

I/O-Efficient Batched Union-Find and Its Applications to Terrain Analysis

Pankaj K. Agarwal

Duke University

and

Lars Arge

MADALGO, University of Aarhus

and

Ke Yi

Hong Kong University of Science and Technology

In this paper we present an I/O-efficient algorithm for the batched (off-line) version of the union-find problem. Given any sequence of N union and find operations, where each union operation joins two distinct sets, our algorithm uses $O(\text{SORT}(N)) = O(\frac{N}{B} \log_{M/B} \frac{N}{B})$ I/Os, where M is the memory size and B is the disk block size. This bound is asymptotically optimal in the worst case. If there are union operations that join a set with itself, our algorithm uses $O(\text{SORT}(N) + \text{MST}(N))$ I/Os, where $\text{MST}(N)$ is the number of I/Os needed to compute the minimum spanning tree of a graph with N edges. We also describe a simple and practical $O(\text{SORT}(N) \log(\frac{N}{M}))$ -I/O algorithm for this problem, which we have implemented.

We are interested in the union-find problem because of its applications in terrain analysis. A terrain can be abstracted as a height function defined over \mathbb{R}^2 , and many problems that deal with such functions require a union-find data structure. With the emergence of modern mapping technologies, huge amount of elevation data is being generated that is too large to fit in memory, thus I/O-efficient algorithms are needed to process this data efficiently. In this paper, we study two terrain-analysis problems that benefit from a union-find data structure: (i) computing topological persistence and (ii) constructing the contour tree. We give the first $O(\text{SORT}(N))$ -I/O algorithms for these two problems, assuming that the input terrain is represented as a triangular mesh with N vertices.

P. K. Agarwal and K. Yi was supported by ARO grant W911NF-04-1-0278, by NSF under grants CCR-00-86013, EIA-01-31905, CCR-02-04118, and DEB-04-25465, and by a grant from the U.S.–Israel Binational Science Foundation. L. Arge was supported by ARO grant W911NF-04-1-0278, by an Ole Rømer Scholarship from the Danish National Science Research Council, by a NABIT grant from the Danish Strategic Research Council, and by the Danish National Research Foundation. K. Yi was also supported by Hong Kong Direct Allocation Grant (DAG07/08).

Part of this work was done while K. Yi was at Duke University.

A preliminary version of the paper appeared in the 22nd Annual Symposium on Computational Geometry (SoCG), 2006.

Authors addresses: P. K. Agarwal, Department of Computer Science, Duke University, Durham, NC 27708, USA; email: pankaj@cs.duke.edu. L. Arge, Department of Computer Science, MADALGO, University of Aarhus, IT-Parken, Aabogade 34, DK-8200 Aarhus N, Denmark; email: large@daimi.au.dk. K. Yi, Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China; email: yike@cse.ust.hk.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0004-5411/20YY/0100-0001 \$5.00

Categories and Subject Descriptors: F.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: I/O-efficient algorithms, Union-find

1. INTRODUCTION

The *union-find* problem asks for maintaining a partition of a set $U = \{x_1, x_2, \dots\}$ (the universe) and a *representative* element of each set in this partition under a sequence Σ of $\text{UNION}(x_i, x_j)$ and $\text{FIND}(x_i)$ operations: $\text{UNION}(x_i, x_j)$ joins the set containing x_i and the set containing x_j , and provides a new representative element for the new set; $\text{FIND}(x_i)$ returns the representative element of the set containing x_i . In the *on-line* version of the problem, Σ is given one operation at a time, whereas in the *batched* (or *off-line*) version, the entire sequence Σ is known in advance. The union-find problem is a fundamental algorithmic problem because of its applications in numerous problems across different domains, from programming languages to graph and geometric algorithms, and from computational topology to computational biology. In many of these algorithms only the batched version of the problem is required; see [Carr et al. 2003; Cormen et al. 2001; Edelsbrunner et al. 2000; Gabow and Tarjan 1985] for a sample of applications.

The main motivation for our study of the union-find problem arises from terrain modeling and analysis. A terrain can be abstracted as a height function defined over \mathbb{R}^2 , and there is a rich literature on the study of such functions [Milnor 1963]. We are interested in two broad problems in terrain analysis, namely flow and contour-line analysis. A key step in the flow analysis of a terrain is to modify the height function so that “small” depressions on the terrain (sinks) disappear. We use the notion of *topological persistence*, introduced in [Edelsbrunner et al. 2000], to address this problem. In the contour-line analysis, the notion of *contour tree* is critical [Carr et al. 2003; Tarasov and Vyalıy 1998; van Kreveld et al. 1997]. Most existing topological persistence and contour-tree algorithms rely on efficient data structures for the batched union-find problem.

With the emergence of high-resolution terrain-mapping technologies, huge amount of data is being generated that is too large to fit in memory and has to reside on disks. Existing algorithms cannot handle such massive data sets, mainly because they optimize CPU running time while optimizing disk access is much more important. Motivated by these factors we propose efficient algorithms for the batched union-find problem in the *I/O-model* [Aggarwal and Vitter 1988] (also known as the *external memory model*), and use them to develop I/O-efficient algorithms for computing topological persistence and contour trees.

Related results. In the I/O model, the machine consists of an infinite-size external memory (disk) and a main memory of size M . 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 occur on elements in main memory, and the complexity of an algorithm is measured in terms of the number of I/Os it uses to solve a problem. Many fundamental problems have been solved in the I/O model. For example, sorting N elements takes $\text{SORT}(N) = \Theta\left(\frac{N}{B} \log_{M/B} \frac{N}{B}\right)$ I/Os, and permuting N

elements takes $\text{PERM}(N) = \Theta(\min\{N, \text{SORT}(N)\})$ I/Os. Note that the lower bounds for sorting and permuting hold in the *indivisible model* [Arge 1995; Zeh 2002]. See the surveys by Vitter [2001] and Arge [2002] for other results, including I/O-efficient algorithms for various geometric and graph problems.

The study of the union-find problem (in internal memory computation models) started early in the sixties [Galler and Fisher 1964] and continues even today [Seidel and Sharir 2005]. Tarjan [1975] proved that an on-line sequence of N union and find operations can be performed in time $O(N\alpha(N))$, where $\alpha(N)$ is the inverse Ackermann function. This bound is tight in the worst case [Fredman and Saks 1989; Tarjan 1979]. A simpler proof of the upper bound and various generalizations have been proposed in [Seidel and Sharir 2005; Tarjan and Leeuwen 1984]. The batched version of the problem has a lower bound of $\Omega(N\alpha(N))$ in the pointer-machine model [Tarjan 1979], but it can be solved in linear time in the RAM model [Gabow and Tarjan 1985]. However, despite of its importance, no I/O-efficient algorithms have been developed for the union-find problem. Note that the naïve use of the RAM algorithms in the I/O model results in $O(N\alpha(N))$ and $O(N)$ I/O-algorithms, in the on-line and batched case, respectively. For realistic values of N , B , and M , $\text{SORT}(N) \ll N$, and the difference in running time between an algorithm performing N I/Os and one performing $\text{SORT}(N)$ I/Os can be very significant.

Topological persistence (see Section 4 and [Edelsbrunner et al. 2001; Edelsbrunner and Harer 2008] for its definition) of a height function is a measure of its topological attributes. Over the last few years, it has been successfully applied to a variety of problems, including topological simplification [Edelsbrunner and Harer 2008], identifying features on a surface [Agarwal et al. 2004], and removing topological noises [Bremer et al. 2004]; see also [Zomorodian 2005]. Efficient algorithms for computing persistence in internal memory are developed in [Edelsbrunner et al. 2001; Edelsbrunner et al. 2000]. The efficient algorithms for computing the topological persistence of a two-dimensional height function rely on a union-find data structure, so no I/O-efficient algorithm is known for this problem.

Contour trees are widely used to represent the topological changes in the contours of a height function (see Section 5 and [Carr et al. 2003; Tarasov and Vyalys 1998; van Kreveld et al. 1997] for a formal definition). van Kreveld et al. [1997] gave an $O(N \log N)$ -time algorithm for constructing the contour tree of a piecewise-linear height function on \mathbb{R}^2 , represented by a mesh with N vertices. The algorithm was later extended to 3D by Tarasov and Vyalys [1998], and to arbitrary dimensions by Carr et al. [2003]. The last two algorithms rely on a union-find data structure. No I/O-efficient algorithm is known for computing contour trees even in two dimensions.

