $\Pi(u)$ both pass through $s$. In other words, all saddles in $\Pi(s)$ contribute equally to $\Delta(u)$ and $\Delta\left(v^{\prime}\right)$. Since $\alpha(s)=u \notin \Pi(s), s \notin \Sigma(s)$ while $s \in \Sigma(u)$ and therefore $s$ is the only saddle that contributes to $\Delta(u)$ but not to $\Delta\left(v^{\prime}\right)$. Let $S_{0}$ be the saddles on $\Pi\left(v^{\prime}\right)$ between $v^{\prime}$ and $s$, i.e., on $\mathcal{P}^{-}(s)$, that are contained in $\Sigma\left(v^{\prime}\right)$. Note that since $s$ is a red positive saddles, all the saddles on $\mathcal{P}^{-}(s)$ are red and therefore $\delta\left(s^{\prime}\right)>0$ for any $s^{\prime} \in S_{0}$. Since $s$ is a red saddle as well, $\delta(s)$ is also positive. Therefore, as in the previous case, to prove the lemma we show that

$$
n(s)>\sum_{s^{\prime} \in S_{0}} n\left(s^{\prime}\right)
$$

Since saddles in $S_{0}$ are on an up-going red path, i.e., $\mathcal{P}^{-}(s)$, Lemma 4.3 implies that for any $s^{\prime}, s^{\prime \prime} \in S_{0}$, $\mathcal{C}^{\beta\left(s^{\prime}\right)}$ and $\mathcal{C}^{\beta\left(s^{\prime \prime}\right)}$ are disjoint. On the other hand, $u^{\prime}=\beta(s)$ and therefore $\mathcal{C}^{\beta\left(s^{\prime}\right)}$ is a subtree of $\mathcal{C}^{\beta(s)}$ for any $s^{\prime} \in S_{0}$. This means that $n(s) \geq \sum_{s^{\prime} \in S_{0}} n\left(s^{\prime}\right)$. Again, the inequality follows from observing that $v^{\prime}$ cannot be in $\beta\left(s^{\prime}\right)$ for any $s^{\prime} \in S_{0}$ and therefore only contributes to $n(s)$.

Corollary 4.2. Let $s$ and $s^{\prime}$ be distinct positive (resp. negative) saddles. If head( $s^{\prime}$ ) is a descendent of tail(s) in $\hat{T}$ (resp. $\check{T}$ ), then for any vertex $u \in$ $\mathcal{P}^{-}\left(s^{\prime}\right) \cup \mathcal{P}^{+}\left(s^{\prime}\right)$ and any vertex $v \in \mathcal{P}^{+}(s)$ with $v \neq s$, $h^{\prime}(u)<h^{\prime}(v)\left(\right.$ resp. $\left.h^{\prime}(u)>h^{\prime}(v)\right)$.

Proof. Since head $\left(s^{\prime}\right)$ is a descendent of tail $(s)$, there is a path from tail $\left(s^{\prime}\right)$ to head $(s)$ in $\hat{T}$. Let $s^{\prime}=s_{0}, s_{1}, \ldots, s_{k}=s$ be labels of the edges on this path. For each $i=0, \ldots, k$, let $u_{i}=\operatorname{tail}\left(s_{i}\right)$ and $v_{i}=\operatorname{head}\left(s_{i}\right)$. Since edges labeled $s_{0}, \ldots, s_{k}$ make a path in $\hat{T}, u_{i}=v_{i-1}$ for each $i=1, \ldots, k$. By Lemma 4.4, $h^{\prime}\left(v_{i}\right)>h^{\prime}\left(u_{i}\right)$ for all $i=1, \ldots, k-1$ and by the same lemma, $h^{\prime}(u)<h^{\prime}\left(v_{0}\right)$ and $h^{\prime}\left(v_{k-1}\right)<h^{\prime}(v)$. Thus $h^{\prime}(u)<h^{\prime}(v)$. The proof for negative saddles is similar.

