

Time-Space Trade-offs for Triangulations and Voronoi Diagrams

Matias Korman^{a,1}, Wolfgang Mulzer^{d,1}, André van Renssen^{b,c}, Marcel Roeloffzen^{b,c}, Paul Seiferth^{d,1}, Yannik Stein^{d,1}

^a*Tohoku University, Sendai, Japan.*

mati@dais.is.tohoku.ac.jp.

^b*National Institute of Informatics (NII), Tokyo, Japan.*

{andre,marcel}@nii.ac.jp

^c*JST, ERATO, Kawarabayashi Large Graph Project.*

^d*Institut für Informatik, Freie Universität Berlin, Germany.*

{mulzer,pseiferth,yannikstein}@inf.fu-berlin.de

Abstract

Let S be a planar n -point set. A *triangulation* for S is a maximal plane straight-line graph with vertex set S . The *Voronoi diagram* for S is the subdivision of the plane into cells such that all points in a cell have the same nearest neighbor in S . Classically, both structures can be computed in $O(n \log n)$ time and $O(n)$ space. We study the situation when the available workspace is limited: given a parameter $s \in \{1, \dots, n\}$, an s -workspace algorithm has read-only access to an input array with the points from S in arbitrary order, and it may use only $O(s)$ additional words of $\Theta(\log n)$ bits for reading and writing intermediate data. The output should then be written to a write-only structure. We describe a deterministic s -workspace algorithm for computing an arbitrary triangulation of S in time $O(n^2/s + n \log n \log s)$ and a randomized s -workspace algorithm for finding the Voronoi diagram of S in expected time $O((n^2/s) \log s + n \log s \log^* s)$.

1. Introduction

Since the early days of computer science, a major concern has been to cope with strong memory constraints. This started in the '70s [22] when memory was expensive. Nowadays, a major motivation comes from a proliferation of small embedded devices where large memory is neither feasible nor desirable (e.g., due to constraints on budget, power, size, or simply to make the device less attractive to thieves).

Even when memory size is not an issue, we might want to limit the number of write operations: one can read flash memory quickly, but writing (or

¹MK was supported in part by the ELC project (MEXT KAKENHI No. 24106008). WS and PS were supported in part by DFG Grants MU 3501/1 and MU 3501/2. YS was supported by the DFG within the research training group "Methods for Discrete Structures" (GRK 1408).

even reordering) data is slow and may reduce the lifetime of the storage system; write-access to removable memory may be limited for technical or security reasons (e.g., when using read-only media such as DVDs or to prevent leaking information about the algorithm). Similar problems occur when concurrent algorithms access data simultaneously. A natural way to address this is to consider algorithms that do not modify the input.

The exact setting may vary, but there is a common theme: the input resides in read-only memory, the output must be written to a write-only structure, and we can use $O(s)$ additional variables to find the solution (for a parameter s). The goal is to design algorithms whose running time decreases as s grows, giving a *time-space trade-off* [23]. One of the first problems considered in this model is *sorting* [19, 20]. Here, the time-space product is known to be $\Omega(n^2)$ [8], and matching upper bounds for the case $b \in \Omega(\log n) \cap O(n/\log n)$ were obtained by Pagter and Rauhe [21] (b denotes the available workspace in *bits*).

Our current notion of memory constrained algorithms was introduced to computational geometry by Asano *et al.* [4], who showed how to compute many classic geometric structures with $O(1)$ workspace (related models were studied before [9]). Later, time-space trade-offs were given for problems on simple polygons, e.g., shortest paths [1], visibility [6], or the convex hull of the vertices [5].

We consider a model in which the set S of n points is in an array such that random access to each input point is possible, but we may not change or even reorder the input. Additionally, we have $O(s)$ variables (for a parameter $s \in \{1, \dots, n\}$). We assume that each variable or pointer contains a data word of $\Theta(\log n)$ bits. Other than this, the model allows the usual word RAM operations. In this setting we study two problems: computing an arbitrary triangulation for S and computing the Voronoi diagram $\text{VD}(S)$ for S . Since the output cannot be stored explicitly, the goal is to report the edges of the triangulation or the vertices of $\text{VD}(S)$ successively, in no particular order. Dually, the latter goal may be phrased in terms of Delaunay triangulations. We focus on Voronoi diagrams, as they lead to a more natural presentation.

Both problems can be solved in $O(n^2)$ time with $O(1)$ workspace [4] or in $O(n \log n)$ time with $O(n)$ workspace [7]. However, to the best of our knowledge, no trade-offs were known before. Our triangulation algorithm achieves a running time of $O(n^2/s + n \log n \log s)$ using $O(s)$ variables. A key ingredient is the recent time-space trade-off by Asano and Kirkpatrick for triangulating a special type of simple polygons [3]. This also lets us obtain significantly better running times for the case that the input is sorted in x -order; see Section 2. For Voronoi diagrams, we use random sampling to find the result in expected time $O((n^2 \log s)/s + n \log s \log^* s)$; see Section 3. Together with recent work of Har-Peled [16], this appears to be one of the first uses of random sampling to obtain space-time trade-offs for geometric algorithms. The sorting lower bounds also apply to triangulations and Voronoi diagrams (since we can reduce the former to the latter). This implies that our second algorithm is almost optimal.

2. A Time-Space Trade-Off for General Triangulations

In this section we describe an algorithm that outputs the edges of a triangulation for a given point set S in arbitrary order. For ease in the presentation we first assume that S is presented in sorted order. In this case, a time-space trade-off follows quite readily from known results. We then show how to generalize this for arbitrary inputs, which requires a careful adaptation of the existing data structures.

2.1. Sorted Input

Suppose the input points $S = \{q_1, \dots, q_n\}$ are stored in increasing order of x -coordinate and that all x -coordinates are distinct, i.e., $x_i < x_{i+1}$ for $1 \leq i < n$, where x_i denotes the x -coordinate of q_i .

A crucial ingredient in our algorithm is a recent result by Asano and Kirkpatrick for triangulating *monotone mountains*² (or *mountains* for short). A mountain is a simple polygon with vertex sequence v_1, v_2, \dots, v_k such that the x -coordinates of the vertices increase monotonically. The edge v_1v_k is called the *base*. Mountains can be triangulated very efficiently with bounded workspace.

