

Linear-Time Triangulation of a Simple Polygon Made Easier Via Randomization *

Nancy M. Amato [†]

Michael T. Goodrich [‡]

Edgar A. Ramos [§]

Abstract

We describe a randomized algorithm for computing the trapezoidal decomposition of a simple polygon. Its expected running time is linear in the size of the polygon. By a well-known and simple linear time reduction, this implies a linear time algorithm for triangulating a simple polygon. Our algorithm is considerably simpler than Chazelle's (1991) celebrated optimal deterministic algorithm and, hence, positively answers his question of whether a simpler randomized algorithm for the problem exists. The new algorithm can be viewed as a combination of Chazelle's algorithm and of non-optimal randomized algorithms due to Clarkson *et al.* (1991) and to Seidel (1991), with the essential innovation that sampling is performed on subchains of the initial polygonal chain, rather than on its edges. It is also essential, as in Chazelle's algorithm, to include a bottom-up preprocessing phase previous to the top-down construction phase.

1 Introduction

Polygon triangulation is a classic problem in computational geometry, as it was one of the first problems studied in the field [15]. Furthermore, there are sev-




Figure 1: A triangulated simple polygon

eral other problems in computational geometry dealing with polygons that have efficient solutions that begin with polygon triangulation as a preprocessing step (e.g., see [18, 20]). Thus, there has been considerable interest in finding efficient algorithms for this problem.

1.1 Related Prior Work

Garey *et al.* [15] were the first to provide a non-trivial algorithm for the polygon triangulation problem. Their algorithm runs in $O(n \log n)$ time and is based on an elegant plane-sweeping paradigm. Asano *et al.* [2] show that this bound is in fact optimal for polygons that may contain holes. For simple polygons without holes, the lower bound of Asano *et al.* does not hold, however. This fact, and the importance of the polygon triangulation problem, in turn prompted several researchers to work on methods for beating $O(n \log n)$ time for this problem.

Fournier and Montuno [14] and Chazelle and Incerci [6] showed, even prior to the Asano *et al.* lower bound result, that to triangulate a simple polygon in linear time it is sufficient to produce a trapezoidal decomposition (*trapezoidation*) of a simple polygon. In addition, Yap [31] showed that a similar result holds in a parallel computing model. A trapezoidation is formed by shooting a vertical ray through each vertex of the polygon, stopping each ray as soon as it hits another segment on the polygon. Since this early work

*To appear in *Proceedings of the 16th Annual ACM Symposium on Computational Geometry (SoCG'00)*, June, 2000.

[†]Texas A&M University, College Station, TX. E-mail: amato@cs.tamu.edu

[‡]The Johns Hopkins University, Baltimore, MD. E-mail: goodrich@jhu.edu

[§]Max-Planck-Institut für Informatik, Saarbrücken, Germany. E-mail: ramos@mpi-sb.mpg.de

showing the importance of trapezoidation for triangulation, every published triangulation algorithm has concentrated on improving the running time of producing a trapezoidation of a simple polygon. For example, Tarjan and Van Wyk [30] and Kirkpatrick *et al.* [21] showed that the trapezoidation step can be performed in $O(n \log \log n)$ time, resulting in a similar running time for the polygon triangulation problem. Using randomization, Clarkson *et al.* [10], Clarkson *et al.* [8, 7], and Seidel [28] gave simple randomized algorithms that run in $O(n \log^* n)$ expected time. Finally, in a much celebrated and anticipated result, Chazelle [4] showed that one could, in fact, triangulate a polygon in linear time. Goodrich [16] subsequently showed that a similar result can be proven for a parallel computation model. Unfortunately, the trapezoidation methods utilized by these optimal deterministic algorithms are quite complex. Indeed, this conceptual complexity has led many researchers, including Chazelle [4] himself, to ask whether there is a simple randomized algorithm for triangulating a polygon in linear time. To our knowledge, no simple linear-time randomized algorithm has been presented previously.

1.2 Our Results

We describe a randomized algorithm for computing the trapezoidation of a simple polygon. The expected running time of our algorithm is linear in the size of the polygon. As already mentioned, from the trapezoidation, a triangulation of the polygon can be obtained in linear time using well-known methods [6, 14]. Thus, our algorithm provides a randomized algorithm for polygon triangulation that runs in linear expected time. In addition, our algorithm is considerably simpler than Chazelle’s celebrated optimal deterministic algorithm; hence, it addresses the open problem posed by Chazelle and others as to the existence of a simple randomized triangulation algorithm that runs in linear expected time.

The general approach of our algorithm for computing a trapezoidation of a simple polygon P follows that of the non-optimal randomized algorithms by Clarkson *et al.* [8, 7] and Seidel [28]. That is, we compute the trapezoidation of a successively finer sample from P , using an algorithm for arbitrary edges (thus with nonlinear running time), in $O(\log^* n)$ rounds. The fact that the edges come from a simple polygonal chain is used to efficiently perform the computation of the *conflict lists* of the trapezoidation of the sample, once in each round, by *walking* along the original polygonal chain in the trapezoidation. Unfortunately, an approach that maintains the lists of

edge conflicts for the trapezoidation of the sample is doomed to spend at least linear time per round. To avoid this, we decompose the original polygonal chain into smaller subchains, sample from the resulting set of subchains and, taking advantage of the *coherence* between edges in the polygonal chain, maintain lists of *subchain conflicts* for the resulting subproblems, rather than edge conflicts.

A technical difficulty in this approach is the definition of the subproblems defined by a set of subchains. For the approach to work, one needs a decomposition with a size that is proportional to the number of subchains involved, and with *faces* (subproblems) of *bounded complexity*. The later requirement is originated in the need to be able to derive appropriate bounds for the sizes of the conflict lists, and in the need to have a decomposition that can be traversed efficiently as one *walks* along the polygonal chain. This concept also appears in Chazelle’s algorithm; following him, we call this bounded-complexity property *conformality*. Fortunately, our problem is simpler; we describe a simple procedure that computes a conformal decomposition in time linear in the number of edges in the set of subchains. This is actually sub-linear in the size of the input chain because it is performed for a small sample. In order to traverse the decomposition efficiently, we need a data structure for each subchain that answers intersection queries between a vertical edge, called a *portal*, and the subchain. Thus, as in Chazelle’s algorithm, we need a preprocessing phase that constructs these data structures prior to the actual construction phase. These phases proceed bottom-up and top-down respectively. Randomization also plays an important role in the preprocessing phase. Chazelle has *argued* that such a combination of bottom-up and top-down approaches is indispensable.

A final technicality is the proof of appropriate *sampling bounds* for the sizes of the *chain-conflict* lists of our conformal decomposition: such bounds are known under *locality* or *monotonicity* properties that our decomposition does not satisfy [1, 9, 11, 22, 24]. Fortunately, we can prove appropriate bounds using the fact that, although the faces in the decomposition do not satisfy a locality property, they are chosen from a relatively small “pool” of candidates that satisfy a locality property.

This paper is organized as follows. First, for comparison purposes, we present a detailed outline of a non-optimal randomized algorithm (Sec. 2). We then describe our procedure to compute a conformal decomposition for a set of chains (Sec. 3) and our linear time algorithm (Sec. 4). Finally, we obtain appropriate sampling bounds (Sec. 5) for analyzing our algo-

rithm (Sec. 6). We conclude with some remarks and state some open problems (Sec. 7).