Remark. Lemma 4.4 ensures that in $\boldsymbol{M}^{\prime}$ a contour $K$ carried by an edge in $\mathcal{P}^{+}(s)$ of a saddle $s$ is not fragmented by a contours carried by edges of $\mathcal{P}^{-}(s)$. Corollary 4.2 extends this by saying that $K$ is not fragmented by contours carried by edges of $\mathcal{P}\left(s^{\prime}\right)$ for any saddle $s^{\prime}$ for which head $\left(s^{\prime}\right)$ is a descendent of $\operatorname{tail}(s)$ in $\hat{T}$ or $\check{T}$.

Bounding fragmentation. We are now ready to bound the number of fragments in the up- or downcontour of any vertex in $\boldsymbol{M}^{\prime}$.

ThEOREM 4.1. Let $K$ be the red down-contour (resp. blue up-contour) of a vertex $v$ in $\boldsymbol{M}^{\prime}$ and let $\ell=$ $h^{\prime}(v)-\varepsilon$ (resp. $\left.\ell=h^{\prime}(v)+\varepsilon\right)$. Then $K$ has at most three fragments in the level-ordered list of triangles in $F\left(\boldsymbol{M}_{\ell^{\prime}}^{\prime}\right)$.

Proof. Let $K$ be a red down counter of $v$ in $\boldsymbol{M}_{\ell}^{\prime}$ and let $K^{\prime}$ be another contour in $\boldsymbol{M}_{\ell}^{\prime}$ such that fragments from $K$ and $K^{\prime}$ appear consecutively in $F_{\rho}\left(\boldsymbol{M}_{\ell}^{\prime}\right)$. By Lemma 4.1 this means that there is either a path in $\hat{\boldsymbol{T}} \cap \boldsymbol{M}_{\leq \ell}^{\prime}$ or a path in $\check{\boldsymbol{T}} \cap \boldsymbol{M}_{\geq \ell}^{\prime}$ that connects $K$ and $K^{\prime}$. Assume that a path $Q$ of the former type exists. Projecting $Q$ under the map $\pi_{M^{\prime}}$ into $\mathcal{C}_{M^{\prime}}$, gives us a walk on $\mathcal{C}_{M^{\prime}}$ between the points representing $K$ and $K^{\prime}$ (both points on edges of $\mathcal{C}_{M}^{\prime}$ ) that stays below level $\ell$ throughout. In particular, the points representing $K$ and $K^{\prime}$ must be connected by a path in $\hat{\mathcal{T}}$ in which all vertices have heights less than $\ell$.

Consider any positive saddle $s$ for which $\mathcal{P}^{+}(s)$ passes through $v$ (and therefore the point representing $K$ ) in $\mathcal{C}_{\boldsymbol{M}}^{\prime}$. By Lemma 4.4 any contour of $\boldsymbol{M}^{\prime}$ carried by the edges of $\mathcal{P}^{-}(s)$ belongs to a level strictly smaller than $\ell$. Furthermore, by Corollary 4.2, any contour carried by an edge in $\mathcal{P}\left(s^{\prime}\right)$, where $s^{\prime}$ is a positive saddle for which head $\left(s^{\prime}\right)$ is a descendent of $\operatorname{tail}(s)$ in $\hat{T}$, also belongs to a level strictly lower than $\ell$. Therefore, $K^{\prime}$ can be no such contour. By definition, $v$ can be contained in $\mathcal{P}^{-}\left(s_{0}\right)$ for only one positive saddle $s_{0}$. This is because tail $\left(s_{0}\right)$ has outdegree at most one in $\hat{T}$. Let $u$ be the neighbor of $s_{0}$ in $\mathcal{P}^{+}\left(s_{0}\right)$, i.e., $u=\alpha\left(s_{0}\right)$. By Lemma $4.4 h^{\prime}(u)>\ell$. Thus $\pi_{M^{\prime}}(Q)$ cannot include $u$ or any other vertex
of $\hat{\mathcal{T}}$ reachable through $u$. Therefore, $K^{\prime}$ can only be carried by the edge $s_{0} u$ of $\mathcal{C}_{M}^{\prime}$. Thus there can be at most one such contour $K^{\prime}$. Similarly, there is at most one contour $K^{\prime \prime}$ of the latter type (reachable from $K$ by a path in $\check{\boldsymbol{T}} \cap \boldsymbol{M}_{\geq l}^{\prime}$. This bounds the number of fragments of $K$ to no more than three.