Theorem 2.1 (Lemma 3 in [3], rephrased). *Let H be a mountain with n vertices, stored in sorted x -order in read-only memory. Let $s \in \{2, \dots, n\}$. We can report the edges of a triangulation of H in $O(n \log_s n)$ time and using $O(s)$ words of space.*

Since S is given in x -order, the edges q_iq_{i+1} , for $1 \leq i < n$, form a monotone simple polygonal chain. Let $\text{Part}(S)$ be the subdivision obtained by the union of this chain with the edges of the convex hull of S (denoted by $\text{conv}(S)$). A convex hull edge is *long* if the difference between its indices is at least two (i.e., the endpoints are not consecutive). The following lemma (illustrated in Fig. 1) lets us decompose the problem into smaller pieces.




Figure 1: Any face of $\text{Part}(S)$ is a mountain that is uniquely associated with a long convex hull edge.

²Also known as *unimotone polygons* [15].

Lemma 2.2. *Any bounded face of $\text{Part}(S)$ is a mountain whose base is a long convex hull edge. Moreover, no point of S lies in more than four faces of $\text{Part}(S)$.*

Proof. Any point $q_i \in S$ has at most four neighbors in $\text{Part}(S)$: q_{i-1} , q_{i+1} , its predecessor and its successor along the convex hull (if q_i lies on $\text{conv}(S)$). Thus, no point of S belongs to more than four faces of $\text{Part}(S)$.

Next we show that every face F of $\text{Part}(S)$ is a mountain with a long convex-hull edge as its base. The boundary of F contains at least one long convex-hull edge $e = (q_i, q_j)$ ($i < j$), as other edges connect only consecutive vertices. Since the monotone path q_i, \dots, q_j forms a cycle with the edge e and since the boundary of F is a simple polygon, we conclude that e is the only long convex-hull edge bounding F . Recall that e is a convex hull edge, and thus all points q_{i+1}, \dots, q_{j-1} lie on one side of e and form a monotone chain (and in particular F is a mountain with base e). \square

The algorithm for sorted input is now very simple. We compute the edges of the convex hull (starting from the leftmost point and proceeding in clockwise order). Whenever a long edge would be reported, we pause the convex hull algorithm, and we triangulate the corresponding mountain. Once the mountain has been triangulated, we resume with the convex hull algorithm until all convex hull edges have been computed. The trade-off now follows from already existing trade-offs in the various subroutines.

Theorem 2.3. *Let S be a set of n points, sorted in x -order. We can report the edges of a triangulation of S in $O(n^2)$ time using $O(1)$ variables, in $O(n^2 \log n / 2^s)$ time using $O(s)$ variables (for any $s \in \Omega(\log \log n) \cap o(\log n)$), and in $O(n \log_p n)$ time using $O(p \log_p n)$ variables (for any $2 \leq p \leq n$).*

Proof. Correctness follows from Lemma 2.2, so we focus on the performance analysis. The main steps are: (i) computing the convex hull of a point set given in x -order; and (ii) triangulating a mountain.

By Theorem 2.1, we can triangulate a mountain F_i with n_i vertices in time $O(n_i \log_s n_i)$ with $O(s)$ variables. We do not need to store F_i explicitly, since its vertices constitute a consecutive subsequence of S and can be specified by the two endpoints of the base. No vertex appears in more than four mountains by Lemma 2.2, so the total time for triangulating the mountains is $\sum_i O(n_i \log_s n_i) = O(n \log_s n)$. By reusing space, we can ensure that the total space requirement is $O(s)$.

Now we bound the time for computing $\text{conv}(S)$. This algorithm is paused to triangulate mountains, but overall it is executed only once. There are several convex hull algorithms for sorted point sets under memory constraints. If $s \in \Theta(1)$, we can use gift-wrapping (Jarvis march [17]), which runs in $O(n^2)$ time. Barba *et al.* [5] provided a different algorithm that runs in $O(n^2 \log n / 2^s)$ time using $O(s)$ variables (for any $s \in o(\log n)$).³ This approach is desirable for $s \in \Omega(\log \log n) \cap o(\log n)$. As soon as $s = \Omega(\log n)$, we can use the approach of Chan

³In fact, Barba *et al.* show how to compute the convex hull of a simple polygon, but also

and Chen [10]. This algorithm runs in $O(n \log_p n)$ time and uses $O(p \log_p n)$ variables, for any $2 \leq p \leq n$. Regardless of the size of the workspace, the time for computing the convex hull dominates the time needed for triangulating all mountains. \square

A similar approach is unlikely to work for the Delaunay triangulation, since knowing the x -order of the input does not help in computing it [14].

2.2. General Input

The algorithm from Section 2.1 uses the sorted order in two ways. Firstly, the convex-hull algorithms of Barba *et al.* [5] and of Chan and Chen [10] work only for simple polygons (e.g., for sorted input). Instead, we use the algorithm by Darwish and Elmasry [13] that gives the upper (or lower) convex hull of any sequence of n points in $O(n^2/(s \log n) + n \log n)$ time with $O(s)$ variables⁴, matching known lower bounds. Secondly, and more importantly, the Asano-Kirkpatrick (AK) algorithm for triangulating a mountain requires the input to be sorted. To address this issue, we simulate sorted input using multiple heap structures. This requires a close examination of how the AK-algorithm accesses its input.

Let F be a mountain with n vertices. Let F^\uparrow and F^\downarrow denote the vertices of F in ascending and in descending x -order. The AK-algorithm has two phases, one focused on F^\uparrow and the other one on F^\downarrow .⁵ Each pass computes a portion of the triangulation edges, uses $O(s)$ variables, and scans the input $\Theta(\log_s n)$ times. We focus on the approach for F^\uparrow .

As mentioned, the algorithm uses $\Theta(\log_s n)$ rounds. In round i , it partitions F into blocks of $O(|F|/s^i)$ consecutive points that are processed from left to right. Each block is further subdivided into $O(s)$ sub-blocks b_1, \dots, b_k of size $O(|F|/s^{i+1})$. The algorithm does two scans over the sub-blocks. The first scan processes the elements from left to right. Whenever the first scan finishes reading a sub-block b_i , the algorithm makes b_i *active* and creates a pointer l_i to the rightmost element of b_i . The second scan goes from right to left and is concurrent to the first scan. In each step, it reads the element at l_i in the rightmost active sub-block b_i , and it decreases l_i by one. If l_i leaves b_i , then b_i becomes inactive. As the first scan creates new active sub-blocks as it proceeds, the second scan may jump between sub-blocks. The interested reader may find a more detailed description in Appendix A.

To provide the input for the AK-algorithm, we need the heap by Asano *et al.* [2]. For completeness, we briefly restate its properties here.

show that both problems are equivalent. The monotone chain can be completed to a polygon by adding a vertex with a very high or low y -coordinate.

⁴Darwish and Elmasry [13] state a running time of $O(n^2/s + n \log n)$, but they measure workspace in bits, while we use words.