2 A Non-optimal Algorithm

For the purpose of comparison with our algorithm, and as a warm-up, we outline a non-optimal randomized algorithm which is an adaptation from those in [8, 7] and [28]. Let ℓ_0 be a simple polygonal chain, S be the corresponding set of polygon edges, and let $n = |S|$. We make the nondegeneracy assumption that no two vertices have the same horizontal coordinate; this can be simulated symbolically [13, 32, 33], e.g., through lexicographic ordering. For $R \subseteq S$, let $\mathcal{T}(R)$ denote the usual (vertical) *trapezoidal decomposition* or *trapezoidation* of the plane induced by R (obtained by introducing *vertical visibility rays* from the endpoints of edges in R).




Figure 2: The trapezoidation of a simple polygon.

For $\Delta \in \mathcal{T}(R)$, let $S_{|\Delta}$ denote the *conflict list* of Δ , that is, the set of those edges in S that intersect (the interior of) Δ , and let $n_\Delta = |S_{|\Delta}|$. We adopt the following sampling model: For p with $0 \leq p \leq 1$, a p -sample R from S is obtained by taking each $s \in S$ into R with probability p independently.

2.1 Algorithm Outline

The algorithm constructs the trapezoidation of a successively finer random sample in $O(\log^* n)$ rounds. Formally, let us define a global probability $p_i = 1/\log^{(i)} n$ for round i in the computation, and let R_i be a p_i -sample from S chosen in this round (so each $s \in S$ is taken with probability p_i independently). Furthermore, let $R_i^+ = \bigcup_{j \leq i} R_j$. Note that R_i^+ is a p_i^+ -sample from S where $p_i^+ \leq \sum_{j \leq i} p_j = \Theta(p_i)$. In the i -th round, given $\mathcal{T}(R_{i-1}^+)$ and its conflicts with respect to S (that is, $S_{|\Delta}$ for $\Delta \in \mathcal{T}(R_{i-1}^+)$) the algorithm constructs $\mathcal{T}(R_i^+)$ and its conflicts with respect to S as summarized in Fig. 3.

Step 1.a, for a $\Delta \in \mathcal{T}(R_{i-1}^+)$, involves a simple scan

Non-Optimal-Trapezoidation (i -th round)

Input: $\mathcal{T}(R_{i-1}^+)$ and its conflicts w.r.t. S
Output: $\mathcal{T}(R_i^+)$ and its conflicts w.r.t. S

1. For each $\Delta \in \mathcal{T}(R_{i-1}^+)$
 - a. Determine $R_{i|\Delta}$
 - b. $T_\Delta \leftarrow \begin{cases} \mathcal{T}(R_{i|\Delta} \cup \{e_1, e_2\}) \text{ restricted to } \Delta, \\ \text{where } e_1, e_2 \text{ are the edges that} \\ \text{bound } \Delta \end{cases}$
2. Merge all the T_Δ , $\Delta \in \mathcal{T}(R_{i-1}^+)$, into $\mathcal{T}(R_i^+)$
3. Compute $S_{|\Delta}$ for all $\Delta \in \mathcal{T}(R_i^+)$, by “walking” along ℓ_0 in T_i

Figure 3: Non-optimal trapezoidation procedure.




Figure 4: The three dotted edges are added to the trapezoidation determined by the two other edges: (a) Local trapezoidations after Step 1, and (b) trapezoidation after merging in Step 2.

of the conflict list and hence takes time $O(n_\Delta)$.¹ Step 1.b computes $\mathcal{T}(R_{i|\Delta} \cup \{e_1, e_2\})$ restricted to Δ , where e_1, e_2 are the (non-vertical) edges bounding Δ . This takes time $O(r_\Delta \log r_\Delta)$, where $r_\Delta = |R_{i|\Delta}|$, using one of several algorithms for computing a trapezoidation with this complexity, e.g., either the randomized algorithm by Clarkson and Shor [9] or the one by Mulmuley [23]. Step 2 involves “stitching” together pieces of trapezoids in $\mathcal{T}(R_i^+)$ that are “chopped” by vertical-rays in $\mathcal{T}(R_{i-1}^+)$. In other words, vertical-rays that are no longer necessary are removed. See Fig. 4. It takes total time linear in the sizes of the T_Δ ’s, and hence, the time required is dominated by that of Step 1. Since each trapezoid in $\mathcal{T}(R_i^+)$ has at most four neighbors, then, assuming that an appropriate data structure is used, Step 3 can be performed in time proportional to the size of ℓ_0 , which is n , plus the total number of segment-trapezoid conflicts found.

¹Alternatively, one can maintain for each $s \in S$ the list of trapezoids it intersects, and then the scan is not necessary.

2.2 Sampling Bound and Analysis

The algorithm can be analyzed with the use of the following *sampling bound*. Let R be a p -sample from S . Then, for any function f such that $f(x) = O(e^{x/2})$,

$$\mathbf{E} \left[\sum_{\Delta \in \mathcal{T}(R)} f(p n_{\Delta}) \right] = O(r), \quad (1)$$

where $r = pn$ is the expected size of R (see [9, 24] or Section 5). Using $f(x) = x$ in Eqn. (1), the expected total size of the conflict lists is $O(n)$. This implies a bound of $O(n)$ for the expected time required by all steps in a round, except Step 1.b. Denoting the expectation with respect to the first i samples by $\mathbf{E}_{\leq i}$, the total expected time required by Step 1.b is

$$\begin{aligned} & \mathbf{E}_{\leq i} \left[\sum_{\Delta \in \mathcal{T}(R_{i-1}^+)} O(r_{\Delta} \log r_{\Delta}) \right] \\ &= \mathbf{E}_{\leq i-1} \left[\sum_{\Delta \in \mathcal{T}(R_{i-1}^+)} O((p_i n_{\Delta}) \log(p_i n_{\Delta})) \right] \\ &= \left(\frac{p_i}{p_{i-1}} \right) \log \left(\frac{p_i}{p_{i-1}} \right) \cdot O(p_{i-1} n) \\ &= \frac{1}{\log^{(i)} n} \log \left(\frac{\log^{(i-1)} n}{\log^{(i)} n} \right) \cdot O(n) = O(n), \end{aligned}$$

where we have used both $f(x) = x \log x$ and $f(x) = x$ in Eqn. (1). Thus, since the number of rounds is $O(\log^* n)$, the total expected time required by the algorithm is $O(n \log^* n)$.

3 Conformal Decomposition

Our algorithm considers subchains of the original polygonal chain, rather than individual edges, and applies sampling to subchains. In order to effectively deal with such samples, we need a method for defining subproblems of constant descriptive complexity. Consider a set L of \tilde{n} chains with a corresponding set S of n edges. Let $K \subseteq L$ be a subset of chains of L and let $R \subseteq S$ be the corresponding set of edges. For convenience, we write $\mathcal{T}(K)$ to denote the trapezoidation $\mathcal{T}(R)$. We suppose that we are given a *planar subdivision* representation (e.g., see [3, 19, 26]) of $\mathcal{T}(K)$. This planar subdivision has $O(|R|)$ faces and each face has at most 2 edges and 4 vertical rays on its boundary. For our application, we need a planar subdivision with $O(|K|)$ faces, each of which is *conformal* [4], that is, bounded by portions of at most




Figure 5: (a) Subdivision after Step 2, and (b) subdivision after Step 3.

$O(1)$ chains in K and at most $O(1)$ vertical rays determined by their vertices. We obtain this subdivision *retraction* by selecting certain rays of $\mathcal{T}(K)$. More precisely, our candidate rays are those *ray-pairs* (one ray upward and one ray downward) incident to a locally extreme vertex of a chain (a vertex without incident polygonal edges either on its left or on its right side).