Corollary 4.3. For a contour represented by $(f, \ell)$ on $\boldsymbol{M}$, let $\left(f, \ell^{\prime}\right)$ be the pair returned by the algorithm TranslateQuery. Then $K_{M^{\prime}}\left(f, \ell^{\prime}\right)$ has at most three fragments in $F_{\rho}\left(\boldsymbol{M}_{\ell^{\prime}}^{\prime}\right)$.

## 5 I/O-Efficient Implementation

In this section we describe how the stretched terrain $\boldsymbol{M}^{\prime}$ can be constructed I/O-efficiently and how the data structure in [2] can be adapted to answer contour queries on $\boldsymbol{M}^{\prime}$ I/O-efficiently. We rely on a few standard I/O-efficient techniques such as list ranking, Euler tour computation of a tree, and time-forward processing. We do not describe these techniques in detail and refer the reader to [18].
5.1 Algorithms for computing the stretching Here we explain how the individual steps of the algorithm of Section 4 can be implemented using $O(\operatorname{Sort}(N))$ I/Os.

Computing the colored contour tree. The algorithm of Agarwal et al. [3] computes the augmented contour tree $\mathcal{C}$ of $\boldsymbol{M}$ in $O(\operatorname{Sort}(N))$ I/Os. If $\mathcal{C}$ is rooted, i.e., if the parent of each vertex with respect to the root $v_{\infty}$ of $\mathcal{C}$ is determined, finding the colors of the edges of $\mathcal{C}$ can be achieved through a simple scan of the edges of $\mathcal{C}$. This is because an upward edge between a vertex $v$ and a child of $v$ in $\mathcal{C}$ is always red and a downward edge between $v$ and a child of $v$ is always blue. Rooting $\mathcal{C}$ can be achieved in $O(\operatorname{Sort}(N))$ I/Os as described above. Once the color of the edges incident to a vertex are determined the color of the vertex itself follows the local rules described in Section 2.

Computing and rooting $\hat{T}$ and $\check{T}$. Agarwal et al. [2] have described an algorithm that computes the cut trees $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$ in $O(\operatorname{Sort}(N))$ I/Os, so we only describe how to compute $\hat{T}$ and $\check{T}$ from $\hat{\boldsymbol{T}}$ and $\check{\boldsymbol{T}}$, respectively. First note that, using the algorithm in [3], we can compute in $O(\operatorname{Sort}(N))$ I/Os the type of each vertex in $\boldsymbol{M}$, and also the counterclockwise first vertex of each connected component of the upper link of every vertex.

Recall that $\hat{T}=(\hat{V}, \hat{E})$ where $\hat{V}$ is the set of maxima in $\boldsymbol{M}$ and $u v \in \hat{E}$ if there is a positive saddle $s$ for which the path $P(s)$ ends in $u$ and $v$
(see Section 4 for details). $\hat{T}$ can be constructed from $\hat{\boldsymbol{T}}$ using the time-forward-processing technique as follows. For the purpose of this algorithm, we assume that the edges of $\hat{\boldsymbol{T}}$ are oriented in decreasing order of the heights of their endpoints. We scan the vertices of $\hat{\boldsymbol{T}}$ in the decreasing order of their heights and compute a labeling $\varphi$ of each edge of $\hat{\boldsymbol{T}}$ that will be used to construct $\hat{T}$. During the scan we maintain the set of edges whose upper endpoints have been visited in an external priority queue $Q$; the priority of an edge is the ID of its lower endpoint.