Our results. Our main result is the first I/O-efficient algorithms for the batched union-find problem. In Section 2, we present an algorithm that uses $O(\text{SORT}(N))$ I/Os on a sequence of N union and find operations, provided that none of the union operations is *redundant*, that is, for each $\text{UNION}(x_i, x_j)$, x_i and x_j are in different sets. An interesting feature of our algorithm is that it reduces the problem to two geometric problems. Note that if $N < \text{SORT}(N)$ we can simply use the $O(N)$ -time RAM algorithm [Gabow and Tarjan 1985] to obtain an algorithm using $O(N)$ I/Os. Thus we in total solve the problem in $O(\text{PERM}(N))$ I/Os, which is

optimal. If redundant union operations are allowed, we describe an algorithm that uses $O(\text{SORT}(N) + \text{MST}(N))$ I/Os, where $\text{MST}(N)$ is the number of I/Os needed to compute the minimum spanning tree of a graph with N edges. Currently the best bound for $\text{MST}(N)$ is $O(\text{SORT}(N) \log \log B)$ I/Os [Arge et al. 2004] by a deterministic algorithm or expected $O(\text{SORT}(N))$ I/Os by a randomized algorithm [Chiang et al. 1995]. We also present a very simple divide-and-conquer algorithm for the batched union-find problem in Section 3 that uses $O(\text{SORT}(N) \cdot \log(\frac{N}{M}))$ I/Os. It does not have to invoke an MST algorithm to handle redundant union operations. We have implemented this algorithm. Our empirical results show that our algorithm gives order-of-magnitude improvement over previous methods on large data sets.

We develop the first I/O-efficient algorithms for topological persistence and contour trees using our union-find algorithm. In Section 4, we describe an $O(\text{SORT}(N))$ -I/O algorithm for computing topological persistence of a terrain represented as a triangular mesh with N vertices. The algorithm is obtained by simply using our union-find algorithm in the previous internal memory algorithm of Edelsbrunner et al. [2001]. We have implemented the algorithm and our empirical results show that it works well in practice. In Section 5, we describe an $O(\text{SORT}(N))$ -I/O algorithm for computing the contour tree of a terrain represented as a triangular mesh with N vertices. For this problem it is not enough to just apply our union-find algorithm to the previous internal memory algorithms. We prove a property of contour trees (cf. Lemma 5.4), which is interesting in its own right. Using this property, we reduce the problem to a geometric problem and design an I/O-efficient algorithm for this problem. As in Carr et al. [2003], our algorithm extends to higher dimensions.

2. I/O-OPTIMAL BATCHED UNION-FIND

Let $U = \{x_1, x_2, \dots\}$ be the universe of elements, and let $\Sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_N \rangle$ be a sequence of union and find operations. Let the index of the operation σ_i denote its *time stamp*. Assume without loss of generality that each element in U has been used in at least one operation of Σ . In this section we describe an I/O-efficient algorithm for the batched union-find problem, assuming that all union operations in Σ are non-redundant. At the end of the section we discuss how to handle redundant union operations.

Our algorithm consists of two stages. In the first stage, we compute a total ordering on the elements of U so that as Σ is performed, each set consists of a contiguous sequence of elements in this ordering and each union operation joins two adjacent sets. More precisely, if sets A and B are joined together by any union operation in Σ , then there exist some $x_i \in A$ and $x_j \in B$ such that x_i and x_j are adjacent in this ordering. We call this special case the *interval union-find* problem. In the second stage, we solve an instance of the batched interval union-find problem. We formulate each of these two subproblems as a geometric problem and present an $O(\text{SORT}(N))$ -I/O algorithm for each of them. We also note that the assumption on non-redundant union operations is needed only in the first stage.

2.1 From union-find to interval union-find

Given a sequence Σ , let the *union graph* $G(\Sigma)$ be a weighted undirected graph whose vertices are the elements of U . There is an edge e between x_i and x_j with

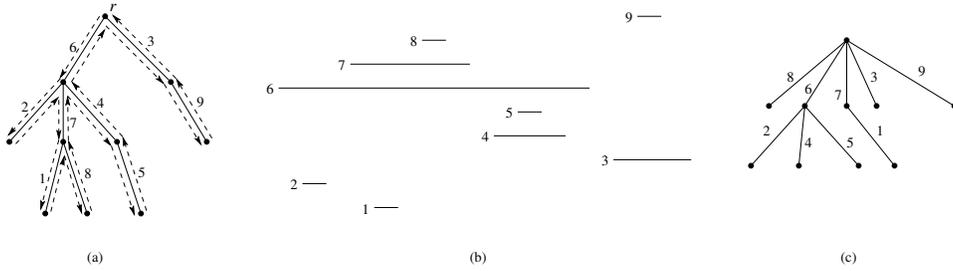


Fig. 1. (a) The union tree T , with weights (time stamps) associated with each edge. The dashed lines show the Euler tour. (b) A horizontal segment is built for each edge e of T , with $\omega(e)$ as its y -coordinate and the positions of e in the Euler tour as its left and right x -coordinates. (c) The equivalent union tree T' after transformation.

weight $\omega(e) = t$ for each $\text{UNION}(x_i, x_j)$ operation with time stamp t . If all union operations are non-redundant, $G(\Sigma)$ is a forest. For simplicity we assume it to be a tree; otherwise, we connect them by edges of weight ∞ . We pick an arbitrary node of $G(\Sigma)$ as its root and call this rooted tree the *union tree*, denoted by T . As a convention, when we refer to an edge (u, v) in a rooted tree, u is always the parent of v .

A union tree T' on the same set of nodes as T is *equivalent* to T if for any t , the connected components of the forest formed by the edges of T and T' with weight at most t are the same. In order to reduce the union-find problem to the interval union-find problem, we first transform T into an equivalent union tree T' , with the property that the weights along any leaf-to-root path are increasing. Then we show how a certain in-order traversal of T' defines an ordering of the elements of U that results in Σ being an interval union-find instance.

Transforming T to T' . We transform T into T' by repeatedly applying the following “short-cut” operation, which simulates the *path compression* technique: For any $w \in T$, let v be w ’s parent and u be v ’s parent. If $\omega(u, v) < \omega(v, w)$, we promote w to be a child of u , by removing the edge (v, w) and adding the edge (u, w) with $\omega(u, w) = \omega(v, w)$. It is easy to see that the new tree resulted after this operation is equivalent to the old tree, since at the time $\text{UNION}(v, w)$ is issued, u and v are already in the same set. When this procedure stops, we obtain a tree T' in which the weights along any leaf-to-root path are increasing.

In order to perform all these short-cut operations efficiently, for each node v , we need to find its ancestor in T that becomes the parent of v in T' . We cast this problem in a geometric setting. We construct an Euler tour on T , which starts and ends at r and traverses every edge of T exactly twice, once in each direction (see Fig. 1(a)). Such a tour can be constructed in $O(\text{SORT}(N))$ I/Os [Chiang et al. 1995]. For each edge $e \in T$, we map it to a horizontal line segment whose y -coordinate is $\omega(e)$, and whose left (resp. right) x -coordinate is the position where e appears for the first (resp. second) time in the Euler tour. Note that the x -spans of these segments are nested. Refer to Fig. 1(a) and (b). These segments can be easily constructed in $O(\text{SORT}(N))$ I/Os given the Euler tour.

Abusing the notation, we also use e to denote the segment that corresponds to

the edge $e \in T$. Let E be this set of segments. Let $\gamma(e)$ be the shortest segment in E that lies above e and whose x -projection contains that of e . We claim that if $e = (u, v)$ and $\gamma(e) = (w, z)$, then (z, v) is an edge of T' , i.e., z is the parent of v in T' . Indeed, the second condition ($\gamma(e)$ above e) ensures that v cannot be shortcut to w , and the first condition ($\gamma(e)$ being the shortest) and the third condition (the x -projection of $\gamma(e)$ containing that of e) ensure that z is the lowest ancestor of v that satisfies the second condition. If $\gamma(e)$ does not exist for some $e = (u, v)$, we know that v must be a child of the root r . Refer to Fig. 1(c).