If we start with the set of chains K and introduce all of these ray-pairs, the plane is divided into faces bounded by at most two chains and at most two ray-pairs which we call *chain-trapezoids*; however, the number of faces may be more than the desired bound $O(|K|)$. Let us therefore consider the *augmented adjacency graph* $\tilde{\mathcal{G}}(K)$ of this decomposition defined as follows: the *nodes* correspond to both chain-trapezoids and (locally extreme) ray-pairs, and there is an *arc* between a chain-trapezoid and a ray-pair if they are incident (Fig. 5(a) illustrates a portion of this graph). Note that the degree of a trapezoid node is two and the degree of a ray-pair node is three. This graph can be easily obtained from the usual adjacency graph $\mathcal{G}(K)$ of $\mathcal{T}(K)$.² The procedure **conformal**, in Fig. 7, selects $O(|K|)$ ray-pairs which induce a conformal decomposition of size $O(|K|)$.

The procedure selects first all the extreme ray-pairs, that is, those originating from the leftmost and rightmost vertices of each chain. The faces of the resulting subdivision $\bar{\mathcal{T}}$ are simply connected (see Fig. 5(a)). Step 3 selects other ray-pairs in a “non-local” manner, so as to obtain a subdivision whose faces are conformal (see Fig. 5(b)). More precisely, each face is bounded by at most two chains and at most two ray-pairs. Note that the portion of a chain bounding one of these faces does not need to be monotone, and also that one of the bounding chains itself can also determine one or both bounding ray-pairs.

²It is not really necessary to determine $\tilde{\mathcal{G}}(K)$ explicitly, it is sufficient to use $\mathcal{G}(K)$. However, the introduction of $\tilde{\mathcal{G}}(K)$ simplifies the description of the algorithm.




Figure 6: Some examples of chain-trapezoids. A chain-trapezoid is determined by up to four distinct chains, but can be determined by a single one as in the bottom-left example. Also, because ray-pairs from endpoints are not necessarily selected, the situation in the bottom-right example can happen.

See Fig. 6.

This algorithm clearly runs in $O(|R|)$ time: given input $\mathcal{T}(K)$, all steps can be performed by simple traversals of $\mathcal{G}(K)$, $\tilde{\mathcal{G}}(K)$, and the τ_f 's. Moreover, the number of selected ray-pairs is $O(|K|)$. It is clear that $\tilde{\mathcal{T}}$ has $O(|K|)$ selected ray-pairs; furthermore, since we select only ray-pairs with real-degree 3 in each tree τ_f , this in turn implies that at most $O(|K|)$ additional rays are selected in Step 3.

Let $\tilde{\mathcal{T}}(K)$ be the collection of all the conformal faces. Therefore, we have the following:

Lemma 3.1 *Let K be a set of chains and let $R \subseteq S$ be the corresponding set of edges, and suppose that we are given a planar subdivision representation of the trapezoidation $\mathcal{T}(R)$ of the edges in R . Then we can construct in $O(|R|)$ time a conformal subdivision $\tilde{\mathcal{T}}(K)$ containing $O(|K|)$ faces.*

We refer to the selected (single) rays as *portals*, to the conformal faces as *chain-trapezoids* (as they are defined by chains rather than by edges), and to the conformal decomposition $\tilde{\mathcal{T}}(K)$ as the *chain-trapezoidal decomposition* or *chain-trapezoidation*.

4 The Linear-Time Algorithm

Our new algorithm can be viewed as a refinement of the non-optimal algorithm of the previous section, in which sampling is applied to subchains of the original chain ℓ_0 rather than to edges. More precisely, the chain ℓ_0 is divided into a set L of subchains of length λ , and then a p -sample $K \subseteq L$ is obtained by taking each $\ell \in L$ into K with probability p independently. For each chain-trapezoid $\tilde{\Delta}$ in the chain-

conformal($K, \mathcal{T}(K)$)

Input: A set of chains K and the trapezoidation $\mathcal{T}(K)$ of its edges.

Output: Conformal decomposition $\tilde{\mathcal{T}}(K)$, and its adjacency graph.

1. Obtain the augmented adjacency graph $\tilde{\mathcal{G}}(K)$ from $\mathcal{T}(K)$.
2. Select the extreme ray-pairs of each chain. Let $\overline{\mathcal{T}}$ be the planar subdivision (which is simply connected) induced by the chains in K and these selected ray-pairs. (See Fig. 5(a).)
3. For each face f of $\overline{\mathcal{T}}$, do the following:
 - (a) Let the tree τ_f be the subgraph of $\tilde{\mathcal{G}}(K)$ corresponding to f (it includes leaves corresponding to ray-pairs that bound f). (See Fig. 8.)
 - (b) Let the *real-degree* of a vertex in τ_f be the number of incident edges corresponding to branches that contain already selected ray-pairs as leaves.
 - (c) Select any ray-pair with real-degree 3.
4. Construct $\tilde{\mathcal{T}}(K)$, the decomposition induced by all the selected ray-pairs, and its adjacency graph. (See Fig. 5(b).)

Figure 7: Conformal decomposition procedure.




Figure 8: The tree τ_f for the upper middle face in Fig. 5. The nodes corresponding to chain-trapezoids are represented by small thick circles, and the nodes corresponding to ray-pairs are represented by black circles at the corresponding locally extreme vertices.

trapezoidation $\tilde{\mathcal{T}}(K)$, let $L_{|\tilde{\Delta}|}$ denote the set of subchains in L that intersect (the interior of) $\tilde{\Delta}$. Let $n = |L|$, $\tilde{n}_{\tilde{\Delta}} = |L_{|\tilde{\Delta}|}|$. In analogy with Eqn. (1), one would conjecture a bound

$$\mathbf{E} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K)} f(p\tilde{n}_{\tilde{\Delta}}) \right] = O(p\tilde{n}). \quad (2)$$

In particular, the expected total chain-conflict size would be $O(\tilde{n}) = O(n/\lambda)$; thus, it can be made appropriately sublinear by choosing λ sufficiently large. This is a first step in obtaining a linear time algorithm. Unfortunately, we cannot proof such a bound; however, in Sect. 5, we proof a bound that is sufficient for our purpose.

At the same time, our algorithm can also be viewed as a simplification of Chazelle's algorithm, as it considers a subdivision of the input chain into successively finer subchains, which we call a *gradation*. However, while Chazelle's algorithm computes the chain-trapezoidation of *all* the subchains in each level starting with the coarser level, our algorithm does the same but for a random sample of the subchains at each level. As the subchains become finer, the random sample also becomes finer (the probability approaches 1). At the last level, the chain-trapezoidation of the sample coincides with the trapezoidation of the complete chain, the desired result. In this section, we first define precisely the gradation of subchains and its corresponding probabilities, then we give an outline of the two phases of the algorithm, and finally describe the top-down construction phase and the bottom-up preprocessing phase.

4.1 Gradation of Subchains

The sampling in our algorithm is performed on a *gradation* of subchains with $O(\log^* n)$ levels defined as follows (Chazelle uses $O(\log n)$ levels). Let ℓ_0 be the initial simple polygonal chain of size n . We decompose ℓ_0 into collections L_i of subchains of length λ_i , $i = 0, \dots, k$, starting with $L_0 = \{\ell_0\}$ and $\lambda_0 = n$, and with L_i $i > 1$, obtained by decomposing each chain $\ell \in L_{i-1}$ into a set L_i^ℓ of subchains each of size $\lambda_i = \log^2 \lambda_{i-1}$, and ending with $k = O(\log^* n)$ so that $\lambda_k = O(1)$. Thus, the subchains in the i -th gradation are $L_i = \bigcup_{\ell \in L_{i-1}} L_i^\ell$. We denote the total number of subchains in L_i by $\tilde{n}_i = |L_i| = n/\lambda_i$.