Suppose we reach a vertex $v$ of $\hat{\boldsymbol{T}}$. We first delete all incoming edges at $v$, i.e., the edges for which $v$ is the lower endpoint, using the delete-min operation repeatedly. If $v$ is a maximum (in which case there are no incoming edges at $v$ ), then we label each outgoing edge $e$ of $\hat{\boldsymbol{T}}$ at $v$ with $\varphi(e)=v$. If $v$ is a regular vertex or a negative saddle, then there is exactly one incoming edge $e^{-}$at $v$, and we label each outgoing edge $e$ at $v$ with $\varphi(e)=\varphi\left(e^{-}\right)$. If $v$ is a positive saddle, then there are two incoming edges $e_{1}$ and $e_{2}$ at $v$. We report $\varphi\left(e_{1}\right) \varphi\left(e_{2}\right)$ as an edge of $\hat{T}$. For each outgoing edge $e$ at $v$, we label $e$ with $\varphi\left(e_{1}\right)$ or $\varphi\left(e_{2}\right)$ depending on whichever is $\lambda^{+}(e)$. Finally, we insert all outgoing edges at $v$ into $Q$. Omitting further details, we note that this procedure requires $O(\operatorname{Sort}(N))$ I/Os. Next, we pick a maximum $\hat{r}$ that is reachable from $v_{\infty}$ through an up-going path in the contour tree $\mathcal{C}$. This can be done in $O(\operatorname{Sort}(N))$ I/Os using a similar approach - a label from $v_{\infty}$ is propagated upward in $\mathcal{C}$ until it is picked up by some maximum which is appointed $\hat{r}$. Once $\hat{r}$ is chosen, the edges of $\hat{T}$ can be oriented toward $\hat{r}$ as described above for the contour tree.

Similarly $\check{T}$ can be computed and rooted at $v_{\infty}$ in $O(\operatorname{Sort}(N))$ I/Os.

Computing $\alpha(s)$ and $\beta(s)$ for a each $s$. We perform an Euler tour on $\mathcal{C}$, initiated at $v_{\infty}$, and record for each vertex $v$ the steps of the tour in which $v$ is visited. The number of steps recorded for a vertex $v$ is its degree in $\mathcal{C}$. Thus extrema get a single number while saddles record three numbers. Let $s$ be a saddle in $\mathcal{C}$ and let $u$ and $v$ be its two children. Assume that the Euler tour reaches $u$ before $v$ for the first time. Since the tour is started at $v_{\infty}$, if the three numbers recorded at a saddles $s$, i.e., the steps of the tours in which $s$ is visited, are $a_{1}<a_{2}<a_{3}$, then all the step numbers associated to all vertices in $\mathcal{C}^{u}$ are between $a_{1}$ and $a_{2}$ and those associated to vertices in $\mathcal{C}^{v}$ are between $a_{2}$ and $a_{3}$. All other vertices of $\mathcal{C}$ are visited in steps before $a_{1}$ or in steps after $a_{3}$. Thus if $s$ is, say a positive saddle, one can determine the neighbors $\alpha(s)$ and $\beta(s)$ on the paths from $s$ to
$\operatorname{head}(s)$ and tail $(s)$ by comparing $a_{1}, a_{2}$, and $a_{3}$ to the steps in which each of $u$ and $v$ are visited. This can therefore be done in $O(\operatorname{Sort}(N))$ I/Os.

Computing $n(s)$ for each saddle. One can compute in $O(\operatorname{Sort}(N))$ I/Os the size of all subtrees rooted at all vertices of a rooted tree of size $N$. This will be enough to determine $n(s)$ for all saddles as $n(s)$ is by definition the number of leaves in $\mathcal{C}^{\beta(s)}$.