Thus the problem of computing the parent in T' of each node $u \in T$ reduces to computing $\gamma(e)$ for each e . Since the segments in E are nested, the x -span of a segment e contains that of another segment e' if and only if the x -span of e contains the x -coordinate of the left endpoint of e' . Let P be the set of left endpoints of the segments in E . For a point $p \in P$ and a subset $X \subseteq E$, let $\gamma(p, X)$ denote the shortest segment of X that is above p and whose x -span contains the x -coordinate of p . The problem now reduces to computing $\gamma(p, E)$ for each $p \in P$. We use the distribution sweeping technique [Goodrich et al. 1993] to solve this problem. We first sort E and P by y -coordinates. We then divide the plane into $m = M/B$ vertical slabs L_1, \dots, L_m , each of which contains roughly the same number of segment endpoints of E . Let $P_i = P \cap L_i$, let E_i be the set of segments with at least one endpoint inside L_i , and E_i^* the set of segments that completely span L_i . Note that for any point $p \in P_i$, $\gamma(p, E)$ is the shorter segment of $\gamma(p, E_i)$ and $\gamma(p, E_i^*)$. We compute the sets E_i , P_i , and the segment $\gamma(p, E_i^*)$ for each $p \in P_i$, by a sweep-line algorithm, and compute $\gamma(p, E_i)$ for each $p \in P_i$ and for $1 \leq i \leq m$ recursively. See Fig. 2. We sweep the plane in the $(-y)$ -direction starting from $y = \infty$. For each slab L_i , we maintain in memory the shortest segment λ_i among the segments swept so far that completely span L_i . When the sweep line reaches a point $p \in P$, we determine the slab L_i that contains p , set $\gamma(p, E_i^*)$ to λ_i , and add p to P_i . When the sweep line reaches a segment e of E , we add e to E_i if L_i contains at least one of the endpoints of e , and update λ_j if e spans L_j and e is shorter than λ_j . We recursively solve the problem for each (E_i, P_i) . The sweep can be implemented by a scan of P and E . Since the depth of recursion is $O(\log_{M/B} \frac{N}{B})$, the overall cost is $O(\text{SORT}(N))$ I/Os.

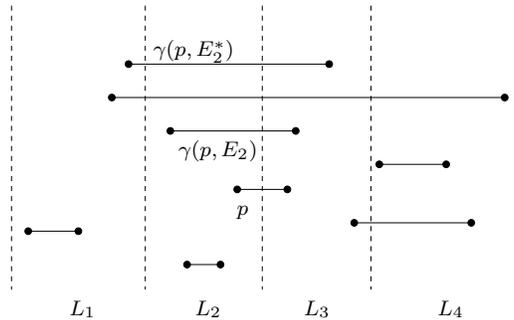


Fig. 2. Use the distribution sweeping technique to compute $\gamma(e)$ for each e . For each $p \in P$, we compute $\gamma(p, E_i^*)$ during the sweep, and compute $\gamma(p, E_i)$ recursively.

Traversing T' . To obtain an ordering of the elements of U that results in Σ being an interval union-find instance, we perform a weight-guided, in-order traversal on T' starting from the root. When we reach a node u of T' with children u_1, \dots, u_k where $w(u, u_1) < \dots < w(u, u_k)$, first we recursively visit the subtree rooted at u_1 , then visit u , and then recursively visit the subtrees rooted at u_2, \dots, u_k . To perform this traversal I/O-efficiently we first order each node's children in weight order, then compute an Euler tour. A leaf of T' appears once in the tour, and an internal node with k children appears $k + 1$ times. It is easy to verify that, if we delete all but the second appearance of an internal node u from the tour, then the desired order is obtained. To do so we first associate with each appearance of u with a node id and its position in the Euler tour, then sort them by the node id, and finally do a scan to remove all but the second appearance of each internal node.

LEMMA 2.1. *We can convert a sequence of N union and find operations without redundant union operations into an instance of the batched interval union-find problem in $O(\text{SORT}(N))$ I/Os.*

PROOF. It only remains to show that this ordering indeed produces an instance of the interval union-find problem. Let u be any node of T' , and let u_1, \dots, u_k be u 's children where $w(u, u_1) < \dots < w(u, u_k)$. For any $1 \leq i \leq k$, when the union operation corresponding to (u, u_i) is to be performed, in the sequence Σ , all union operations corresponding to the whole subtree of u_i must have already been performed, since they have smaller weights than $w(u, u_i)$. If $i = 1$, then u must not have been joined with any other node at the time. So the union operation corresponding to the edge (u, u_1) joins u_1 's subtree and the singleton set $\{u\}$, which is immediately after u_1 's subtree in the ordering. If $i \geq 2$, then all the subtrees of u_1, \dots, u_{i-1} must have been joined together with u , and u_i 's subtree is immediately after these subtrees in the ordering, so the union operation corresponding to the edge (u, u_i) also joins adjacent sets. \square

2.2 Solving interval union-find

We solve the interval union-find problem by formulating it as another geometric problem. Let $x_1 < x_2 < \dots$ be the sequence of ordered elements of U . For each union operation $\sigma_t = \text{UNION}(x_i, x_j)$ in Σ , we create a horizontal line segment with y -coordinate t , x -coordinate x_i , and right x -coordinate x_j . For each find operation $\sigma_t = \text{FIND}(x_i)$ in Σ , we create a query point (x_i, t) . For $\sigma_t = \text{FIND}(x_i)$ we return the smallest element of the set as its representative at time t that contains x_i . For each query point q , consider the union of the x -projections of all the segments lower than q . Each interval in the union corresponds to a set at time t (when $\text{FIND}(x_i)$ was performed), so we need to return the left endpoint of the interval that contains x_i .

We solve this geometric problem by solving a series of *batched orthogonal ray-shooting* problems, in which we are given a set of horizontal segments and a set of query points, the goal is to find the first segment hit by a vertical ray shooting upwards from each query point. This is a special case of the *endpoint dominance* problem [Arge et al. 2007] and can be solved in $O(\text{SORT}(N))$ I/Os. First, from the left endpoint of each segment e we shoot a ray downwards, and collect all the segments whose rays do not hit any other segments; these are the only endpoints

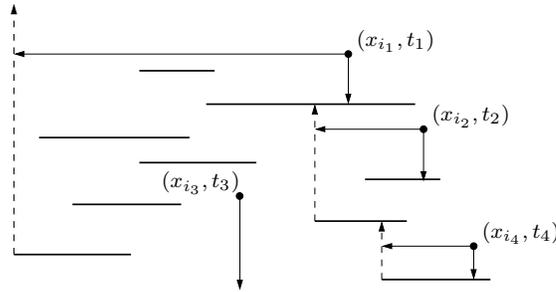


Fig. 3. Answering find queries $\sigma_t = \text{FIND}(x_i)$. If no segment lies below a left endpoint, it is connected to the segment lying immediately above it by a dashed vertical line.

that can be representatives in the partition (except singleton sets). Second, we shoot upwards from the left endpoints of these segments. These rays stop at the first horizontal segments they hit, or go to infinity. Let V be the resulting set of vertical segments. See Fig. 3 where these rays are drawn as dashed lines. Finally, from each query point $q = (x_i, t)$, we shoot a ray downwards and hit a horizontal segment. If the ray does not hit any segment we simply return x_i as the answer to q ; otherwise, we shoot a ray leftward from q and find the first segment of V hit by the ray. Then the x -coordinate of this vertical segment is returned as the answer to q . It is not difficult to verify that we indeed return the correct representative element for each query point. The whole query answering process is four instances of the batched orthogonal ray-shooting problem, thus can be completed in $O(\text{SORT}(N))$ I/Os.

LEMMA 2.2. *Given a sequence of N union and find operations possibly with redundant union operations, the batched interval union-find problem can be solved in $O(\text{SORT}(N))$ I/Os.*

Combining Lemmas 2.1 and 2.2, we obtain the main result.

THEOREM 2.3. *Given a sequence of N union and find operations with no redundant union operations, the batched union-find problem can be solved in $O(\text{SORT}(N))$ I/Os.*

Remark. In some applications, certain elements are required to be the representative elements of the sets returned by the find operations. For example, each element may be weighted, and the minimum-weight element of a set is required to be the representative of the set (see Section 4 and 5). Our algorithm as described above may pick arbitrary elements as the representatives, but such a requirement can be satisfied by running our algorithm twice, followed by a post-processing step in $O(\text{SORT}(N))$ I/Os: We first answer all the find queries by returning the minimum element in the ordering produced by our reduction from union-find to interval union-find; then by solving the interval union-find problem symmetrically we can also return the maximum elements in the same ordering. Now we have both the minimum and maximum elements of the set for each find query, returning the minimum-weight element in the set becomes a range-minimum query. Again, this batched range-minimum problem can be solved in $O(\text{SORT}(N))$ I/Os using distri-

bution sweeping [Goodrich et al. 1993].

2.3 Lower bound

We can prove an $\Omega(\text{PERM}(N))$ lower bound on the batched union-find problem. Formally, we assume that the input is given as two lists stored on disk. The first list is a sequence of (id, element) pairs, one for each element in the universe U , where the id is a distinct integer from 1 to $|U|$. The second list is the sequence of N union and find operations, where each union operation is represented by the id's of the two elements, while each find operation is represented by an actual element from U . The output should be given as a sequence of element pairs, where each pair consists of an element defining a find operation and the corresponding representative element. We can easily modify our algorithm to conform to this input and output representations, while still using $O(\text{SORT}(N))$ I/Os. Similar to many other graph problems, we will prove the lower bound in the *indivisibility model*, in which an algorithm can only make copies of existing elements (and e.g. not delete elements or make new elements). Our algorithm works in this model; see [Arge 1995; Zeh 2002] for a formal definition of the model.