Instead of attempting to compute $\tilde{\mathcal{T}}(L_i)$ directly (the analog of what Chazelle's algorithm does), our algorithm further simplifies the problem by taking a random sample K_i of subchains from L_i of a size such that one can afford to compute the trapezoidation of




Figure 9: Gradation of subchains.

K_i using an inefficient algorithm [9, 15, 23]. Specifically, for each $i \geq 1$, we choose a global probability $p_i = 1/\log^3 \lambda_{i-1}$, and let K_i be a p_i -sample from L_i . In the i -th round, it is more convenient to deal with the set of subchains K_i^+ that consists of the subchains in K_i and the subchains in L_i contained in all the previous samples K_j , $j < i$. That is, $K_i^+ = K_i \cup \{\ell \mid \ell \in L_i \text{ and } \ell \subset \ell' \text{ where } \ell' \in K_j, j < i\}$. Note that the expected number of the later subchains is

$$\sum_{j < i} \tilde{n}_j \cdot p_j \cdot \frac{\lambda_j}{\lambda_i} = \tilde{n}_i \cdot \sum_{j < i} p_j = \tilde{n}_i \cdot o(p_i).$$

That is, the expected size of K_i^+ is dominated by the expected size of K_i . As a result, from the analysis in Section 5, it will follow that adding the subchains of previous samples does not affect substantially the randomness of the sample K_i .

4.2 Overview of the Algorithm

As mentioned previously, our polygon trapezoidation algorithm consists of two phases. The main phase proceeds top-down constructing the decompositions $\tilde{\mathcal{T}}(K_i^+)$ iteratively. For each chain trapezoid $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$, the algorithm maintains its *chain-conflict* list $L_{|\tilde{\Delta}|} \subseteq L_i$, that is, the set of subchains in L_i that intersect (the interior of) $\tilde{\Delta}$. Maintaining chain-conflict lists, rather than edge-conflict lists, is essential to the efficiency of our algorithm. At the beginning of the i -th round, we have $\tilde{\mathcal{T}}(K_{i-1}^+)$ and the chain-conflict lists $L_{|\tilde{\Delta}|}$ for each $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)$, then the algorithm adds K_i to $\tilde{\mathcal{T}}(K_{i-1}^+)$ to obtain $\tilde{\mathcal{T}}(K_i^+)$, and computes the new chain-conflict lists by following the chain ℓ_0 without actually scanning every edge. In a preprocessing bottom-up phase, the algorithm constructs for each chain $\ell \in L_i$, $i = 1, \dots, k$, a data structure $\mathcal{D}(\ell)$ that supports portal-chain intersection queries: given a portal ρ , determine whether the chain ℓ intersects ρ . These queries are needed for the efficient computation of chain-conflict lists during the construction phase. These data structures also support ray-shooting queries (given a point x , determine the lowest point in ℓ hit by a vertical ray upward

Top-Down (i -th round)

Input: $\tilde{\mathcal{T}}(K_{i-1}^+)$ and its conflicts w.r.t. L_{i-1}

Output: $\tilde{\mathcal{T}}(K_i^+)$ and its conflicts w.r.t. L_i

1. For each $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)$
 - a. Determine $K_{i|\tilde{\Delta}}$
 - b. $T_{\tilde{\Delta}} \leftarrow \begin{cases} \mathcal{T}(K_{i|\tilde{\Delta}} \cup \{\ell_1, \ell_2\}) \text{ restricted to } \tilde{\Delta}, \\ \text{where } \ell_1, \ell_2 \text{ are the chains that} \\ \text{bound } \tilde{\Delta} \end{cases}$
2. Merge all the $T_{\tilde{\Delta}}$, $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)$, into $\mathcal{T}(K_i^+)$
3. Obtain $\tilde{\mathcal{T}}(K_i^+)$ using `conformal`($K_i^+, \mathcal{T}(K_i^+)$)
4. Compute $L_{i|\tilde{\Delta}}$ for all $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$, by “hopping” along L_i in $\tilde{\mathcal{T}}(K_i^+)$

Figure 10: Top-down phase procedure.

from x) which can then be used for testing whether a query point is contained in a chain-trapezoid (perform ray-shooting queries on the two bounding chains and determine if the result corresponds to hitting them from inside the chain-trapezoid).

4.3 Top-Down Construction Phase

Let us now formally describe how our algorithm performs the top-down construction phase. In the i -th round, given the decomposition $\tilde{\mathcal{T}}(K_i^+)$ and its conflict lists with respect to L_{i-1} , the algorithm adds the subchains in K_i^+ to $\tilde{\mathcal{T}}(K_{i-1}^+)$ as summarized in Fig. 10.

Step 1.a determines the conflict list $K_{i|\tilde{\Delta}}$ by checking for each $\ell \in L_{i-1|\tilde{\Delta}}$ and $\ell' \in L_i^\ell \cap K_i$, whether ℓ' intersects $\tilde{\Delta}$: if so either ℓ' intersects one of the portals of $\tilde{\Delta}$, or its endpoints are inside $\tilde{\Delta}$. Both queries are solved by the same data structures \mathcal{D} constructed during the preprocessing phase. However, note that the point location query is on chains in L_{i-1} and, consequently, it is more expensive (we could afford to construct a faster standard point location data structure for $\tilde{\Delta}$, but it is not necessary). Let $\tilde{r}_{i,\tilde{\Delta}} = |K_{i|\tilde{\Delta}}|$. Step 1.b computes the trapezoidation $\mathcal{T}(K_{i|\tilde{\Delta}} \cup \{\ell_1, \ell_2\})$ restricted to $\tilde{\Delta}$, where ℓ_1 and ℓ_2 are the two chains bounding $\tilde{\Delta}$. This uses a simple algorithm with running time $O(r_{i,\tilde{\Delta}} \log r_{i,\tilde{\Delta}})$ [9, 15, 23], where $r_{i,\tilde{\Delta}} = O((\tilde{r}_{i,\tilde{\Delta}} + 1)\lambda_i)$ is the number of edges involved in all these chains (the plus 1 accounts for ℓ_1 and ℓ_2). Step 2, with a simple traversal

of all the $T_{\tilde{\Delta}}$ ’s, “stitches” together trapezoids of $\mathcal{T}(K_i^+)$ “chopped” by the portals of $\tilde{\mathcal{T}}(K_{i-1}^+)$; this takes time linear in the total size of the $T_{\tilde{\Delta}}$. The procedure `conformal` in Step 3 was described in Section 3; it takes time linear in the size of $\mathcal{T}(K_i^+)$ and returns the (conformal) chain-trapezoidation $\tilde{\mathcal{T}}(K_i^+)$ (including its adjacency graph). In Step 4, the conflict lists $L_{i|\tilde{\Delta}}$ for $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$ are found chain by chain, using the adjacency graph of $\tilde{\mathcal{T}}(K_i^+)$ and the data structures $\mathcal{D}(\ell)$ to “hop” along L_i in $\tilde{\mathcal{T}}(K_i^+)$, as described next.