Computing $\boldsymbol{M}^{\prime}$. With $n(s)$ and the color of the saddles known, $\delta(s)$ is determined for each saddle. To compute $\boldsymbol{M}^{\prime}$ we perform another Euler tour initiated at $v_{\infty}$. As explained in Section 3, the stretching associated with a saddle $s$ is applied to all vertices in the subtree of $\mathcal{C}$ rooted at the child endpoint of the edge $s \alpha(s)$. In the Euler tour, we maintain the total change in height applicable to the vertices being visited in a variable $L$, initially set to zero. For any saddle $s$, the edge $s \alpha(s)$ is traversed twice in the Euler tour. The first traversal moves from the parent endpoint of the edge to the child endpoint at which time we add $\delta(s)$ to $L$. When $s \alpha(s)$ is visited for the second time, it will be in the direction from its child endpoint to its parent endpoint in $\mathcal{C}$. We thus subtract $\delta(s)$ from $L$. Thus the value of $L$ when visiting a vertex $v$ is precisely $\Delta(v)$. To compute $\boldsymbol{M}^{\prime}$, all one needs to do is to add $L$ to $h(v)$ when it is visited for the first time.
5.2 Answering contour queries We now briefly review the level-set query algorithm of Agarwal et al. [2] and explain how one can answer contour queries using their structure. We refer to the original paper for further details.


Figure 17: Locating a triangle $f$ in the appropriate version of the $B$-tree $\mathcal{B}$ and scanning the leafs for the extraction of the contour $K$ containing $f$. Fragments of $K$ are shown in grey. At the end of a fragment, one locates the subsequent fragment by searching the tree for the rank of the neighbor $g^{\prime}$ of the last triangle $g$ of the previous fragment.

Level-set queries. Let $v_{1}, \ldots, v_{N}$ be the vertices of $\boldsymbol{M}$ in the order of height, i.e., $h\left(v_{1}\right)<\cdots<$
$h\left(v_{N}\right)$. For simplicity we assume that all queries are performed at heights $\ell \notin h(V)$. A level-set query at height $\ell$ returns $\bar{F}[K]$ for every contour $K$ in $\boldsymbol{M}_{\ell}$. The resulting collection of lists is exactly the same for two level-set queries at levels $\ell_{1}, \ell_{2}$ with $h\left(v_{i}\right)<\ell_{1}, \ell_{2}<h\left(v_{i+1}\right)$. It is therefore sufficient for our algorithm to generate $M[K]$ for each contour $K$ of $M_{h\left(v_{i}\right)+\varepsilon}$ for $i=1, \ldots, N$; here $\varepsilon$ is a sufficiently small constant as in the definition of up- and downcontours.

The overall data structure is a persistent B-tree $\mathbb{B}$ that holds $N$ versions, each corresponding to one vertex of $\boldsymbol{M}$. The version $\mathbb{B}_{i}$ associated with $v_{i}$ is a B-tree that stores the sequence $F_{\rho}\left(\boldsymbol{M}_{h\left(v_{i}\right)+\varepsilon}\right)$. The search key for a triangle $f$ in any version of $\mathbb{B}$ is its rank $\rho(f)$ as determined by the Agarwal et al. [2] level-ordering $\rho$. To construct $\mathbb{B}$, we scan the list $v_{1}, \ldots, v_{N}$. On arriving at $v_{i}$, we first delete from $\mathbb{B}$ all triangles for which $v_{i}$ is the highest vertex and then insert all those for which $v_{i}$ is the lowest vertex. The resulting version of $\mathbb{B}$ is time-stamped $h\left(v_{i}\right) . \mathbb{B}$ can be constructed in $O(\operatorname{Sort}(N))$ I/Os and uses $O(N)$ space.