We will show that if we have an algorithm \mathcal{A} that solves the batched union-find algorithm using $f(N)$ I/Os, then we can construct an algorithm \mathcal{B} that solves the so-called *split proximate neighbor problem* (SPN) using $O(N/B + f(N))$ I/Os. Since SPN has an I/O lower bound of $O(\text{PERM}(N))$ in the indivisibility model [Arge 1995], we then also obtain a lower bound for the batched union-find problem. In the SPN problem we are given N elements, each containing a key $1 \leq k \leq N/2$: The i -th element has key i for $1 \leq i \leq N/2$, while the last $N/2$ elements may have arbitrary but distinct keys. The goal is to permute the N elements such that the two elements with the same key are stored next to each other on disk.

More precisely, our reduction goes as follows. Given an SPN input instance, the algorithm \mathcal{B} first scans the N elements and constructs a sequence of (id, element) pairs. For the first $N/2$ elements the key k is used as the id, while $k + N/2$ is used for the last $N/2$ elements. Note that to solve the SPN instance the elements with id k and id $k + N/2$ should be output next to each other. Next \mathcal{B} constructs a sequence of $N/2$ union operations, followed by N find operations, represented by id pairs and elements, respectively. More precisely, the sequence consists of a union operation $\text{UNION}(k, k + N/2)$ for each $1 \leq k \leq N/2$, followed by a find operation for each input element. The second sequence can easily be constructed in a scan of the input elements. Thus \mathcal{B} has constructed the two input lists for a batched union-find problem in $O(N/B)$ I/Os. It now invokes algorithm \mathcal{A} on this input. After \mathcal{A} finishes, it then scans the N find results using $O(N/B)$ I/Os and outputs x and y consecutively if a $\text{FIND}(x)$ operation returned $y \neq x$. It is easy to see that this way \mathcal{B} solves the original SPN instance. Since \mathcal{B} uses $O(N/B + f(N))$ I/Os the lower bound follows.

2.4 Handling redundant union operations

So far we have assumed that there are no redundant union operations in Σ . If this assumption does not hold, then $G(\Sigma)$ contains cycles. We first compute the MST (minimum spanning forest in general) on the union graph $G(\Sigma)$, and delete all edges that are not in the MST. Because for any non-MST edge (u, v) , there must

be another path from u to v , consisting edges with weights lower than that of (u, v) , the union operation $\text{UNION}(u, v)$ must be redundant. For a general graph with N edges, computing the MST takes $O(\text{SORT}(N) \cdot \log \log B)$ I/Os deterministically [Arge et al. 2004] or expected $O(\text{SORT}(N))$ I/Os by a randomized algorithm [Chiang et al. 1995]; if the graph is planar, computing the MST takes deterministic $O(\text{SORT}(N))$ I/Os [Chiang et al. 1995]. So we have the following.

THEOREM 2.4. *Given a sequence Σ of N union and find operations possibly with redundant union operations, the batched union-find problem can be solved by a deterministic algorithm with $O(\text{SORT}(N) \log \log B)$ I/Os or a randomized algorithm with expected $O(\text{SORT}(N))$ I/Os. If the union graph of Σ is planar, $O(\text{SORT}(N))$ I/Os are needed in the worst case.*

2.5 A persistent union-find data structure

Suppose we are given a sequence of N union operations, each with a time stamp. A $\text{FIND}(x, t)$ query asks for the representative of the set containing a query element x at a specific time t , and we want to build a data structure that supports these queries in an online fashion. Such a structure that allows queries in the past is also called a *partially persistent* (or simply *persistent*) structure [Becker et al. 1996], and it is useful in situations in which one is not only interested in the data in their latest version, but also in their development over time.

We can design a persistent union-find data structure by combining our algorithm and the persistent B-tree [Becker et al. 1996]. Observe that we have not used the information of the find queries until the very last two steps in our batched union-find algorithm, when we shoot a ray downwards from the query point and then shoot to the left. Such *orthogonal ray-shooting* queries can be handled by a *persistent B-tree* [Becker et al. 1996] in worst-case $O(\log_B N)$ I/Os per query. Thus we run our batched union-find algorithm until we have generated those horizontal and vertical segments of Fig. 3, and then build two persistent B-trees on these segments, respectively. Since a persistent B-tree takes linear space and can be constructed in $O(\text{SORT}(N))$ I/Os [van den Bercken et al. 1997], we have the following.

THEOREM 2.5. *Given a sequence of N union operations, with $O(\text{SORT}(N) + \text{MST}(N))$ I/Os, we can build a linear-size persistent data structure for the union-find problem such that a find query at any specific time can be answered in $O(\log_B N)$ I/Os.*

3. A SIMPLER BATCHED UNION-FIND ALGORITHM

In this section we present a simple divide-and-conquer algorithm that bypasses the need of an MST algorithm. It uses $O(\text{SORT}(N) \cdot \log(\frac{N}{M}))$ I/Os, but has a much smaller constant of proportionality, and can be easily implemented. The input to a recursive call is a sequence Σ of union and find operations. The recursive call outputs the answers of all $\text{FIND}(x_i)$ queries in Σ and returns a set R of $(x, \rho(x))$ pairs, one for each element x involved in any operation in Σ , where $\rho(x)$ is the representative of the set containing x after all union operations in Σ are performed. The basic idea behind a recursive call, outlined in Algorithm 4, is the following. If Σ fits in main memory, we use an internal memory algorithm; otherwise we split Σ into two halves Σ_1 and Σ_2 . We solve Σ_1 recursively. Before solving Σ_2

recursively, we use the element-representative set R_1 , returned by the recursive call for Σ_1 , to pass on “information” to Σ_2 about how the sets are joined in Σ_1 . We do so by replacing each element x involved in any operation in Σ_2 with $\varrho(x)$ if $(x, \varrho(x)) \in R_1$ (line (a)). When the second recursive call on Σ_2 finishes, we need to return the complete and correct element-representative set to the upper level calls. All element-representative pairs in R_2 are correct, but some in R_1 might get updated. We update each $(x, \varrho(x)) \in R_1$ with $(x, \varrho(y))$, if there exists some $(y, \varrho(y)) \in R_2$ such that $\varrho(x) = y$ (line (b)). Finally we return the union of R_1 and R_2 . Both line (a) and (b) can be performed by a constant number of sort and scan steps, so the total cost of the algorithm is $O(\text{SORT}(N) \log(\frac{N}{M}))$ I/Os.

Since this practical algorithm allows redundant union operations, one immediate application is a straightforward implementation of Kruskal’s algorithm for computing MSTs. Given a graph G , we sort all edges of G by weight, and construct a union-find sequence by issuing two FIND and one UNION operations on its two endpoints for each edge in order. After solving this batched union-find problem we check the two FIND results for each edge e . If they are different, we declare that e is in the MST. This simple algorithm runs in deterministic $O(\text{SORT}(N) \log(\frac{N}{M}))$ I/Os. Previously the only known practical external memory MST algorithm [Dementiev et al. 2004] was randomized with $O(\text{SORT}(N) \log(\frac{N}{M}))$ I/Os expected. Similarly we can compute the connected components of a graph I/O-efficiently using our practical union-find algorithm.

Experimental results. To investigate the practical performance of our simple batched union-find algorithm, we implemented it and used it in Kruskal’s MST algorithm [Cormen et al. 2001]. We compared the performance of the algorithm with the performance of Kruskal’s MST algorithm implemented using the internal memory algorithm that uses link-by-rank and path compression [Cormen et al. 2001]; we could have implemented the linear-time off-line union-find algorithm [Gabow and Tarjan 1985], but it is quite complicated and is unlikely to outperform the simple $O(N\alpha(N))$ on-line algorithm in practice. We should also point out that there exist other (theoretical and practical) external memory MST algorithms [Arge et al. 2004; Dementiev et al. 2004] that do not use the union-find algorithm. However, since our focus is on investigating the performance of union-find algorithms, we

input: a sequence Σ of union and find operations;
output: a set R of $(x, \varrho(x))$ pairs for each element x involved in Σ .

if Σ can be processed in main memory **then**
 Call an internal memory algorithm;
else
 Split Σ into two halves Σ_1 and Σ_2 ;
 $R_1 = \text{UNION-FIND}(\Sigma_1)$;
 (a) For $\forall(x, \varrho(x)) \in R_1$, replace all occurrences of x in Σ_2 with $\varrho(x)$;
 $R_2 = \text{UNION-FIND}(\Sigma_2)$;
 (b) For $\forall(x, \varrho(x)) \in R_1$, if $\exists(y, \varrho(y)) \in R_2$ s.t. $y = \varrho(x)$,
 replace $(x, \varrho(x))$ with $(x, \varrho(y))$ in R_1 ;
 return $R_1 \cup R_2$.