Hopping. If a chain is already in K_i^+ , then it is part of the boundary for some chain-trapezoids and it can automatically be recorded as part of their conflict lists. So, consider some $\ell \in L_i \setminus K_i^+$ and suppose that we know a chain trapezoid $\tilde{\Delta}_0 \in \tilde{\mathcal{T}}(K_i^+)$ that contains the first endpoint e of ℓ . Note that the chain-trapezoids in $\tilde{\mathcal{T}}_i$ that conflict with ℓ are connected. Thus, we perform a breadth-first-search traversal of the adjacency graph of $\tilde{\mathcal{T}}(K_i^+)$, starting with $\tilde{\Delta}_0$. When a chain-trapezoid $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$ is visited, it is labeled as a *conflict* and each of the portals that separates $\tilde{\Delta}$ from an unvisited chain-trapezoid $\tilde{\Delta}' \in \tilde{\mathcal{T}}(K_i^+)$ is tested for conflict with ℓ using $\mathcal{D}(\ell)$. If ℓ conflicts with the portal, the traversal visits $\tilde{\Delta}'$. Note that ℓ can zig-zag arbitrarily within the set of chain-trapezoids that it intersects. See Fig. 11.




Figure 11: Computing the chain-conflicts for a chain. Note that only one conflict per portal is found, and that the order the conflicts are discovered (indicated by numbers) does not necessarily reflect their occurrences on the chain. The portals not queried are shown lighter.

This procedure performs $O(1)$ portal-chain conflict queries per conflict actually found. The location of the first endpoint e is given by the location of the second endpoint of the preceding chain ℓ' (known al-

ready if $\ell' \in K_i^+$). The location of ℓ' 's second endpoint e' can be determined by performing a point location query for each chain-trapezoid found to be in conflict with ℓ . This is necessary since the conflicts were computed not by a linear scan of ℓ , but rather by “hopping” between portals.

Running Time. In Sec. 6, using the sampling bounds obtained in the next section, we show that given the data structures $\mathcal{D}(\ell)$ with query time $O(\log^{3+\epsilon} \lambda_i)$, for $\ell \in L_i$, the top-down construction phase is completed in expected time $O(n)$.

4.4 Bottom-Up Preprocessing Phase

In the preprocessing phase, the algorithm constructs data structures for portal-chain conflict queries, to be used to hop along the chains in the top-down phase. Recall the gradation of subchains L_i , $i = 0, \dots, k$, defined above and that for $\ell \in L_{i-1}$, L_i^ℓ is the set of subchains of ℓ in L_i . Let $K_i^\ell = L_i^\ell \cap K_i$. Note that since K_i is a p_i -sample from L_i , then K_i^ℓ is a p_i -sample from L_i^ℓ .³

For each ℓ in L_{i-1} , $i = 1, \dots, k$, we construct a data structure $\mathcal{D}(\ell)$ that consists of:

- (i) $\tilde{\mathcal{T}}(K_i^\ell)$ and a corresponding point location structure with query time $O(\log \lambda_{i-1})$;
- (ii) for each $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^\ell)$, the chain-conflict list $L_{i|\tilde{\Delta}}^\ell$.

We can use for the construction of $\tilde{\mathcal{T}}(K_i^\ell)$ either the randomized algorithm by Clarkson and Shor [9] or the one by Mulmuley [23], which also result in point location data structures with logarithmic query time as needed in (i). Alternatively, other planar point location data structures can be used [12, 17, 25, 27, 29].

A portal-chain conflict query for an arbitrary portal ρ and chain $\ell \in L_{i-1}$ first uses $\mathcal{D}(\ell)$'s point location data structure $\tilde{\mathcal{T}}(K_i^\ell)$ to locate the endpoints of ρ . If ρ 's endpoints are contained in different chain-trapezoids in $\tilde{\mathcal{T}}(K_i^\ell)$, then ρ must intersect ℓ , and a conflict is reported. Otherwise, ρ is entirely contained in some $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^\ell)$, and the query continues recursively in the data structures $\mathcal{D}(\ell')$, for each and every subchain ℓ' that bounds $\tilde{\Delta}$ or in $\tilde{\Delta}$'s conflict list $L_{i|\tilde{\Delta}}^\ell$, which includes the subchains that bound $\tilde{\Delta}$. See Fig. 13. This query procedure is summarized in Fig. 12. The query procedure for ray-shooting, which determines the lowest intersection point, is similar and we omit it.

conflict?($\rho, \ell, i - 1$)

Input: A portal ρ and a chain ℓ in L_{i-1} .

Output: **Yes** or **No**

1. Determine $\tilde{\Delta}, \tilde{\Delta}' \in \tilde{\mathcal{T}}_i^\ell$ which contain the endpoints of ρ
2. If $\tilde{\Delta}$ and $\tilde{\Delta}'$ are different then return **Yes**
3. For each ℓ' that bounds $\tilde{\Delta}$ or in $L_{i|\tilde{\Delta}}^\ell$ do
 if **conflict?**(ρ, ℓ', i) = **Yes** then return **Yes**
4. Return **No**

Figure 12: Portal-chain query procedure.




Figure 13: Two cases in the portal-chain conflict query: The portal endpoints are in different or in the same chain-trapezoid.

The computation of conflict lists for $\tilde{\mathcal{T}}(K_i^\ell)$ during the construction of $\mathcal{D}(\ell)$ also uses “hopping” and, hence, **conflict?** inductively. The bottom-up preprocessing phase is summarized in Fig. 14.

Running Time. In Sec. 6, using the sampling bounds obtained in the next section, we show that the construction of the data structures $\mathcal{D}(\ell)$ is completed in expected $O(n)$ time. Moreover, we show that, even though we recurse on each subchain in a conflict list, the expected query time for a chain $\ell \in L_i$ is $O(\log^{3+\epsilon} \lambda_i)$, where $\epsilon > 0$ is an arbitrary small fraction. Thus, we can summarize our results in the following.

Theorem 4.1 *Our randomized two-phase algorithm constructs the trapezoidation of a simple polygon of size n in expected $O(n)$ time.*

5 Sampling Bounds

To analyze the running time of our algorithm, we need bounds on the sizes of the subproblems resulting by taking a random sample from a set of chains and then

³This phase could use a sampling independent of that in the construction phase, but this is not necessary.

Bottom-Up (i -th round)

Input: $\mathcal{D}(\ell)$ for each $\ell \in L_{j-1}$, $j > i$

Output: $\mathcal{D}(\ell)$ for each $\ell \in L_{i-1}$

For each $\ell \in L_{i-1}$ do

1. Compute $\mathcal{T}(K_i^\ell)$ and a corresponding point location structure
2. Compute $\tilde{\mathcal{T}}(K_i^\ell)$ using `conformal`($K_i^\ell, \mathcal{T}(K_i^\ell)$)
3. Compute $L_{i|\tilde{\Delta}}$ for all $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^\ell)$ by “hopping” along L_i^ℓ in $\tilde{\mathcal{T}}(K_i^\ell)$

Figure 14: Bottom-up phase procedure.

constructing its chain-trapezoidation. Let K be a p -sample from L , and recall that for a chain-trapezoid $\tilde{\Delta} \in \tilde{\mathcal{T}}(K)$, $L_{|\tilde{\Delta}}$ denotes the list of conflicts of $\tilde{\Delta}$ in L , and that $\tilde{n}_{\tilde{\Delta}} = |L_{|\tilde{\Delta}}|$. Unfortunately, we cannot prove the bound in Eqn. (2). Such a bound can be proved in the framework of *configuration spaces*, when certain *locality* [9, 24] or *monotonicity* [11, 1] properties hold for the decomposition induced by the sample (see [22] for a survey), but neither of these properties hold for our chain-trapezoidation. Fortunately, we can prove a weaker bound that is only a factor $O(f(\log \lambda))$ larger, and that suffices to verify that our algorithm has expected linear running time. The proof of the bound uses a standard trick [9, 5]: one obtains a nontrivial bound for a p -sample in terms of a trivial bound for a $(p/2)$ -sample. First, we need a fact about the chain-trapezoidation that limits the amount of non-locality in the definition of the chain-trapezoidation.