For any $\ell$, let $\mathbb{B}(\ell)$ be the version $\mathbb{B}_{i}$ of $\mathbb{B}$ corresponding to the highest vertex $v_{i}$ with $h\left(v_{i}\right)<\ell$. $\mathbb{B}(\ell)$ can be computed in $O\left(\log _{B} N\right) \mathrm{I} / \mathrm{Os}$ and its entire content, i.e., $F_{\rho}\left(\boldsymbol{M}_{\ell}\right)$, can be reported, ordered by $\rho$, in $O(T / B)$ I/Os, where $T=\left|F\left(\boldsymbol{M}_{\ell}\right)\right|$. Property C 2 of level-ordering allows us to separate the various contours in this sequence using an I/O-efficient stack in $O(T / B)$ extra I/Os [2].

Contour queries. $\mathbb{B}$ can be adapted to answer contour queries, albeit with possible compromise on the I/O complexity by retrieving various fragments of a contour from the relevant version: Given a query triangle $f$ an a level $\ell$, to extract the contour $K$ in $\boldsymbol{M}_{\ell}$ that intersects $f$, as before one first locates $\mathbb{B}(\ell)$ in $O\left(\log _{B} N\right)$ I/Os. Assuming $\mathbb{B}(\ell)$ stores $T_{\ell}$ elements, the leaf of $\mathbb{B}(\ell)$ that stores $f$ is found in $O\left(\log _{B} T_{\ell} / B\right) \mathrm{I} / \mathrm{Os}$ by searching in $\mathbb{B}(\ell)$ using $\rho(f)$ as the search key. The leaves of $\mathbb{B}(\ell)$ are then scanned sequentially from $f$ to the left and right to find the remaining triangles in $K$. Since $K$ may be fragmented in $\mathbb{B}(\ell)$, one may reach the end of the fragment that includes $f$ at a triangles $g$ without having found all of $\bar{F}[K]$ (See Figure 17). If one stores at each triangle of $M$ the rank of its three neighboring triangles, the fragment of $\bar{F}[K]$ that starts with the neighbor $g^{\prime}$ of $g$ can be located by searching $\mathbb{B}(\ell)$ using $\rho\left(g^{\prime}\right)$ as the search key; this step takes in $O\left(\log _{M / B} T_{\ell} / B\right)$ I/Os. Continuing this way, a contour $K$ of length $T$ broken into $J$ fragments
is reported in $O\left(\log _{B} N+J \log _{M / B} T_{\ell} / B+T / B\right)$ I/Os.

If contour queries are issued on the stretched terrain obtained using the algorithm of Section 4, the number $J$ of fragments is constant (Corollary 4.3). Therefore, we obtain the main result of the paper:

Theorem 5.1. A terrain $\boldsymbol{M}$ of size $N$ can be preprocessed, in $O(\operatorname{Sort}(N)) I / O s$, into a data structure of size $O(N)$ so that a contour query $(f, \ell)$ on $\boldsymbol{M}$ can be answered in $O\left(\log _{B} N+T / B\right) I / O s$, where $T$ is the size of the query output.