Fig. 4. Recursive call UNION-FIND(Σ).

did not compare with any of these algorithms. We implemented our algorithm using the TPIE system [Arge et al. 2002] that supports efficient implementation of theoretically I/O-efficient algorithms. For the experiments we used a 2.4GHz Intel XEON machine with hyperthreading, running Linux with kernel 2.4.5-smp, and a local disk system consisting of four 10000RPM 72GB SCSI disks in RAID-0 configuration. In the experiments we performed, we limited the physical memory of the test machine to 128M bytes. While TPIE ensured that our practical algorithm never used more than these 128M bytes, the internal algorithm used much more and we relied on virtual memory (the operating system) to handle I/Os (swapping).

For experimental data we generated random graphs with N vertices and $5N$ edges, similar to the earlier experimental studies of MST algorithms [Dementiev et al. 2004; Moret and Shapiro 1991]: For each edge, we randomly picked its two endpoints from all the N vertices, and also randomly picked its weight uniformly from $(0, 1)$. We tested the internal memory algorithm (IM) and the external memory algorithm (EM) on eight graphs, with N between one million to eight million. The results are shown in Fig. 5, where the running time reported is just the time spent by the union-find phase of the algorithms, that is, it does not include the initial sorting done by both algorithms. The 128M bytes main memory could accommodate the data structure for a graph with roughly up to two million vertices. As can be seen from Fig. 5, the two algorithms performed roughly the same when the graph fits in main memory. In fact, in this case EM is the same as IM except for an additional step to check whether the input can be handled in memory. For larger graphs, the running time of IM immediately deteriorated due to its random memory access pattern. On the other hand, the running time of EM scaled roughly linearly in the size of the input graph. On the graph with three million vertices, IM spent roughly 22 hours, that is, 300 times more than that of EM. We could not run IM on the other data sets in a reasonable amount of time.

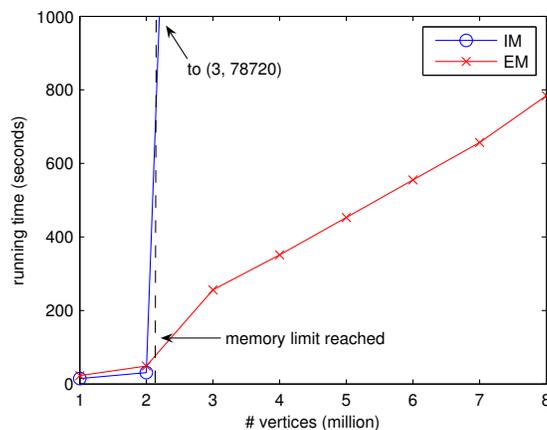


Fig. 5. Running time of the internal memory (IM) and the external memory (EM) union-find algorithms for computing minimum spanning trees.

4. TOPOLOGICAL PERSISTENCE

Let \mathbb{M} be a triangulation of \mathbb{R}^2 in which a height is associated with its every vertex. \mathbb{M} defines a piecewise-linear *height function* $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. It will be convenient to assume $f(u) \neq f(v)$ for all vertices $u \neq v$ in \mathbb{M} . The *link* of a vertex u is the set of edges that are contained in the triangles adjacent to u but that do not contain u . The *lower link* $\text{Lk}^-(u)$ of u is the set of edges (w, z) in the link such that $f(w), f(z) < f(u)$. The *upper link* $\text{Lk}^+(u)$ is defined similarly. There are three types of *critical vertices* on \mathbb{M} : *minima*, *saddles*, and *maxima*, distinguished by the number of connected components in u 's lower link. Refer to Fig. 6. For ease of representation we assume that each saddle is *simple*, i.e., $\text{Lk}^-(u)$ for each saddle u consists of exactly two connected components. If $\text{Lk}^-(u)$ has more than two connected components (e.g. the rightmost case in Fig. 6), we can unfold it into a number of simple saddles [Edelsbrunner et al. 2001].

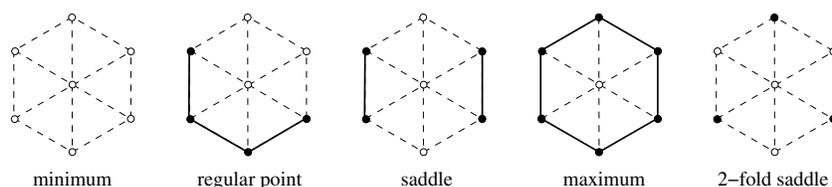


Fig. 6. Classification of a vertex based its lower link. The lower link is marked black.

Let $\mathbb{M}_{<h} = \{x \in \mathbb{R}^2 \mid f(x) < h\}$; define $\mathbb{M}_{>h}$ similarly. Roughly speaking, *topological persistence* of f (or \mathbb{M}) is defined as follows. Suppose we sweep a horizontal plane in the direction of increasing values of f , and keep track of $\mathbb{M}_{<h}$ as we increase h . A component of $\mathbb{M}_{<h}$ starts at a minimum and ends at a saddle when it merges with another, older component. A hole of $\mathbb{M}_{<h}$ starts at a saddle and ends when it is closed off at a maximum. Based on this observation, Edelsbrunner et al. [2001] propose a scheme to pair a critical point that creates a component or a hole with the one that destroys it. For each minimum-saddle or maximum-saddle pair, the absolute difference between their heights is defined as the *persistence* of that pair of critical vertices. See [Edelsbrunner et al. 2001] for a formal definition.

We first sketch the algorithm in [Edelsbrunner et al. 2001], and then describe how to make it I/O-efficient. We focus on the minimum-saddle pairs, as the other case is symmetric. The algorithm first sorts all vertices by their height, then we sweep a horizontal plane bottom-up and maintain the connected components of $\mathbb{M}_{<h}$ when the sweep plane passes through $f(x) = h$. For each component of $\mathbb{M}_{<h}$, we maintain its minimum-height vertex as its representative. When we pass a minimum, a new component is created. When we pass a saddle u , the two components adjacent to u get merged if they were different. Let v and w , with $f(v) < f(w)$, be the minima of these two components. We pair w with u and choose v to be the representative of the merged component. Notice that v is the minimum of the merged component. Passing regular points or maxima does not change the set of components.

We use an I/O-efficient algorithm to sort the vertices by their heights. After sorting, the sweep can be implemented using a batched union-find algorithm, as

follows. When passing a regular vertex u , we pick an arbitrary vertex v in $\text{Lk}^-(u)$ and issue a $\text{UNION}(u, v)$ operation. When passing a saddle u , we first issue two find operations: For each of the two connected components of $\text{Lk}^-(u)$, an arbitrary vertex v in the component is chosen and a $\text{FIND}(v)$ is issued. Next for each such v , we issue the union operation $\text{UNION}(u, v)$. After sweeping over all vertices, we invoke our batched union-find algorithm on the resulting sequence Σ . Using the remark on specific representatives after Theorem 2.3, we can use the vertex of the lowest height as the representative of the component it belongs to. Finally, we scan the results of the find queries. If the answers to the two FIND queries of a saddle are different, we pair the higher one with the saddle. It can be checked that the union graph of Σ is planar, so according to Theorem 2.4, processing the batched sequence of union-find operations takes $O(\text{SORT}(N))$ I/Os.

THEOREM 4.1. *Given a height function over a triangular mesh in \mathbb{R}^2 with N vertices, the topological persistence of all of its critical vertices can be computed using $O(\text{SORT}(N))$ I/Os.*

Experimental results. To investigate the practical performance of the topological persistence algorithm we implemented it using both the internal memory (IM) and our practical (EM) union-find algorithm. We tested the performance of these algorithms using elevation data for the Neuse River Basin of North Carolina obtained from [North Carolina Flood Mapping Program. <http://www.ncfloodmaps.com>]. The data set consists of over 0.5×10^9 points and we created eight smaller data sets by sampling 1 million to 8 million points at random; we converted each of these data sets into a triangular mesh using an I/O-efficient Delaunay triangulation algorithm [Agarwal et al. 2005]. We used the same experimental platform (a machine with 128M bytes of main memory) as described in Section 3.