Recall that a chain-trapezoid is bounded by at most two chains and at most two ray-pairs, each one originating from another chain (but possibly a bounding chain). We say that these at most four chains *determine* the chain-trapezoid. Let $\tilde{\mathcal{T}}^*(L)$ be the set of all chain-trapezoids determined by L , that is, those chain-trapezoids determined by a subset of at most four chains in L (but note that some other chains in L can conflict with such trapezoids). Let $\tilde{\mathcal{T}}^c(K)$ be the set of all *candidate* chain-trapezoids determined by K and with empty conflict list with respect to K . Note that $\tilde{\mathcal{T}}^c(K)$ is bigger than $\tilde{\mathcal{T}}(K)$ as there are candidate chain-trapezoids determined by K that were not chosen in our construction of $\tilde{\mathcal{T}}(K)$. For $\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L)$, let $\delta(\tilde{\Delta}) \subseteq L$ denote the set of those up to four chains that determine $\tilde{\Delta}$. For $\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L)$, we have the *locality* property:

$$\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K) \quad \text{iff} \quad \delta(\tilde{\Delta}) \subseteq K \text{ and } L_{|\tilde{\Delta}} \cap K = \emptyset. \quad (3)$$

Though our chain-trapezoidation lacks locality, or even monotonicity, the following lemma states that we choose it out of a relatively small “pool” of candidates that satisfy the locality property.

Lemma 5.1 *Let K be a set of chains each of length at most λ . Then, $|\tilde{\mathcal{T}}^c(K)| = O(|K|\lambda^2)$.*

Proof. In $\mathcal{T}(K)$, let a *region* be the union of the trapezoids corresponding to a connected subgraph of $\mathcal{G}(K)$. For each $\tilde{\Delta} \in \tilde{\mathcal{T}}(K)$ consider the maximum region $R_{\tilde{\Delta}}$ of $\mathcal{T}(K)$ that contains $\tilde{\Delta}$ and is bounded by the same one or two chains that bound $\tilde{\Delta}$. Note that $R_{\tilde{\Delta}}$ may not be conformal. Any candidate chain-trapezoid is a subregion of $R_{\tilde{\Delta}}$ for some $\tilde{\Delta}$. Since clearly the number of subregions of any $R_{\tilde{\Delta}}$ is $O(\lambda^2)$, and the size of $\tilde{\mathcal{T}}(K)$ is $O(|K|)$, then it follows that the size of $\tilde{\mathcal{T}}^c(K)$ is $O(|K|\lambda^2)$. \square

We use this result now to prove bounds for the chain-conflict list sizes in the chain-trapezoidal decomposition of a random sample of chains.

Lemma 5.2 *Let L be a set of \tilde{n} chains of length λ , and let $n = \lambda\tilde{n}$ be the total number of edges. Let $K \subseteq L$ be a p -sample, $\tilde{r} = p\tilde{n}$ its expected size, and let $\tilde{\mathcal{T}}(K)$ be its chain-trapezoidal decomposition. For $\tilde{\Delta} \in \tilde{\mathcal{T}}(K)$, we write $\tilde{n}_{\tilde{\Delta}} = |L_{|\tilde{\Delta}}|$. Let f be a positive nondecreasing function such that $f(O(x)) = O(f(x))$. Then*

$$\mathbf{E} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K)} f(p\tilde{n}_{\tilde{\Delta}}) \right] = O(\tilde{r} \cdot f(\log \lambda)) \quad (4)$$

Proof. Let $K' \subseteq L$ be a $(p/2)$ -sample. Recall that $\tilde{\mathcal{T}}^c(K)$ is the set of candidate chain-trapezoids for K , and that for these chain-trapezoids the locality property in Eqn. (3) holds. Thus,

$$\begin{aligned} & \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K)\} \\ &= p^{\delta(\tilde{\Delta})} (1-p)^{\tilde{n}_{\tilde{\Delta}}} \\ &= 2^{\delta(\tilde{\Delta})} \cdot \left(\frac{1-p}{1-p/2} \right)^{\tilde{n}_{\tilde{\Delta}}} \cdot (p/2)^{\delta(\tilde{\Delta})} (1-p/2)^{\tilde{n}_{\tilde{\Delta}}} \\ &\leq 16 \cdot e^{-p\tilde{n}_{\tilde{\Delta}}/2} \cdot \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K')\}, \end{aligned}$$

using $\delta(\tilde{\Delta}) \leq 4$ and $(1-p)/(1-p/2) \leq 1-p/2 \leq e^{-p/2}$. Using this upper bound, we obtain

$$\mathbf{E} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K)} f(p\tilde{n}_{\tilde{\Delta}}) \right]$$

$$\begin{aligned}
&= \sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L)} f(p\tilde{n}_{\tilde{\Delta}}) \cdot \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K)\} \\
&= O\left(\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L)} \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K')\}\right) \\
&= O(\mathbf{E}[|\tilde{\mathcal{T}}^c(K')|]) \\
&= O(\tilde{r}\lambda^2),
\end{aligned}$$

using Lemma 5.1. Let $\tau > 0$ be a parameter. Then, using again the upper bound above,

$$\begin{aligned}
&\mathbf{E}\left[\sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L) \\ \tilde{n}_{\tilde{\Delta}} > \frac{2(\log \tau)}{p}}} f(p\tilde{n}_{\tilde{\Delta}})\right] \\
&= \sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L) \\ \tilde{n}_{\tilde{\Delta}} > 2(\log \tau)/p}} f(p\tilde{n}_{\tilde{\Delta}}) \cdot \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K)\} \\
&\leq 16 \cdot \sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L) \\ \tilde{n}_{\tilde{\Delta}} > \frac{2(\log \tau)}{p}}} f(p\tilde{n}_{\tilde{\Delta}}) \cdot e^{-p\tilde{n}_{\tilde{\Delta}}/2} \cdot \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K')\} \\
&\leq \frac{16}{\tau} \cdot \sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L)} f(p\tilde{n}_{\tilde{\Delta}}) \cdot \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K')\} \\
&= \frac{16}{\tau} \cdot \mathbf{E}\left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K')} f(p\tilde{n}_{\tilde{\Delta}})\right] \\
&= \frac{O(1)}{\tau} \cdot \mathbf{E}[|\tilde{\mathcal{T}}^c(K'')|] = O\left(\tilde{r} \cdot \frac{\lambda^2}{\tau}\right),
\end{aligned}$$

where K'' is a $(p/4)$ -sample. Finally, using $\tau = \lambda^2$,

$$\begin{aligned}
&\mathbf{E}\left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K)} f(p\tilde{n}_{\tilde{\Delta}})\right] \\
&= \mathbf{E}\left[\sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}(K) \\ \tilde{n}_{\tilde{\Delta}} > \frac{2(\log \tau)}{p}}} f(p\tilde{n}_{\tilde{\Delta}})\right] + \mathbf{E}\left[\sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}(K) \\ \tilde{n}_{\tilde{\Delta}} \leq \frac{2(\log \tau)}{p}}} f(p\tilde{n}_{\tilde{\Delta}})\right] \\
&\leq \mathbf{E}\left[\sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K) \\ \tilde{n}_{\tilde{\Delta}} > \frac{2(\log \tau)}{p}}} f(p\tilde{n}_{\tilde{\Delta}})\right] + f(2 \log \tau) \cdot \mathbf{E}[|\tilde{\mathcal{T}}(K)|] \\
&= O\left(\tilde{r} \cdot \frac{\lambda^2}{\tau}\right) + O(f(2 \log \tau) \cdot \tilde{r}) = O(f(\log \lambda) \cdot \tilde{r}).
\end{aligned}$$

□

For the analysis of the query time of the ray-shooting data structure, we also need a bound for the expectation of the conflict list size of the chain-trapezoid that contains a *fixed* point x .