⁵AK reduce triangulation to the *next smaller right neighbor* (NSR) and the *next smaller left neighbor* (NSL) problem and present an algorithm for NSR if the input is in x -order. This implies an NSL-algorithm by reading the input in reverse.

Lemma 2.4 ([2]). *Let S be a set of n points. There is a heap that supports insert and extract-min (resp. extract-max) in $O((n/(s \log n) + \log s)D(n))$ time using $O(s)$ variables, where $D(n)$ is the time to decide whether a given element currently resides in the heap (is alive).*⁶

Proof. We first describe the data structure. Then we discuss how to perform insertions and extract-min operations.

We partition the input into $s \log n$ consecutive *buckets* of equal size, and we build a complete binary tree T over the buckets. Let v be a node of T with height h . Then, there are 2^h buckets below v in T . We store $2h$ *information bits* in v to specify the minimum alive element below v . The first h bits identify the bucket containing the minimum. We further divide this bucket into 2^h consecutive parts of equal size, called *quantiles*. The second h bits in v specify the quantile containing the minimum. If $2h > \log n$, we use $\log n$ bits to specify the minimum directly. Hence, the total number of bits is bounded by

$$\sum_{h=0}^{\log(s \log n)} \frac{s \log n}{2^h} \min\{2h, \log n\} = O(s \log n).$$

Therefore we need $O(s)$ variables in total.

Let v be a node with height h . To find the minimum alive element in T below v , we use the $2h$ information bits stored in v . First, we identify the bucket containing the minimum and the correct quantile within this bucket. This quantile contains $O(\frac{n}{2^h s \log n})$ elements. For each element in the quantile, we decide in $D(n)$ time whether it is alive, and we return the minimum such element. This takes $O(\frac{n}{2^h s \log n} D(n))$ time in total.

insert: Assume we want to insert an element x that is at position i in the input array. Let v be the parent of the leaf of T corresponding to the bucket that contains x . We update the information bits at each node u on the root path starting at v . To do so, we use the information bits in u to find the minimum element in the buckets covered by u , as described above. Then we compare it with x . If x is larger, we are done and we stop the insertion. Otherwise, we update the information bits at u to the bucket and quantile that contain x . If we reach and update the root node, we also update the pointer that points to the minimum element in the heap. The work per node is dominated by the costs for finding the minimum, which is $O(\frac{n}{2^h s \log n} D(n))$. Thus, the total cost for insertion is bounded by

$$\sum_{h=0}^{\log(s \log n)} \frac{n}{2^h s \log n} D(n) = O\left(\frac{n}{s \log n} D(n)\right).$$

⁶The bounds in [2] do not include the factor $D(n)$ since the authors studied a setting similar to Lemma 2.5 where it takes $O(1)$ time to decide whether an element is alive.

extract-min: First we use the pointer to the minimum alive element to determine the element x to return. Then we use a similar update strategy as for insertions. Let v be the leaf node corresponding to the bucket of x . We first update the information bits of v by scanning through the whole bucket of v and determining the smallest alive element. Since a bucket contains $O(n/s \log n)$ elements, this needs time $O(n/(s \log n)D(n))$. Then we update the information bits of each node u on the path for v as follows: let v_1 and v_2 be the two children of u . We determine the minimum alive element in the buckets covered by v_1 and v_2 , take the smaller one, and use it to update the information bits at u . Once we reach the root, we also update the pointer to the minimum element of the heap to the new minimum element of the root. The total time again is bounded by $O(\frac{n}{s \log n}D(n))$.

□

Lemma 2.5 ([2]). *Let S be a set of n points. We can build a heap with all elements in S in $O(n)$ time that supports extract-min in $O(n/(s \log n) + \log n)$ time using $O(s)$ variables.*

Proof. The construction time is given in [2]. To decide in $O(1)$ time if some $x \in S$ is alive, we store the last extracted minimum m and test whether $x > m$. □

We now present the complete algorithm. We show how to subdivide S into mountains F_i and how to run the AK-algorithm on each F_i^\uparrow . By reversing the order, the same discussion applies to F_i^\downarrow . Sorted input is emulated by two heaps H_1, H_2 for S according to x -order. By Lemma 2.5, each heap uses $O(s)$ space, can be constructed in $O(n)$ time, and supports extract-min in $O(n/(s \log n) + \log n)$ worst-case time. We will use H_1 to determine the size of the next mountain F_i and H_2 to process the points of F_i .

We execute the convex hull algorithm with $\Theta(s)$ space until it reports the next convex hull edge pq . Throughout the execution of the algorithm, heaps H_1 and H_2 contain exactly the points to the right of p . We repeatedly extract the minimum of H_1 until q becomes the minimum element. Let k be the number of removed points.

If $k = 1$, then pq is short. We extract the minimum of H_2 , and we continue with the convex hull algorithm. If $k \geq 2$, then Lemma 2.2 shows that pq is the base of a mountain F that consists of all points between p and q . These are exactly the $k + 1$ smallest elements in H_2 (including p and q). If $k \leq s$, we extract them from H_2 , and we triangulate F in memory. If $k > s$, we execute the AK-algorithm on F using $O(s)$ variables. At the beginning of the i th round, we create a copy $H_{(i)}$ of H_2 , i.e., we duplicate the $O(s)$ variables that determine the state of H_2 . Further, we create an empty max-heap $H_{(ii)}$ using $O(s)$ variables to provide input for the second scan. To be able to reread a sub-block, we create a further copy $H'_{(i)}$ of H_2 . Whenever the AK-algorithm requests the next point in the first scan, we simply extract the minimum of $H_{(i)}$. When a sub-block is fully read, we use $H'_{(i)}$ to reread the elements and insert them into $H_{(ii)}$. Now, the

rightmost element of all active sub-blocks corresponds exactly to the maximum of $H_{(ii)}$. One step in the second scan is equivalent to an extract-max on $H_{(ii)}$.

At the end of a round, we delete $H_{(i)}$, $H'_{(i)}$, and $H_{(ii)}$, so that the space can be reused in the next round. Once the AK-algorithm finishes, we repeatedly extract the minimum of H_2 until we reach q .

Theorem 2.6. *We can report the edges of a triangulation of a set S of n points in time $O(n^2/s + n \log n \log s)$ using $O(s)$ additional variables.*

Proof. Similarly as before, correctness directly follows from Lemma 2.2 and the correctness of the AK-algorithm. The bound on the space usage is immediate.

Computing the convex hull now needs $O(n^2/(s \log n) + n \log n)$ time [13]. By Lemma 2.5, the heaps H_1 and H_2 can be constructed in $O(n)$ time. During execution, we perform n extract-min operations on each heap, requiring $O(n^2/(s \log n) + n \log n)$ time in total.