The running time for the two algorithms on each of the eight data sets are shown in Fig. 7; the running time reported is just the time spent in the union-find phase of the algorithms (note the log scale on the running time). As can be seen from Fig. 7, we see similar trends of both IM and EM as the MST experiments discussed in Section 3, except that IM deteriorates slower in this case. This can be explained by the fact that the union and find operations have better locality in the topological persistence case than in the MST case due to the underlying geometry. Still, the difference in running time between IM and EM gets more significant as the data set gets larger. On the last data set of 8 million points, the difference is more than two orders of magnitude.

We also ran an experiment on the entire Neuse River Basin data set (with 0.5 billion points) using a machine with 1G bytes of main memory. In this case, EM finishes in about 5.22 hours, while IM crashed after running for about 7 hours, because the 32-bit address space is not enough to accommodate such a large data set: The size of this data set is around 18G bytes, far beyond the limit a 32-bit pointer could address.

5. CONTOUR TREES

Let \mathbb{M} be a triangular mesh in \mathbb{R}^2 as defined in Section 4. A *contour* in \mathbb{M} is a connected component of the level set of f at some height h (i.e., $\{x \in \mathbb{R}^2 \mid f(x) = h\}$). As we vary h , the contours vary continuously but they appear and disappear

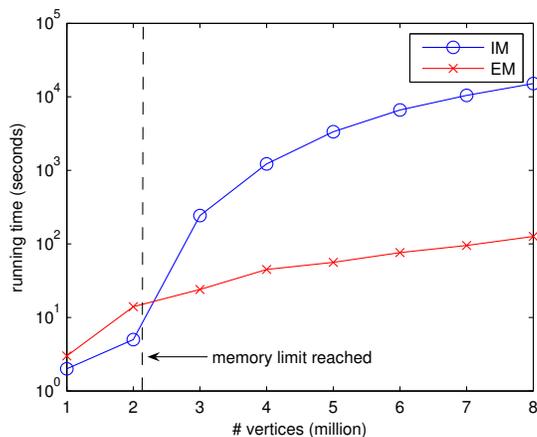


Fig. 7. Running time of the internal memory (IM) and the external memory (EM) union-find algorithms for computing topological persistence (running time is shown on log scale).

only at the critical vertices of \mathbb{M} . If we increase the height, then a new contour appears at a minimum, a contour disappears at a maximum, two contours join at a saddle, or a contour splits into two at a saddle. The *contour tree* is a graph (actually, a tree) that tracks these changes. Its nodes are the critical vertices in \mathbb{M} . Each edge (u, v) in the tree corresponds to a contour that is created at v and destroyed at u . See Fig. 8 for an example. In the *augmented* contour tree, we add all the regular vertices to the tree as well. More precisely, for each edge (u, v) in the contour tree, if its corresponding contour passes through regular vertices w_1, \dots, w_k as we vary h from $f(u)$ to $f(v)$, then (u, v) is replaced by a chain of edges $(u, w_1), (w_1, w_2), \dots, (w_k, v)$. We only describe how we build the augmented contour tree \mathcal{C} , but our algorithm can be easily modified to build the contour tree. We omit the word “augmented” from now on.

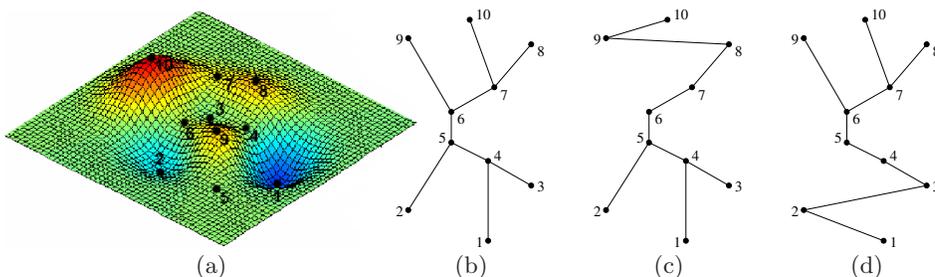


Fig. 8. (a) A height function f , and its (b) contour tree, (c) join tree, (d) split tree.

We follow the same approach as in [Carr et al. 2003]. We first construct the so-called *join* and *split* trees of \mathbb{M} and then merge them to construct \mathcal{C} . The join tree $\mathcal{J}_{\mathbb{M}}$ represents all joins of contours in \mathcal{C} , and the split tree $\mathcal{S}_{\mathbb{M}}$ represents all

splits of contours in \mathcal{C} . $\mathcal{J}_{\mathbb{M}}$ is a tree in which (u, v) , with $f(u) > f(v)$, is an edge if v is the highest vertex in a connected component of $\mathbb{M}_{<f(u)}$ that contains a vertex of $\text{Lk}^-(u)$. The tree $\mathcal{S}_{\mathbb{M}}$ is defined symmetrically by negating the function f . See Fig. 8(c) and (d) for an example. For any edge $(u, v) \in \mathcal{J}_{\mathbb{M}}$ with $f(u) > f(v)$, we declare u to be the parent of v , so $\mathcal{J}_{\mathbb{M}}$ is a tree growing downward with the highest node being the root. Similarly for any edge in the split tree, we declare the lower node to be parent, and it becomes a tree growing upward with the lowest node being the root.

The topological persistence algorithm can be adapted to compute $\mathcal{J}_{\mathbb{M}}$. We sweep a plane bottom-up and maintain the connected components of $\mathbb{M}_{<h}$. When the sweep plane reaches $f(u)$, for each $v \in \text{Lk}^-(u)$, we connect u with the highest node in the connected component of $\mathbb{M}_{<h}$ that contains v . Duplicate edges are removed. This can be done using our batched union-find algorithm in the same way as the topological persistence case, except that here we require the highest node to be the representative of a component, rather than the lowest. The split tree $\mathcal{S}_{\mathbb{M}}$ is built symmetrically. As in the persistence case, the union graph is planar, so the algorithm takes $O(\text{SORT}(N))$ I/Os.

Since the definition of lower links applies to any graph in which height is associated with its vertices, the notion of join and split trees can be extended to \mathcal{C} (i.e., replace \mathbb{M} with \mathcal{C} verbatim in the definition) and the construction algorithm makes sense even when applied to \mathcal{C} instead of \mathbb{M} . The following lemma is proved by Carr et al. [2003].

LEMMA 5.1 [CARR ET AL. 2003]. *The contour tree \mathcal{C} and the mesh \mathbb{M} have the same join tree and split tree.*

Carr et al. define a linear-time algorithm to construct \mathcal{C} from $\mathcal{J}_{\mathbb{M}} = \mathcal{J}_{\mathcal{C}}$ and $\mathcal{S}_{\mathbb{M}} = \mathcal{S}_{\mathcal{C}}$ using the following observation. A node v is *qualified* if it is a leaf of $\mathcal{J}_{\mathcal{C}}$ (resp. $\mathcal{S}_{\mathcal{C}}$) and has only one child in $\mathcal{S}_{\mathcal{C}}$ (resp. $\mathcal{J}_{\mathcal{C}}$), and we call the edge (u, v) a *qualified edge*, where u is the parent of v in $\mathcal{J}_{\mathcal{C}}$ (resp. $\mathcal{S}_{\mathcal{C}}$). A qualified node must be a leaf of \mathcal{C} , and any leaf of \mathcal{C} must also be a qualified node. For a qualified node v , let $\mathcal{J}_{\mathcal{C}} \ominus v$ (or $\mathcal{S}_{\mathcal{C}} \ominus v$) be the tree obtained after deleting the node v and the edges adjacent to it; if v is not a leaf, connect u to the only child of v by an edge.

LEMMA 5.2 [CARR ET AL. 2003]. *A qualified edge (u, v) is an edge of \mathcal{C} . If v is a qualified, then $\mathcal{J}_{\mathcal{C}} \ominus v$ and $\mathcal{S}_{\mathcal{C}} \ominus v$ are the join and split trees of $\mathcal{C} \ominus v$.*

Using this lemma, Carr et al. [2003] repeatedly delete a qualified node from $\mathcal{J}_{\mathcal{C}}$ and $\mathcal{S}_{\mathcal{C}}$, and add the corresponding qualified edge to \mathcal{C} . Fig. 9(a)–(d) illustrate one such iteration, and Fig. 9(e) shows the final contour tree constructed. In internal memory, each step takes $O(1)$ time, but we cannot afford even one I/O per step. It is difficult to batch the steps together since the set of qualified nodes depends on the choices of which qualified nodes to delete in previous steps. We therefore present a different characterization of contour tree edges, formulate the problem as a geometric problem, and solve it using $O(\text{SORT}(N))$ I/Os.