Lemma 5.3 *Let $K \subseteq L$ be a p -sample, and for a fixed point x , let $\tilde{\Delta}_x$ be the chain-trapezoid in $\tilde{\mathcal{T}}(K)$ that contains x . Then*

$$\mathbf{E}[|L_{|\tilde{\Delta}_x}|] = O\left(\frac{1}{p} \cdot \log \lambda\right) = O\left(\frac{\tilde{n}}{\tilde{r}} \cdot \log \lambda\right).$$

Proof. A similar proof as for the previous lemma applies. Only note that the number of candidate chain-trapezoids in $\tilde{\mathcal{T}}^c(K)$ that contain x is $O(\lambda^2)$, and so

$$\begin{aligned}
&\sum_{\substack{\tilde{\Delta} \in \tilde{\mathcal{T}}^*(L) \\ x \in \tilde{\Delta}}} \text{Prob}\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K)\} \\
&= \mathbf{E}[|\{\tilde{\Delta} \in \tilde{\mathcal{T}}^c(K) : x \in \tilde{\Delta}\}|] = O(\lambda^2).
\end{aligned}$$

□

6 Running Time Analysis

First, we note that the sampling bound of Lemma 5.2 holds for K_i^+ even though it contains, in addition to the true random p_i -sample K_i , all the subchains of previous samples K_j , $j < i$. This is because, as noted in Section 4, the size is dominated by K_i and so, from the proof of the lemma, the right hand side of Eqn. (4) is not affected.

In the analysis below, two specific sums appear that can be bounded using Eqn. (4):

$$\mathbf{E}\left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K)} \tilde{n}_{\tilde{\Delta}}\right] = O(\tilde{n} \cdot \log \lambda), \quad (5)$$

and for any $\alpha > 0$,

$$\begin{aligned}
&\mathbf{E}\left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K)} \tilde{n}_{\tilde{\Delta}} \log(\alpha \tilde{n}_{\tilde{\Delta}})\right] \\
&= O\left(\tilde{n} \cdot \log \lambda \cdot \left(\log\left(\frac{\alpha}{p} \log \lambda\right)\right)\right). \quad (6)
\end{aligned}$$

6.1 Preprocessing Phase

First, we verify the query time. The expected query time $Q(\lambda_i)$ is the sum of the time needed for a point location query, plus the expected time needed for all the recursive queries. Thus, we have

$$\begin{aligned}
Q(\lambda_{i-1}) &= O(\log \lambda_{i-1}) + O((1/p_i) \log \lambda_i) \cdot Q(\lambda_i) \\
&= O(\log \lambda_{i-1}) + O((\log \lambda_{i-1})^3 \log \lambda_i) \cdot Q(\lambda_i) \\
&= O(\log^{3+\epsilon} \lambda_{i-1}),
\end{aligned}$$

where ϵ is any positive fraction. We have used Lemma 5.3 to bound the expected number of chain-conflicts as $O((1/p_i) \log \lambda_i)$. Note that it is valid to use the bound for a *fixed* point, as the random choices in the i -th and later rounds are independent of the random choices in earlier rounds. Next, we estimate the expected construction time for $\ell \in L_i$. From the previous description, using Eqn. (5), the expected time is

$$\begin{aligned} & O \left((p_i \lambda_{i-1}) \log(p_i \lambda_{i-1}) + \left(\frac{\lambda_{i-1}}{\lambda_i} \log \lambda_i \right) \cdot Q(\lambda_i) \right) \\ &= O \left(\frac{\lambda_{i-1}}{\log^3 \lambda_{i-1}} \log \left(\frac{\lambda_{i-1}}{\log^3 \lambda_{i-1}} \right) + \right. \\ &\quad \left. \frac{\lambda_{i-1}}{\log^2 \lambda_{i-1}} \cdot \log(\log^2 \lambda_{i-1}) \cdot \log^{3+\epsilon}(\log^2 \lambda_{i-1}) \right) \\ &= O \left(\frac{\lambda_{i-1}}{\log \lambda_{i-1}} \right), \end{aligned}$$

where the first term accounts for the construction of $\tilde{\mathcal{T}}(K_i^\ell)$, and the second term accounts for $O(1)$ portal-chain intersection queries per chain-conflict. Adding over all $\ell \in L_{i-1}$, the construction time in the $(i-1)$ -st level is $O(n/\log \lambda_{i-1})$, and adding over all i , the total construction time is $O(n)$.

6.2 Construction Phase

Consider the i -th round, and let $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$. For the purpose of analysis, let us write $\tilde{n}_{i,\tilde{\Delta}} = |L_{i,\tilde{\Delta}}|$ and $\tilde{r}_{i,\tilde{\Delta}} = |K_{i,\tilde{\Delta}}|$. Recall that $\tilde{n}_i = |L_i| = n/\lambda_i$. Note that the expected value of $\tilde{r}_{i,\tilde{\Delta}}$ is bounded by $p_i \cdot \tilde{n}_{i-1,\tilde{\Delta}} \cdot \frac{\lambda_{i-1}}{\lambda_i}$. In Step 1.a, for each $\ell \in L_{i-1} \cap \tilde{\Delta}$, checking whether $\ell' \in L_i^\ell \cap K_i$ intersects $\tilde{\Delta}$ takes expected time is $O(\log^{3+\epsilon} \lambda_{i-1})$, using the data structure \mathcal{D} for the chains bounding $\tilde{\Delta}$. Thus, using Eqn. (5), the total expected time for this step is big-O of

$$\begin{aligned} & \mathbf{E}_{\leq i-1} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)} p_i \cdot \tilde{n}_{i-1,\tilde{\Delta}} \cdot \frac{\lambda_{i-1}}{\lambda_i} \cdot \log^{3+\epsilon} \lambda_{i-1} \right] \\ &= O \left(p_i \cdot (\tilde{n}_{i-1} \log \lambda_{i-1}) \cdot \frac{\lambda_{i-1}}{\lambda_i} \cdot \log^{3+\epsilon} \lambda_{i-1} \right) \\ &= O \left(\frac{n}{\log^{1-\epsilon} \lambda_{i-1}} \right). \end{aligned}$$

(Here, the term $\log^{3+\epsilon} \lambda_{i-1}$ could be reduced to $\log \lambda_{i-1} + \log^{3+\epsilon} \lambda_i$, if we construct a point location data structure for each $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)$.) In Step 1.b, $T_{\tilde{\Delta}}$ is computed using an algorithm with running time $O(r_{i,\tilde{\Delta}} \log r_{i,\tilde{\Delta}})$ where $r_{i,\tilde{\Delta}} = O((\tilde{r}_{i,\tilde{\Delta}} + 1)\lambda_i)$. Using