Let F_j be a mountain with n_j vertices that is discovered by the convex hull algorithm. If $n_j \leq s$, then F_j is triangulated in memory in $O(n_j)$ time, and the total time for such mountains is $O(n)$. If $n_j > s$, then the AK-algorithm runs in $O(n_j \log_s n_j)$ time. We must also account for providing the input for the algorithm. For this, consider some round $i \geq 1$. We copy H_2 to $H_{(i)}$ in $O(s)$ time. This time can be charged to the first scan, since $n_j > s$. Furthermore, we perform n_j extract-min operations on $H_{(i)}$. Hence the total time to provide input for the first scan is $O(n_j n / (s \log n) + n_j \log n)$.

For the second scan, we create another copy $H'_{(i)}$ of H_2 . Again, the time for this can be charged to the scan. Also, we perform n_j extract-min operations on $H'_{(i)}$ which takes $O(n_j n / (s \log n) + n_j \log n)$ time. Additionally, we insert each fully-read block into $H_{(ii)}$. The main problem is to determine if an element in $H_{(ii)}$ is alive: there are at most $O(s)$ active sub-blocks. For each active sub-block b_i , we know the first element y_i and the element z_i that l_i points to. An element is alive if and only if it is in the interval $[y_i, z_i]$ for some active b_i . This can be checked in $O(\log s)$ time. Thus, by Lemma 2.4, each insert and extract-max on $H_{(ii)}$ takes $O((n/(s \log n) + \log s) \log s)$ time. Since each element is inserted once, the total time to provide input to the second scan is $O(n_j \log(s)(n/(s \log n) + \log s))$. This dominates the time for the first scan. There are $O(\log_s n_j)$ rounds, so we can triangulate F_j in time $O(n_j \log_s n_j + n_j \log(n_j)(n/(s \log n) + \log s))$. Summing over all F_j , the total time is $O(n^2/s + n \log n \log s)$. \square

3. Voronoi Diagrams

Given a planar n -point set S , we would like to find the vertices of $\text{VD}(S)$. Let $K = \{p_1, p_2, p_3\}$ be a triangle with $S \cap K = \emptyset$, $S \subseteq \text{conv}(K)$, and so that all vertices of $\text{VD}(S)$ are vertices of $\text{VD}(S \cup K)$. For example, we can set $K = \{(-\kappa, -\kappa), (-\kappa, \kappa), (0, \kappa)\}$ for some large $\kappa > 0$. Since the desired properties hold for all large enough κ , we do not need to find an explicit value for it. Instead, whenever we want to evaluate a predicate involving points from K , we can take the result obtained for $\kappa \rightarrow \infty$.

Our algorithm relies on random sampling. First, we show how to take a random sample from S with small workspace. One of many possible approaches is the following one that ensures a worst-case guarantee:

Lemma 3.1. *We can sample a uniform random subset $R \subseteq S$ of size s in time $O(n + s \log s)$ and space $O(s)$.*

Proof. The sampling algorithm consists of two phases. In the first phase, we sample a random sequence I of s distinct numbers from $[n]$.⁷ The phase proceeds in s rounds. At the beginning of round k , for $k = 1, \dots, s$, we have already sampled a sequence I of $k - 1$ numbers from $[n]$, and we would like to pick an element from $[n] \setminus I$ uniformly at random. We store I in a binary search tree T . We maintain the invariant that T stores with each element $x \in [n - k + 1] \cap I$ a replacement $\rho_x \in \{n - k + 2, \dots, n\} \setminus I$ such that $[n] \setminus I = ([n - k + 1] \setminus I) \cup \{\rho_x \mid x \in [n - k + 1] \cap I\}$, see Figure 2. In round k , we sample a random number x




Figure 2: Sampling a random sequence I from $[n]$. At the beginning of round k , we have already sampled $k - 1$ elements (shown in gray). Each element $x \in [n - k + 1] \cap I$ has a replacement $\rho_x \in \{n - k + 2, \dots, n\} \setminus I$ (indicated by the arrows). In round k , we pick a random number $x \in [n - k + 1]$. If x is already contained in I , we add ρ_x to I . Otherwise, we add x .

from $[n - k + 1]$, and we check in T whether $x \in I$. If not, we add x to I (and T), otherwise, we add ρ_x to I (and T). By the invariant, we add a uniform random element from $[n] \setminus I$ to I .

It remains to update the replacements, see Figure 3. If $x = n - k + 1$, we do not need a replacement for x . Now suppose $x < n - k + 1$. If $n - k + 1 \notin I$, we set $\rho_x = n - k + 1$. Otherwise, we set $\rho_x = \rho_{n-k+1}$. This ensures that the invariant holds at the beginning of round $k + 1$, and it takes $O(\log s)$ time and $O(s)$ space. We continue for s rounds. At the end of the first phase, any sequence of s distinct numbers in $[n]$ is sampled with equal probability. Furthermore, the phase takes $O(s \log s)$ time and $O(s)$ space.

In the second phase, we scan through S to obtain the elements whose positions correspond to the numbers in I . This requires $O(n)$ time and $O(s)$ space. \square

We use Lemma 3.1 to find a random sample $R \subseteq S$ of size s . We compute $\text{VD}(R \cup K)$, triangulate the bounded cells and construct a planar point location structure for the triangulation. This takes $O(s \log s)$ time and $O(s)$ space [18]. By our choice of K , all Voronoi cells for points in R are bounded, and every

⁷We write $[n]$ for the set $\{1, \dots, n\}$.




Figure 3: Finding a replacement for x . If $x = n - k + 1$, we do not need a replacement for x in the next round (top left). If $n - k + 1$ is not sampled yet, we can make it the replacement for x (top right). Otherwise, we make the old replacement for $n - k + 1$ the new replacement for x (bottom).