As a convention, we say a node is both an ancestor and a descendant of itself. For any edge $e = (u, v) \in \mathcal{J}_{\mathbb{M}}$, let $\sigma(e)$ be the highest node that is a descendant of v in $\mathcal{J}_{\mathbb{M}}$ and at the same time an ancestor of u in $\mathcal{S}_{\mathbb{M}}$.

LEMMA 5.3. *For any edge $e \in \mathcal{J}_{\mathbb{M}}$, the node $\sigma(e)$ is defined.*

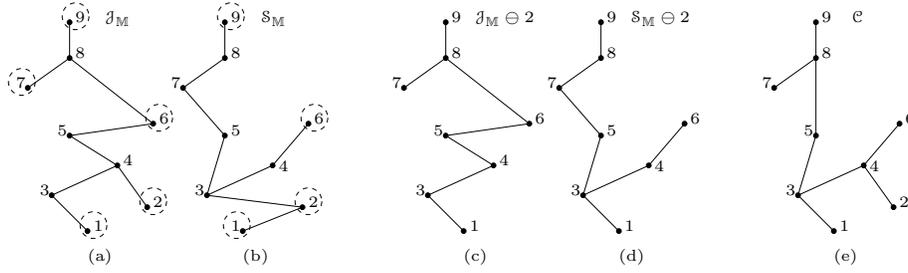


Fig. 9. Construction of a contour tree from the join tree and the split tree using the Carr et al. [2003] algorithm: (a) The join tree \mathcal{J}_M of a mesh \mathbb{M} , where the number beside each node is their height. (b) The split tree \mathcal{S}_M . All qualified nodes are circled. (c)-(d) One step of Carr et al. algorithm: node 2 is removed and the edge between node 2 and node 4 is added to \mathcal{C} . (e) The final contour tree \mathcal{C} .

PROOF. it can be checked that the descendants of a node w in \mathcal{J}_M (resp. \mathcal{S}_M) are exactly those lying in the connected component of the closure of $\mathbb{M}_{<f(w)}$ (resp. $\mathbb{M}_{>f(w)}$) that contains w . Now consider any edge $e = (u, v) \in \mathcal{J}_M$. During the construction of \mathcal{J}_M , u is connected to v because there is a vertex $w \in \text{Lk}^-(u)$ such that v is the highest node in the connected component of $\mathbb{M}_{<f(u)}$ that contains w . So w is an descendant of v in \mathcal{J}_M .

During the construction of \mathcal{S}_M , when the edge (u, w) of \mathbb{M} is processed at w (recall that the vertices are swept in decreasing order of their heights), w is connected to the lowest node w' in the connected component of $\mathbb{M}_{>f(w)}$ that contains u . Hence w' , and thus w , is an ancestor of u in \mathcal{S}_M . This shows that there exists a vertex w that is a descendant of u in \mathcal{J}_M and an ancestor of u in \mathcal{S}_M , thereby implying that $\sigma(e)$ is defined. \square

The following lemma characterizes the edges of \mathcal{C} .

LEMMA 5.4. *For each edge $e = (u, v) \in \mathcal{J}_M$, $(u, \sigma(e))$ is an edge of \mathcal{C} .*

PROOF. Please refer to Fig. 10(a). Let $e = (u, v)$ be any edge of \mathcal{J}_M . We know that $\sigma(e)$ is the highest node that is both a descendant of v in \mathcal{J}_M and an ancestor of u in \mathcal{S}_M . Let P_1 be the set of nodes on the path of \mathcal{J}_M from u to $\sigma(e)$, and P_2 the set of nodes on the path of \mathcal{S}_M from $\sigma(e)$ to u ; u and $\sigma(e)$ are not included in P_1 and P_2 . Note that P_1 and P_2 are disjoint, otherwise we would have chosen a different $\sigma(e)$.

We follow the procedure in [Carr et al. 2003], i.e., repeatedly delete qualified nodes one by one except u until $\sigma(e)$ becomes qualified. This strategy always works since at any time, the remaining part of the contour tree yet to be constructed has at least two leaves. So there are at least two qualified nodes, and we can always find one other than u to remove. Since u is not deleted and u was originally a descendant of $\sigma(e)$ in \mathcal{S}_M , u remains a descendant of $\sigma(e)$ when the latter becomes qualified. Since $\sigma(e)$ is not a leaf in \mathcal{S}_M when it becomes qualified, by definition, it must be a leaf in \mathcal{J}_M and has one child in \mathcal{S}_M .

Next we prove that when $\sigma(e)$ becomes qualified, P_1 must be empty, i.e., u is the parent of $\sigma(e)$ in \mathcal{J}_M , and thus $(u, \sigma(e))$ is added to \mathcal{C} . Suppose on the contrary $P_1 \neq \emptyset$ when $\sigma(e)$ becomes qualified, and $w \in P_1$ is the parent of $\sigma(e)$ in \mathcal{J}_M . Since

such that u is an ancestor of v if and only if $x_1(u) < x_1(v)$ and $x_2(u) > x_2(v)$. Note that these intervals are nested. We do the same with \mathcal{S}_M , and compute another set of nested intervals $[y_1(u), y_2(u)]$. We map each node u to a vertical segment $s(u)$ with x -coordinate $x_1(u)$ and y -span $[y_1(u), y_2(u)]$, and map each edge $e = (u, v) \in \mathcal{J}_M$ to a horizontal segment $s(e)$ with x -span $[x_1(v), x_2(v)]$ and y -coordinate $y_1(u)$. It is easy to see that for any node w and any $e = (u, v) \in \mathcal{J}_M$, w is a descendant of v in \mathcal{J}_M and an ancestor of u in \mathcal{S}_M if and only if $s(w)$ intersects $s(e)$. Thus the problem of finding $\sigma(e)$ can be formulated as follows: for the horizontal segment $s(e)$, find the shortest vertical segment that intersects $s(e)$.

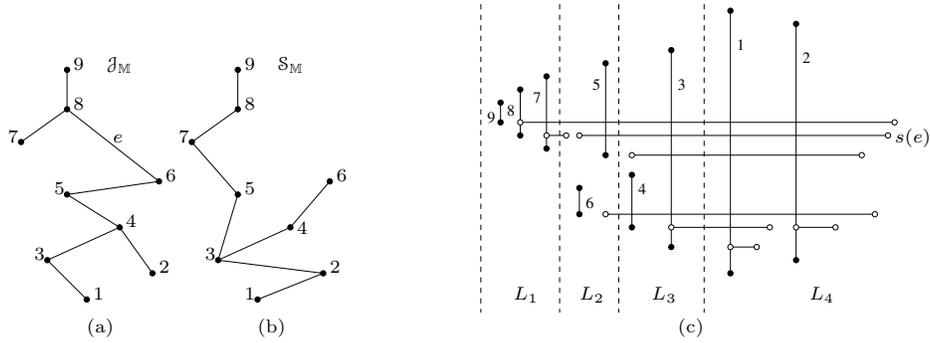


Fig. 11. (a) The join tree. (b) The split tree. (c) Map each node to a vertical segment and map each edge in the join tree to a horizontal segment, and then solve the segment-intersecting problem using distribution sweep. The horizontal segment $s(e)$ is broken into three pieces: the middle piece spans slab L_3 , the left leftover piece falls inside L_2 , and the right leftover piece falls inside L_4 .

We solve this batched problem using the distribution sweeping technique. During the execution, with each horizontal segment we store its shortest intersecting vertical segment found so far. Initially we sort the endpoints of all the segments by the y -coordinate. Then we divide the plane into $\Theta(M/B)$ vertical slabs, each of which contains roughly equal number of vertical segments. Next we sweep a horizontal line top-down. During the sweep, for each slab L we maintain an external memory stack that stores all vertical segments that fall inside L . Whenever we reach the top endpoint of a vertical segment, we push it into the corresponding stack; whenever we reach the bottom endpoint of a vertical segment, we pop one from the corresponding stack. Since there are $O(M/B)$ stacks, we can allocate one memory block for each stack, so that all stack operations take linear I/Os. When we reach a horizontal segment, we look at all the slabs that are completely spanned by it, and check the top element of each corresponding stack as a candidate answer.

During the sweep, we also form the instances of the subproblems for each slab. The vertical segments are simply distributed to their corresponding slabs. For each horizontal segment, we break it into at most three pieces: a middle piece that completely spans a number of slabs, a left leftover piece and a right leftover piece that fall inside a slab (see Fig. 11). We throw away its middle piece, and distribute the two leftover pieces into the corresponding slab. Note that at any level in the

recursion, a horizontal segment has at most two pieces. Thus employing standard distribution sweeping analysis yields an $O(\text{SORT}(N))$ -I/O bound for the algorithm.