Eqn. (6), the total expected time is big-O of (to abbreviate, we are dropping the 1 in the bound for $r_{i,\tilde{\Delta}}$; a more complete calculation leads to the same result)

$$\begin{aligned} & \mathbf{E}_{\leq i} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)} r_{i,\tilde{\Delta}} \log r_{i,\tilde{\Delta}} \right] = O(p_i \cdot \lambda_{i-1}) \cdot \\ & \mathbf{E}_{\leq i-1} \left[\sum_{\tilde{\Delta} \in \tilde{\mathcal{T}}(K_{i-1}^+)} \tilde{n}_{i-1,\tilde{\Delta}} \cdot \log(p_i \cdot \lambda_{i-1} \cdot \tilde{n}_{i-1,\tilde{\Delta}}) \right] \\ &= O \left(p_i \cdot \lambda_{i-1} \cdot \tilde{n}_{i-1} \cdot \log \lambda_{i-1} \cdot \right. \\ &\quad \left. \log \left(\frac{p_i \lambda_{i-1}}{p_{i-1}} \log \lambda_{i-1} \right) \right) \\ &= O(n \cdot p_i \cdot \log^2 \lambda_{i-1}) = O \left(\frac{n}{\log \lambda_{i-1}} \right). \end{aligned}$$

The conformal decomposition in Step 3 is constructed as described in Section 3 and requires expected time big-O of

$$\mathbf{E}_{\leq i}[\|\mathcal{T}(K_i^+)\|] = O(p_i \cdot \tilde{n}_i \cdot \lambda_i) = O \left(\frac{n}{\log^3 \lambda_{i-1}} \right).$$

In Step 4, the conflict lists for regions $\tilde{\Delta} \in \tilde{\mathcal{T}}(K_i^+)$ with subchains in L_i are found using the data structures \mathcal{D} , $O(1)$ queries per conflict determined, so the expected time is

$$\begin{aligned} & O(\tilde{n}_i \cdot \log \lambda_i \cdot Q(\lambda_i)) = O \left(\frac{n}{\lambda_i} \cdot \log \lambda_i \cdot \log^{3+\epsilon} \lambda_i \right) \\ &= O \left(n \cdot \frac{\log^{4+\epsilon} \lambda_i}{\lambda_i} \right) = O \left(\frac{n}{\log \lambda_{i-1}} \right). \end{aligned}$$

The sum of all these contributions over all the rounds is $O(n)$.

7 Concluding Remarks

We have presented a randomized algorithm for computing the trapezoidation of a simple polygon, and hence a triangulation, that runs in expected time that is linear with the number of edges. The algorithm is considerably simpler than Chazelle's algorithm. On the other hand, it is comparatively more complicated than the non-optimal randomized algorithm and, since for any practical value of n , $\log^* n$ is a small constant, our algorithm is not likely to be of practical value.

We conclude by mentioning some questions that remain open. Is the conjectured tighter sampling bound for our conformal decomposition true? Is it

possible to combine our polygon trapezoidation algorithm with a segment intersections algorithm to obtain an algorithm that can report the k intersections of a chain of n segments in time $O(n + k)$? Can our linear time algorithm be parallelized? Can the approach of sampling on subchains lead to efficient algorithms for other problems on simple polygons? Finally, does a deterministic algorithm simpler than Chazelle's exist?

References

- [1] P. K. Agarwal, M. de Berg, J. Matoušek, and O. Schwarzkopf. Constructing levels in arrangements and higher order Voronoi diagrams. *SIAM J. Comput.*, 27:654–667, 1998.
- [2] T. Asano, T. Asano, and R. Y. Pinter. Polygon triangulation: Efficiency and minimality. *J. Algorithms*, 7:221–231, 1986.
- [3] B. G. Baumgart. A polyhedron representation for computer vision. In *Proc. AFIPS Natl. Comput. Conf.*, volume 44, pages 589–596. AFIPS Press, Arlington, Va., 1975.
- [4] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6(5):485–524, 1991.
- [5] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- [6] B. Chazelle and J. Incerpi. Triangulation and shape-complexity. *ACM Trans. Graph.*, 3(2):135–152, 1984.
- [7] K. L. Clarkson, R. Cole, and R. E. Tarjan. Erratum: Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2(3):341–343, 1992.
- [8] K. L. Clarkson, R. Cole, and R. E. Tarjan. Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2(2):117–133, 1992.
- [9] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [10] K. L. Clarkson, R. E. Tarjan, and C. J. Van Wyk. A fast Las Vegas algorithm for triangulating a simple polygon. *Discrete Comput. Geom.*, 4:423–432, 1989.
- [11] M. de Berg, K. Dobrindt, and O. Schwarzkopf. On lazy randomized incremental construction. *Discrete Comput. Geom.*, 14:261–286, 1995.
- [12] H. Edelsbrunner, L. J. Guibas, and J. Stolfi. Optimal point location in a monotone subdivision. *SIAM J. Comput.*, 15(2):317–340, 1986.
- [13] H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, 1990.
- [14] A. Fournier and D. Y. Montuno. Triangulating simple polygons and equivalent problems. *ACM Trans. Graph.*, 3(2):153–174, 1984.
- [15] M. R. Garey, D. S. Johnson, F. P. Preparata, and R. E. Tarjan. Triangulating a simple polygon. *Inform. Process. Lett.*, 7(4):175–179, 1978.
- [16] M. T. Goodrich. Planar separators and parallel polygon triangulation. *J. Comput. Syst. Sci.*, 51(3):374–389, 1995.
- [17] M. T. Goodrich, M. Orletsky, and K. Ramaier. Methods for achieving fast query times in point location data structures. In *Proc. 8th ACM-SIAM Sympos. Discrete Algorithms*, pages 757–766, 1997.
- [18] L. J. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987.
- [19] L. J. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graph.*, 4(2):74–123, Apr. 1985.
- [20] J. Hershberger. Optimal parallel algorithms for triangulated simple polygons. *Internat. J. Comput. Geom. Appl.*, 5:145–170, 1995.
- [21] D. G. Kirkpatrick, M. M. Klawe, and R. E. Tarjan. Polygon triangulation in $O(n \log \log n)$ time with simple data structures. In *Proc. 6th Annu. ACM Sympos. Comput. Geom.*, pages 34–43, 1990.
- [22] J. Matoušek. Derandomization in computational geometry. *J. Algorithms*, 20:545–580, 1996.
- [23] K. Mulmuley. A fast planar partition algorithm, I. *J. Symbolic Comput.*, 10(3-4):253–280, 1990.
- [24] K. Mulmuley. *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [25] F. P. Preparata. A new approach to planar point location. *SIAM J. Comput.*, 10(3):473–482, 1981.
- [26] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, NY, 1985.
- [27] N. Sarnak and R. E. Tarjan. Planar point location using persistent search trees. *Commun. ACM*, 29(7):669–679, July 1986.
- [28] R. Seidel. A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Comput. Geom. Theory Appl.*, 1(1):51–64, 1991.
- [29] R. Seidel. On the exact query complexity of planar point location. In *Abstracts 14th European Workshop Comput. Geom.*, pages 7–8, 1998.
- [30] R. E. Tarjan and C. J. Van Wyk. An $O(n \log \log n)$ -time algorithm for triangulating a simple polygon. *SIAM J. Comput.*, 17:143–178, 1988. Erratum in 17 (1988), 106.
- [31] C. K. Yap. Parallel triangulation of a polygon in two calls to the trapezoidal map. *Algorithmica*, 3:279–288, 1988.
- [32] C. K. Yap. A geometric consistency theorem for a symbolic perturbation scheme. *J. Comput. Syst. Sci.*, 40(1):2–18, 1990.
- [33] C. K. Yap. Symbolic treatment of geometric degeneracies. *J. Symbolic Comput.*, 10:349–370, 1990.