point in S lies in a bounded Voronoi cell. Given a vertex $v \in \text{VD}(R \cup K)$, the *conflict circle* of v is the largest circle with center v and no point from $R \cup K$ in its interior. The *conflict set* B_v of v contains all points from S that lie in the conflict circle of v , and the *conflict size* b_v of v is $|B_v|$. We scan through S to find the conflict size b_v for each vertex $v \in \text{VD}(R \cup K)$: every Voronoi vertex has a counter that is initially 0. For each $p \in S \setminus (R \cup K)$, we use the point location structure to find the triangle Δ of $\text{VD}(R \cup K)$ that contains it. At least one vertex v of Δ is in conflict with p . Starting from v , we walk along the edges of $\text{VD}(R \cup K)$ to find all Voronoi vertices in conflict with p (recall that these vertices induce a connected component in $\text{VD}(R \cup K)$). We increment the counters of all these vertices. This may take a long time in the worst case, so we impose an upper bound on the total work. For this, we choose a *threshold* M . When the sum of the conflict counters exceeds M , we start over with a new sample R . The total time for one attempt is $O(n \log s + M)$, and below we prove that for $M = \Theta(n)$, the success probability is at least $3/4$. Next, we pick another threshold T , and we compute for each vertex v of $\text{VD}(R \cup K)$ the *excess* $t_v = b_v s / n$. The excess measures how far the vertex deviates from the desired conflict size n/s . We check if $\sum_{v \in \text{VD}(R \cup K)} t_v \log t_v \leq T$. If not, we start over with a new sample. Below, we prove that for $T = \Theta(s)$, the success probability is at least $3/4$. The total success probability is $1/2$, and the expected number of attempts is 2. Thus, in expected time $O(n \log s + s \log s)$, we can find a sample $R \subseteq S$ with $\sum_{v \in \text{VD}(R \cup K)} b_v = O(n)$ and $\sum_{v \in \text{VD}(R \cup K)} t_v \log t_v = O(s)$.

We now analyze the success probabilities, using the classic Clarkson-Shor method [12]. We begin with a variant of the Chazelle-Friedman bound [11].

Lemma 3.2. *Let X be a planar point set of size m , and let $Y \subset \mathbb{R}^2$ with $|Y| \leq 3$ and $X \cap Y = \emptyset$. For fixed $p \in (0, 1]$, let $R \subseteq X$ be a random subset of size pm and let $R' \subseteq X$ be a random subset of size $p'm$, for $p' = p/2$. Suppose that $p'm \geq 4$. Fix $\mathbf{u} \subset X \cup Y$ with $|\mathbf{u}| = 3$, and let $v_{\mathbf{u}}$ be the Voronoi vertex defined by \mathbf{u} . Let $b_{\mathbf{u}}$ be the number of points from $X \cup Y$ in the interior of the circle with center $v_{\mathbf{u}}$ and with the points from \mathbf{u} on the boundary. Then,*

$$\Pr[v_{\mathbf{u}} \in \text{VD}(R \cup Y)] \leq 64e^{-pb_{\mathbf{u}}/2} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup Y)].$$

Proof. Let $\sigma = \Pr[v_{\mathbf{u}} \in \text{VD}(R \cup Y)]$ and $\sigma' = \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup Y)]$. The vertex $v_{\mathbf{u}}$ is in $\text{VD}(R \cup Y)$ precisely if $\mathbf{u} \subseteq R \cup Y$ and $B_{\mathbf{u}} \cap (R \cup Y) = \emptyset$, where $B_{\mathbf{u}}$ are the points from $X \cup Y$ inside the circle with center $v_{\mathbf{u}}$ and with the points from \mathbf{u} on the boundary. If $B_{\mathbf{u}} \cap Y \neq \emptyset$, then $\sigma = \sigma' = 0$, and the lemma holds. Thus, assume that $B_{\mathbf{u}} \subseteq X$. Let $d_{\mathbf{u}} = |\mathbf{u} \cap X|$, the number of points in \mathbf{u} from X . There are $\binom{m-b_{\mathbf{u}}-d_{\mathbf{u}}}{pm-d_{\mathbf{u}}}$ ways to choose a pm -subset from X that avoids all points in $B_{\mathbf{u}}$ and contains all points of $\mathbf{u} \cap X$, so

$$\begin{aligned} \sigma &= \binom{m-b_{\mathbf{u}}-d_{\mathbf{u}}}{pm-d_{\mathbf{u}}} \bigg/ \binom{m}{pm} \\ &= \frac{\prod_{j=0}^{pm-d_{\mathbf{u}}-1} (m-b_{\mathbf{u}}-d_{\mathbf{u}}-j)}{\prod_{j=0}^{pm-d_{\mathbf{u}}-1} (pm-d_{\mathbf{u}}-j)} \bigg/ \frac{\prod_{j=0}^{pm-1} (m-j)}{\prod_{j=0}^{pm-1} (pm-j)} \\ &= \prod_{j=0}^{d_{\mathbf{u}}-1} \frac{pm-j}{m-j} \cdot \prod_{j=0}^{pm-d_{\mathbf{u}}-1} \frac{m-b_{\mathbf{u}}-d_{\mathbf{u}}-j}{m-d_{\mathbf{u}}-j} \\ &\leq p^{d_{\mathbf{u}}} \prod_{j=0}^{pm-d_{\mathbf{u}}-1} \left(1 - \frac{b_{\mathbf{u}}}{m-d_{\mathbf{u}}-j}\right). \end{aligned}$$

Similarly, we get

$$\sigma' = \prod_{i=0}^{d_{\mathbf{u}}-1} \frac{p'm-i}{m-i} \prod_{j=0}^{p'm-d_{\mathbf{u}}-1} \left(1 - \frac{b_{\mathbf{u}}}{m-d_{\mathbf{u}}-j}\right),$$

and since $p'm \geq 4$ and $i \leq 2$, it follows that

$$\sigma' \geq \left(\frac{p'}{2}\right)^{d_{\mathbf{u}}} \prod_{j=0}^{p'm-d_{\mathbf{u}}-1} \left(1 - \frac{b_{\mathbf{u}}}{m-d_{\mathbf{u}}-j}\right).$$

Therefore, since $p' = p/2$,

$$\frac{\sigma}{\sigma'} \leq \left(\frac{2p}{p'}\right)^{d_{\mathbf{u}}} \prod_{j=p'm-d_{\mathbf{u}}}^{pm-d_{\mathbf{u}}-1} \left(1 - \frac{b_{\mathbf{u}}}{m-d_{\mathbf{u}}-j}\right) \leq 64 \left(1 - \frac{b_{\mathbf{u}}}{m}\right)^{pm/2} \leq 64e^{-pb_{\mathbf{u}}/2}.$$

□

We can now bound the total expected conflict size.