Putting all the pieces together, we obtain the following.

THEOREM 5.5. *The contour tree of a triangular mesh with N vertices in \mathbb{R}^2 can be computed using $O(\text{SORT}(N))$ I/Os.*

Remark. Our algorithm extends to higher dimensions, with N replaced by the number of simplices in the mesh [Carr et al. 2003], except that in $d \geq 3$ dimensions the union graph when building the join and split trees may not be planar, and the $O(\text{SORT}(N) \log \log B)$ -I/O MST algorithm or the randomized MST algorithm with expected $O(\text{SORT}(N))$ I/Os needs to be used.

6. CONCLUSION

In this paper we gave the first I/O-efficient algorithms for the off-line version of the classical union-find problem. We presented two algorithms: a theoretically optimal algorithm, and a simpler algorithm that is a log factor from being I/O-optimal. Using our new union-find algorithm, we also presented I/O-efficient topological persistence and contour trees algorithms. We have used the topological persistence algorithm for computing flow networks on terrains [Danner et al. 2007]. We have also recently simplified the algorithm for computing contour trees and implemented it, and are currently using it for contour-line analysis on large terrains.

A natural question to ask is whether we can develop an I/O-efficient algorithm for the *online* union-find algorithm. However, in this case we cannot hope for anything below $\Omega(N)$ I/Os for a sequence of N operations, since an adversary can always issue the next operation on an element (or elements) that are not currently in main memory. Using the internal memory algorithm [Cormen et al. 2001] directly results in an $O(N\alpha(N))$ -I/O algorithm. It would be a very interesting problem to close this gap. Ideally one would hope for an algorithm that uses, say $O(N\lceil\alpha(N)/B\rceil)$ I/Os, but this will be only of theoretical interests since in practice $\alpha(N)$ is much smaller than B .

REFERENCES

- AGARWAL, P. K., ARGE, L., AND YI, K. 2005. I/O-efficient construction of constrained Delaunay triangulations. In *Proc. 13th Annual European Symposium on Algorithms*. 355–366.
- AGARWAL, P. K., EDELSBRUNNER, H., HARER, J., AND WANG, Y. 2006. Extreme elevation on a 2-manifold. In *Discrete and Computational Geometry 36*, 553–572.
- AGGARWAL, A. AND VITTER, J. S. 1988. The input/output complexity of sorting and related problems. *Communications of the ACM* 31, 9, 1116–1127.
- ARGE, L. 1995. The I/O-complexity of ordered binary-decision diagram manipulation. In *Proc. 6th International Symposium on Algorithms and Computation*. 82–91.
- ARGE, L. 2002. External memory data structures. In *Handbook of Massive Data Sets*, J. Abello, P. M. Pardalos, and M. G. C. Resende, Eds. Kluwer Academic Publishers, 313–358.
- ARGE, L., BRODAL, G. S., AND TOMA, L. 2004. On external memory MST, SSSP and multi-way planar graph separation. *Journal of Algorithms* 53,, 186–206.
- ARGE, L., PROCOPIUC, O., AND VITTER, J. S. 2002. Implementing I/O-efficient data structures using TPIE. In *Proc. 10th European Symposium on Algorithms*. 88–100.
- ARGE, L., VENGROFF, D. E., AND VITTER, J. S. 1995. External-memory algorithms for processing line segments in geographic information systems. *Algorithmica* 47, 1–25.

- BECKER, B., GSCHWIND, S., OHLER, T., SEEGER, B., AND WIDMAYER, P. 1996. An asymptotically optimal multiversion B-tree. *VLDB Journal* 5, 264–275.
- BREMER, P. T., PASCUCCI, V., EDELSBRUNNER, H., AND HAMANN, B. 2004. A topological hierarchy for functions on triangulated surfaces. *IEEE Trans. Visualization and Computer Graphics* 10, 385–396.
- CARR, H., SNOEYINK, J., AND AXEN, U. 2003. Computing contour trees in all dimensions. *Computational Geometry: Theory and Applications* 24, 75–94.
- CHIANG, Y.-J., GOODRICH, M. T., GROVE, E. F., TAMASSIA, R., VENGROFF, D. E., AND VITTER, J. S. 1995. External-memory graph algorithms. In *Proc. 6th Annual ACM-SIAM Symposium on Discrete Algorithms*. 139–149.
- CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L., AND STEIN, C. 2001. *Introduction to Algorithms, 2nd Edition*. The MIT Press, Cambridge, Mass.
- DANNER, A., MØLHAVE, T., YI, K., AGARWAL, P. K., ARGE, L., AND MITASOVA, H. 2007. Ter-rastream: From elevation data to watershed hierarchies. In *Proc. 9th International Symposium on Advances in Geographic Information Systems*.
- DEMENTIEV, R., SANDERS, P., SCHULTES, D., AND SIBEYN, J. 2004. Engineering an external memory minimum spanning tree algorithm. In *Proc. 3rd IFIP International Conference on Theoretical Computer Science*. 195–208.
- EDELSBRUNNER, H. AND HARER, J. 2008. Persistent homology—a survey. In *Surveys on Discrete and Computational Geometry* (Goodman, J. E., Pach, J., and Pollack, R. eds.), American Mathematical Society, Providence, RI. 257–282.
- EDELSBRUNNER, H., HARER, J., AND ZOMORODIAN, A. 2001. Hierarchical Morse complexes for piecewise linear 2-manifolds. In *Proc. 17th Annual Symposium on Computational Geometry*. 70–79.
- EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. 2000. Topological persistence and simplification. In *Proc. 41st Annual IEEE Symposium on Foundations of Computer Science*. 454–463.
- FREDMAN, M. L. AND SAKS, M. E. 1989. The cell probe complexity of dynamic data structures. In *Proc. 21st Annual ACM Symposium on Theory of Computing*. 345–354.
- GABOW, H. AND TARJAN, R. 1985. A linear time algorithm for a special case of disjoint set union. *Journal of Computer and Systems Sciences* 30, 209–221.
- GALLER, B. A. AND FISHER, M. J. 1964. An improved equivalence algorithm. *Communications of the ACM* 7, 5, 301–303.
- GOODRICH, M. T., TSAY, J.-J., VENGROFF, D. E., AND VITTER, J. S. 1993. External-memory computational geometry. In *Proc. 34th IEEE Symposium on Foundations of Computer Science*. 714–723.
- MILNOR, J. W. 1963. *Morse Theory*. Princeton University Press, Princeton, NJ.
- MORET, B. M. E. AND SHAPIRO, H. D. 1991. An empirical analysis of algorithms for constructing a minimum spanning tree. In *Proc. 2nd Workshop Algorithms Data Structures*. 400–411.
- NORTH CAROLINA FLOOD MAPPING PROGRAM. [HTTP://WWW.NCFLOODMAPS.COM](http://www.ncfloodmaps.com).
- SEIDEL, R. AND SHARIR, M. 2005. Top-down analysis of path compression. *SIAM Journal on Computing* 34, 3, 515–525.
- TARASOV, S. P. AND VYALYI, M. N. 1998. Construction of contour trees in 3D in $O(n \log n)$ steps. In *Proc. 14th Annual Symposium on Computational Geometry*. 68–75.
- TARJAN, R. E. 1975. Efficiency of a good but not linear set union algorithm. *Journal of the ACM* 22, 215–225.
- TARJAN, R. E. 1979. A class of algorithms that require nonlinear time to maintain disjoint sets. *Journal of Computer and System Sciences* 18, 110–127.
- TARJAN, R. E. AND LEEUWEN, J. V. 1984. Worst-case analysis of set union algorithms. *Journal of the ACM* 31, 245–281.
- VAN DEN BERCKEN, J., SEEGER, B., AND WIDMAYER, P. 1997. A generic approach to bulk loading multidimensional index structures. In *Proc. 23rd International Conference on Very Large Databases*. 406–415.

- VAN KREVELD, M., VAN OOSTRUM, R., BAJAJ, C., PASCUCCI, V., AND SCHIKORE, D. 1997. Contour trees and small seed sets for isosurface traversal. In *Proc. 13th Annual Symposium on Computational Geometry*. 212–219.
- VITTER, J. S. 2001. External memory algorithms and data structures: Dealing with MASSIVE data. *ACM Computing Surveys* 33, 2, 209–271.
- ZEH, N. 2002. I/O-efficient graph algorithms. *EEF Summer School on Massive Data Sets*.
- ZOMORODIAN, A. 2005. *Topology for Computing*. Cambridge University Press, New York.