## References

[1] P. K. Agarwal, L. Arge, T. M. Murali, K. Varadarajan, and J. S. Vitter. I/O-Efficient Algorithms for Contour Line Extraction and Planar Graph Blocking. In Proc. ACM-SIAM Sympos. Discrete Algorithms, pages 117-126, 1998.
[2] Pankaj K. Agarwal, Lars Arge, Thomas Mølhave, and Bardia Sadri. I/O-Efficient Algorithms for Computing Contours on a Terrain. In $S C G$ ' 08 : Proceedings of the twenty-fourth annual symposium on Computational geometry, pages 129-138, New York, NY, USA, 2008. ACM.
[3] Pankaj K. Agarwal, Lars Arge, and Ke Yi. I/OEfficient Batched Union-Find and its Applications to Terrain Analysis. In SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry, pages 167-176, New York, NY, USA, 2006. ACM.
[4] A. Aggarwal and J. S. Vitter. The input/output complexity of sorting and related problems. Commun. ACM, 31:1116-1127, 1988.
[5] L. Arge. External memory data structures. In J. Abello, P. M. Pardalos, and M. G. C. Resende, editors, Handbook of Massive Data Sets, pages 313358. Kluwer Academic Publishers, 2002.
[6] L. Arge, A. Danner, and S-H. Teh. I/O-efficient point location using persistent B-trees. In Proc. Workshop on Algorithm Engineering and Experimentation, 2003.
[7] Surender Baswana and Sandeep Sen. Planar graph blocking for external searching. Algorithmica, 34(3):298-308, 2002.
[8] Hamish Carr, Jack Snoeyink, and Ulrike Axen. Computing contour trees in all dimensions. Comput. Geom., 24(2):75-94, 2003.
[9] Y.-J. Chiang and C. T. Silva. I/O-Optimal Isosurface Extraction. In Proc. IEEE Visualization, pages 293-300, 1997.
[10] Yi-Jen Chiang, Michael T. Goodrich, Edward F. Grove, Roberto Tamassia, Darren Erik Vengroff, and Jeffrey Scott Vitter. External-memory graph algorithms. In SODA '95: Proceedings of the
sixth annual ACM-SIAM symposium on Discrete algorithms, pages 139-149, 1995.
[11] Kree Cole-McLaughlin, Herbert Edelsbrunner, John Harer, Vijay Natarajan, and Valerio Pascucci. Loops in Reeb graphs of 2-manifolds. Discrete $\mathcal{E}$ Computational Geometry, 32(2):231-244, 2004.
[12] A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge, and H. Mitásová. Terrastream: From elevation data to watershed hierarchies. In Proc. ACM Symposium on Advances in Geographic Information Systems, page 28, 2007.
[13] Herbert Edelsbrunner, John Harer, and Afra Zomorodian. Hierarchical morse - smale complexes for piecewise linear 2-manifolds. Discrete \& Computational Geometry, 30(1):87-107, 2003.
[14] C. Hutton. An account of the calculations made from the survey and measures taken at schehallien, in order to ascertain the mean density of the earth. Philosophical Transactions of the Royal Society of London, 68:689-845, 1779.
[15] M. H. Nodine, M. T. Goodrich, and J. S. Vitter. Blocking for external graph searching. Algorithmica, 16(2):181-214, 1996.
[16] P. J. Varman and R. M. Verma. An efficient multiversion access structure. IEEE Transactions on Knowledge and Data Engineering, 9(3):391-409, 1997.
[17] J. S. Vitter. External memory algorithms and data structures. In External memory algorithms, pages 1-38. American Mathematical Society, Boston, MA, USA, 1999.
[18] Norbert Zeh. I/O-efficient graph algorithms. In EEF Summer School on Massive Data Sets, to appear.


[^0]:    *Duke University, Durham, NC. Supported by NSF under grants CNS-05-40347, CCF-06-35000, IIS-07-13498, and CCF-09-40671, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, 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.
    $\dagger$ Duke University, Durham, NC. Supported by NSF under grants CNS-05-40347, CCF-06-35000, IIS-07-13498, and CCF-09-40671, by ARO grants W911NF-07-1-0376 and W911NF-$08-1-0452$, 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.
    $\ddagger$ University of Toronto, Toronto, ON, Canada

[^1]:    ${ }^{1}$ If $\boldsymbol{M}$ is obvious from the context or not important, we will simply use the notation $K(f, \ell)$.

[^2]:    ${ }^{2}$ As earlier, $\boldsymbol{M}$ will be dropped from the subscript when it is not important or clear from the context.

[^3]:    ${ }^{3}$ In the literature, contour, join, and split trees are often defined as trees on the set of critical vertices and regular vertices are regarded as points along the edges and not as vertices of the tree. As defined here, these trees are usually called the "augmented" versions of themselves.

[^4]:    ${ }^{4}$ The asymmetry here is due to the fact that we cannot lower the height of any vertex below that of $v_{\infty}$. The choice of $v_{\infty}$ as the root of $\check{T}$ then puts restrictions on the choice of $\hat{r}$.