Lemma 3.3. *We have $\mathbf{E} \left[\sum_{v \in \text{VD}(R \cup K)} b_v \right] = O(n)$.*

Proof. By expanding the expectation, we get

$$\mathbf{E} \left[\sum_{v \in \text{VD}(R \cup K)} b_v \right] = \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} \Pr[v_{\mathbf{u}} \in \text{VD}(R \cup K)] b_{\mathbf{u}},$$

with $v_{\mathbf{u}}$ being the Voronoi vertex of \mathbf{u} and $b_{\mathbf{u}}$ its conflict size. By Lemma 3.2 with $X = S$, $Y = K$ and $p = s/n$, this is

$$\leq \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} 64e^{-pb_{\mathbf{u}}/2} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] b_{\mathbf{u}},$$

where $R' \subseteq S$ is a sample of size $s/2$. We bound this as

$$\begin{aligned} &\leq \sum_{i=0}^{\infty} \sum_{\substack{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3 \\ b_{\mathbf{u}} \in [\frac{i}{p}, \frac{i+1}{p})}} \frac{64e^{-i/2}(i+1)}{p} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] \\ &\leq \frac{1}{p} \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] \sum_{i=0}^{\infty} 64e^{-i/2}(i+1) \\ &= O(s/p) = O(n), \end{aligned}$$

since $\sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] = O(s)$ is the size of $\text{VD}(R' \cup K)$ and $\sum_{i=0}^{\infty} e^{-i/2}(i+1) = O(1)$. \square

By Lemma 3.3 and Markov's inequality, it follows that there is an $M = \Theta(n)$ with $\Pr[\sum_{v \in \text{VD}(R \cup K)} b_v > M] \leq 1/4$.

Lemma 3.4. $\mathbf{E} \left[\sum_{v \in \text{VD}(R \cup K)} t_v \log t_v \right] = O(s)$.

Proof. By Lemma 3.2 with $X = S$, $Y = K$, and $p = s/n$,

$$\begin{aligned} \mathbf{E} \left[\sum_{v \in \text{VD}(R \cup K)} t_v \log t_v \right] &= \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} \Pr[v_{\mathbf{u}} \in \text{VD}(R \cup K)] t_{\mathbf{u}} \log t_{\mathbf{u}} \\ &\leq \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} 64e^{-pb_{\mathbf{u}}/2} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] t_{\mathbf{u}} \log t_{\mathbf{u}} \\ &\leq \sum_{i=0}^{\infty} \sum_{\substack{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3 \\ b_{\mathbf{u}} \in [\frac{i}{p}, \frac{i+1}{p})}} 64e^{-\frac{i}{2}}(i+1)^2 \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] \\ &\leq \sum_{i=0}^{\infty} 64e^{-\frac{i}{2}}(i+1)^2 \sum_{\mathbf{u} \subset S \cup K, |\mathbf{u}|=3} \Pr[v_{\mathbf{u}} \in \text{VD}(R' \cup K)] \\ &= O(s). \end{aligned}$$

□

By Markov's inequality and Lemma 3.4, we can conclude that there is a $T = \Theta(s)$ with $\Pr[\sum_{v \in \text{VD}(R \cup K)} t_v \log t_v \geq T] \leq 1/4$. This finishes the first sampling phase.

The next goal is to sample for each vertex v with $t_v \geq 2$ a random subset $R_v \subseteq B_v$ of size $\min\{\alpha t_v \log t_v, b_v\}$ for large enough $\alpha > 0$ (recall that B_v is the conflict set of v and that $b_v = |B_v|$).

Lemma 3.5. *In total time $O(n \log s)$, we can sample for each vertex $v \in \text{VD}(R \cup K)$ with $t_v \geq 2$ a random subset $R_v \subseteq B_v$ of size $\min\{\alpha t_v \log t_v, b_v\}$.*

Proof. First, we sample for each vertex v with $t_v \geq 2$ a sequence I_v of $\alpha t_v \log t_v$ distinct numbers from $\{1, \dots, b_v\}$. For this, we use the first phase of the algorithm from the proof of Lemma 3.1 for each such vertex, but *without* reusing the space. As explained in the proof of Lemma 3.1, this takes total time

$$O\left(\sum_v (t_v \log t_v) \log(t_v \log t_v)\right) = O\left(\sum_v (t_v \log t_v) \log s\right) = O(s \log s),$$

since $\sum_v t_v \log t_v = O(s)$, and in particular $t_v \log t_v = O(s)$ for each vertex v (note that the constant in the O-notation is independent of v). Also, since $\sum_v t_v \log t_v = O(s)$, the total space requirement is $O(s)$.

After that, we scan through S . For each vertex v , we have a counter c_v , initialized to 0. For each $p \in S$, we find the conflict vertices of p , and for each conflict vertex v , we increment c_v . If c_v appears in the corresponding set I_v , we add p to R_v . The total running time is $O(n \log s)$, as we do one point location for each input point and the total conflict size is $O(n)$. □

We next show that for a *fixed* vertex $v \in \text{VD}(R \cup K)$, with constant probability, all vertices in $\text{VD}(R_v)$ have conflict size n/s with respect to B_v .

Lemma 3.6. *Let $v \in \text{VD}(R \cup K)$ with $t_v \geq 2$, and let $R_v \subseteq B_v$ be the sample for v . The expected number of vertices v' in $\text{VD}(R_v)$ with at least n/s points from B_v in their conflict circle is at most $1/4$.*

Proof. If $R_v = B_v$, the lemma holds, so assume that $\alpha t_v \log t_v < b_v$. Recall that $t_v = b_v s/n$. We have

$$\mathbf{E} \left[\sum_{\substack{v' \in \text{VD}(R_v) \\ b'_{v'} \geq n/s}} 1 \right] = \sum_{\substack{\mathbf{u} \subseteq B_v, |\mathbf{u}|=3 \\ b'_{\mathbf{u}} \geq n/s}} \Pr[v'_{\mathbf{u}} \in \text{VD}(R_v)],$$

where b'_u denotes the number of points from B_v inside the circle with center v'_u and with the points from u on the boundary. Using Lemma 3.2 with $X = B_v$, $Y = \emptyset$, and $p = (\alpha t_v \log t_v)/b_v = \alpha(s/n) \log t_v$, this is

$$\begin{aligned}
&\leq \sum_{\substack{u \subset B_v, |u|=3 \\ b'_u \geq n/s}} 64e^{-pb'_u/2} \Pr[v'_u \in \text{VD}(R'_v)] \\
&\leq 64e^{-(\alpha/2) \log t_v} \sum_{u \subset B_v, |u|=3} \Pr[v'_u \in \text{VD}(R'_v)] \\
&= O(t_v^{-\alpha/2} t_v \log t_v) \leq 1/4,
\end{aligned}$$

for α large enough (remember that $t_v \geq 2$). □




Figure 4: A Voronoi Diagram of the sampled set R (left). The two red square vertices of $\text{VD}(R \cup K)$ are not good and we need to resample within their conflict list (the blue crosses) and compute the new Voronoi Diagram (right).

By Lemma 3.6 and Markov's inequality, the probability that all vertices from $\text{VD}(R_v)$ have at most n/s points from B_v in their conflict circles is at least $3/4$. If so, we call v *good*, see Figure 4. Scanning through S , we can identify the good vertices in time $O(n \log s)$ and space $O(s)$. Let s' be the size of $\text{VD}(R \cup K)$. If we have less than $s'/2$ good vertices, we repeat the process. Since the expected number of good vertices is $3s'/4$, the probability that there are at least $s'/2$ good vertices is at least $1/2$, by Markov's inequality. Thus, in expectation, we need to perform the sampling twice. For the remaining vertices, we repeat the process, but now we take two samples per vertex, decreasing the failure probability to $1/4$. We repeat the process, taking in each round the maximum number of samples that fit into the work space. In general, if we have s'/a_i active vertices in round i , we can take a_i samples per vertex, resulting in a failure probability of 2^{-a_i} . Thus, the expected number of active vertices in round $i+1$ is $s'/a_{i+1} = s'/(a_i 2^{a_i})$. After $O(\log^* s)$ rounds, all vertices are good. To summarize:

Lemma 3.7. *In total expected time $O(n \log s \log^* s)$ and space $O(s)$, we can find sets $R \subseteq S$ and $R_v \subset B_v$ for each vertex $v \in \text{VD}(R \cup K)$ such that (i) $|R| = s$; (ii) $\sum_{v \in \text{VD}(R \cup K)} |R_v| = O(s)$; and (iii) for every R_v , all vertices of $\text{VD}(R_v)$ have at most n/s points from B_v in their conflict circle.*

We set $R_2 = R \cup \bigcup_{v \in \text{VD}(R \cup K)} R_v$. By Lemma 3.7, $|R_2| = O(s)$. We compute $\text{VD}(R_2 \cup K)$ and triangulate its bounded cells. For a triangle Δ of the triangulation, let $r \in R_2 \cup K$ be the site whose cell contains Δ , and v_1, v_2, v_3 the vertices of Δ . We set $B_\Delta = \{r\} \cup \bigcup_{i=1}^3 B_{v_i}$. Using the next lemma, we show that $|B_\Delta| = O(n/s)$.

Lemma 3.8. *Let $S \subset \mathbb{R}^2$ and $\Delta = \{v_1, v_2, v_3\}$ a triangle in the triangulation of $\text{VD}(S)$. Let $x \in \Delta$. Then any circle C with center x that contains no points from S is covered by the conflict circles of v_1, v_2 and v_3 .*

Proof. Let $p \in C$ and let $r \in S$ be the site whose cell contains Δ . We show that p is contained in the conflict circle of v_1, v_2 , or v_3 . Consider the bisector B of p and r . Since C contains p but not r , we have $d(x, p) < d(x, r)$, so x lies on the same side of B as p . Since $x \in \Delta$, at least one of v_1, v_2, v_3 , is on the same side of B as p ; say v_1 . This means that $d(v_1, p) < d(v_1, r)$, so p lies inside the circle around v_1 with r on the boundary. This is precisely the conflict circle of v_1 . \square

Lemma 3.9. *Any triangle Δ in the triangulation of $\text{VD}(R_2 \cup K)$ has $|B_\Delta| = O(n/s)$.*

Proof. Let v be a vertex of Δ . We show that $b_v = O(n/s)$. Let $\Delta_R = \{v_1, v_2, v_3\}$ be the triangle in the triangulation of $\text{VD}(R)$ that contains v . By Lemma 3.8, we have $B_v \subseteq \bigcup_{i=1}^3 B_{v_i}$. We consider the intersections $B_v \cap B_{v_i}$, for $i = 1, 2, 3$. If $t_{v_i} < 2$, then $b_{v_i} = O(n/s)$ and $|B_v \cap B_{v_i}| = O(n/s)$. Otherwise, we have sampled a set R_{v_i} for v_i . Let $\Delta_i = \{w_1, w_2, w_3\}$ be the triangle in the triangulation of $\text{VD}(R_{v_i})$ that contains v . Again, by Lemma 3.8, we have $B_v \subseteq \bigcup_{j=1}^3 B_{w_j}$ and thus also $B_v \cap B_{v_i} \subseteq \bigcup_{j=1}^3 B_{w_j} \cap B_{v_i}$. However, by construction of R_{v_i} , $|B_{w_j} \cap B_{v_i}|$ is at most n/s for $j = 1, 2, 3$. Hence, $|B_v \cap B_{v_i}| = O(n/s)$ and $b_v = O(n/s)$. \square

The following lemma enables us to compute the Voronoi diagram of $R_2 \cup K$ locally for each triangle Δ in the triangulation of $\text{VD}(R_2 \cup K)$ by only considering sites in B_Δ . It is a direct consequence of Lemma 3.8.

Lemma 3.10. *For every triangle Δ in the triangulation of $\text{VD}(R_2 \cup K)$, we have $\text{VD}(S \cup K) \cap \Delta = \text{VD}(B_\Delta) \cap \Delta$.*

Theorem 3.11. *Let S be a planar n -point set. In expected time $O((n^2/s) \log s + n \log s \log^* s)$ and space $O(s)$, we can compute all Voronoi vertices of S .*

Proof. We compute a set R_2 as above. This takes $O(n \log s \log^* s)$ time and space $O(s)$. We triangulate the bounded cells of $\text{VD}(R_2 \cup K)$ and compute a point location structure for the result. Since there are $O(s)$ triangles, we can store the resulting triangulation in the workspace. Now, the goal is to compute simultaneously for all triangles Δ the Voronoi diagram $\text{VD}(B_\Delta)$ and to output all Voronoi vertices that lie in Δ and are defined by points from S . By Lemma 3.10, this gives all Voronoi vertices of $\text{VD}(S)$.

Given a planar m -point set X , the algorithm by Asano et al. finds all vertices of $\text{VD}(X)$ in $O(m)$ scans over the input, with constant workspace [4]. We can

perform a simultaneous scan for all sets B_Δ by determining for each point in S all sets B_Δ that contain it. This takes total time $O(n \log s)$, since we need one point location for each $p \in S$ and since the total size of the B_Δ 's is $O(n)$. We need $O(\max_\Delta |B_\Delta|) = O(n/s)$ such scans, so the second part of the algorithm needs $O((n^2/s) \log s)$ time. \square

As mentioned in the introduction, Theorem 3.11 also lets us report all edges of the Delaunay triangulation of S in the same time bound: by duality, the three sites that define a vertex of $\text{VD}(S)$ also define a triangle for the Delaunay triangulation. Thus, whenever we discover a vertex of $\text{VD}(S)$, we can instead output the corresponding Delaunay edges, while using a consistent tie-breaking rule to make sure that every edge is reported only once.

Acknowledgments. This work began while W. Mulzer, P. Seiferth, and Y. Stein visited the Tokuyama Laboratory at Tohoku University. We would like to thank Takeshi Tokuyama and all members of the lab for their hospitality and for creating a conducive and stimulating research environment.

- [1] T. Asano, K. Buchin, M. Buchin, M. Korman, W. Mulzer, G. Rote, and A. Schulz. Memory-constrained algorithms for simple polygons. *Comput. Geom.*, 46(8):959–969, 2013.
- [2] T. Asano, A. Elmasry, and J. Katajainen. Priority queues and sorting for read-only data. In *Theory and applications of models of computation (TAMC)*, pages 32–41, 2013.
- [3] T. Asano and D. Kirkpatrick. Time-space tradeoffs for all-nearest-larger-neighbors problems. In *Proc. 13th Int. Symp. Algorithms and Data Structures (WADS)*, pages 61–72, 2013.
- [4] T. Asano, W. Mulzer, G. Rote, and Y. Wang. Constant-work-space algorithms for geometric problems. *J. of Comput. Geom.*, 2(1):46–68, 2011.
- [5] L. Barba, M. Korman, S. Langerman, K. Sadakane, and R. I. Silveira. Space–time trade-offs for stack-based algorithms. *Algorithmica*, 72(4):1097–1129, 2015.
- [6] L. Barba, M. Korman, S. Langerman, and R. I. Silveira. Computing the visibility polygon using few variables. *Comput. Geom.*, 47(9):918–926, 2013.
- [7] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational geometry. Algorithms and applications*. Springer-Verlag, third edition, 2008.
- [8] A. Borodin and S. Cook. A time-space tradeoff for sorting on a general sequential model of computation. *SIAM J. Comput.*, 11:287–297, 1982.
- [9] H. Brönnimann, T. M. Chan, and E. Y. Chen. Towards in-place geometric algorithms and data structures. In *Proc. 20th Annu. Sympos. Comput. Geom. (SoCG)*, pages 239–246, 2004.

- [10] T. M. Chan and E. Y. Chen. Multi-pass geometric algorithms. *Discrete Comput. Geom.*, 37(1):79–102, 2007.
- [11] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- [12] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [13] O. Darwish and A. Elmasry. Optimal time-space tradeoff for the 2D convex-hull problem. In *Proc. 22nd Annu. European Sympos. Algorithms (ESA)*, pages 284–295, 2014.
- [14] H. Djidjev and A. Lingas. On computing Voronoi diagrams for sorted point sets. *Internat. J. Comput. Geom. Appl.*, 5(3):327–337, 1995.
- [15] A. Fournier and D. Y. Montuno. Triangulating simple polygons and equivalent problems. *ACM Transactions on Graphics*, 3:153–174, 1984.
- [16] S. Har-Peled. Shortest path in a polygon using sublinear space. *J. of Comput. Geom.*, 7(2):19–45, 2016.
- [17] R. Jarvis. On the identification of the convex hull of a finite set of points in the plane. *Inform. Process. Lett.*, 2(1):18–21, 1973.
- [18] D. Kirkpatrick. Optimal search in planar subdivisions. *SIAM J. Comput.*, 12(1):28–35, 1983.
- [19] J. I. Munro and M. Paterson. Selection and sorting with limited storage. *Theoret. Comput. Sci.*, 12:315–323, 1980.
- [20] J. I. Munro and V. Raman. Selection from read-only memory and sorting with minimum data movement. *Theoret. Comput. Sci.*, 165(2):311–323, 1996.
- [21] J. Pagter and T. Rauhe. Optimal time-space trade-offs for sorting. In *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 264–268, 1998.
- [22] I. Pohl. A minimum storage algorithm for computing the median. Technical Report RC2701, IBM, 1969.
- [23] J. E. Savage. *Models of computation—exploring the power of computing*. Addison-Wesley, 1998.

Appendix A. The Asano-Kirkpatrick Algorithm

We give more details on the algorithm of Asano and Kirkpatrick [3]. Let F be a mountain with vertices q_1, \dots, q_n sorted in x -order and base q_1q_n . We define the *height* $h(q_i)$ of q_i , $i = 1, \dots, n$, as the distance from q_i to the line through the base. Let $A = (q_1, \dots, q_n)$ be the input array. A vertex q_r is the *nearest-smaller-right-neighbor* (NSR) of a vertex q_l if (i) $l < r$; (ii) $h(q_l) > h(q_r)$; and (iii) $h(q_l) \leq h(q_k)$ for $l < k < r$. We call (q_l, q_r) a *NSR-pair*, with *left endpoint* q_l and *right endpoint* q_r . *Nearest-smaller-left-neighbors* (NSL) and NSL-pairs are defined similarly. Let R be the set of all NSR-pairs and L be the set of all NSL pairs. Asano and Kirkpatrick show that the edges $R \cup L$ triangulate F . We describe the algorithm for computing R . The algorithm for L is the same, but it reads the input in reverse.

Let s denote the space parameter. The algorithm runs in $\log_s n$ rounds. In round i , $i = 0, \dots, \log_s n - 1$, we partition A into s^i consecutive *blocks* of size n/s^i . Each block B is further partitioned into s consecutive *sub-blocks* b_1, \dots, b_s of size n/s^{i+1} . In each round, we compute only NSR-pairs with endpoints in different sub-blocks of the same block. We handle each block B individually as follows. The sub-blocks of B are visited from left to right. When we visit a sub-block b_j , we compute all NSR-pairs with a right endpoint in b_j and a left endpoint in the sub-blocks b_1, \dots, b_{j-1} . Initially, we visit the first sub-block b_1 and we push a pointer to the rightmost element in b_1 onto a stack S . We call a sub-block with a pointer in S *active*. Assume now that we have already visited sub-blocks b_1, \dots, b_{j-1} . Let l be the topmost pointer in S , referring to an element q_l in $b_{j'}$, $j' < j$. Furthermore, let r be a pointer to the leftmost element q_r in b_j . If $h(q_l) > h(q_r)$, we output (q_l, q_r) and we decrement l until we find the first element whose height is smaller than the current $h(q_l)$. If l leaves $b_{j'}$, this sub-block becomes inactive and we remove l from S . We continue with the new topmost pointer as our new l . On the other hand, if $h(q_l) \leq h(q_r)$, we increment r by one. We continue until either r leaves b_j or S becomes empty. Then we push a pointer to the rightmost element in b_j onto S and proceed to the next sub-block.

In each round, the algorithm reads the complete input once in x -order. In addition, the algorithm reads at most once each active sub-blocks in reverse order. Note that a sub-block becomes active only once.