Enumerating Triangulation Paths ⋆

Adrian Dumitrescu [a], Bernd Gärtner [b], Samuele Pedroni [b], and Emo Welzl [b]

[a] *Applied Mathematics and Statistics SUNY Stony Brook*

[b] *Institut für Theoretische Informatik, ETH Zürich ETH Zentrum, CH-8092 Zürich, Switzerland*

Abstract

Recently, Aichholzer introduced the remarkable concept of the so-called triangulation path (of a triangulation with respect to a segment), which has the potential of providing efficient counting of triangulations of a point set, and efficient representations of all such triangulations. Experiments support such evidence, although – apart from the basic uniqueness properties – little has been proved so far.

In this paper we provide an algorithm which enumerates all triangulation paths (of all triangulations of a given point set with respect to a given segment) in time $O(t\,n^3 \log n)$ and $O(n)$ space, where n denotes the number of points and t is the number of triangulation paths. For the algorithm we introduce the notion of flips between such paths, and define a structure on all paths such that the reverse search approach can be applied. We also refute Aichholzer's conjecture that points in convex position maximize the number of such paths. There are configurations that allow $\Omega(2^{2n-\Theta(\log n)})$ paths.

1 Introduction

Throughout this paper we let S stand for a set of n points in general position in the plane (i.e. no three on a line, no four on a circle), and we let P stand for a closed and bounded simple polygonal region in the plane, (polygon for short, from now on) with its vertices in general position. A pair (S, P) is a *points-in-polygon-pair*, if S contains all vertices of P, and $S \subseteq P$. For example, S forms such a pair with its convex hull *conv*(S). Given a points-in-polygon-pair (S, P), a triangulation T of

⋆ A preliminary version of this paper has appeared in *Proc. 12th Canadian Conf. on Comput. Geom.* (2000), 233–238.

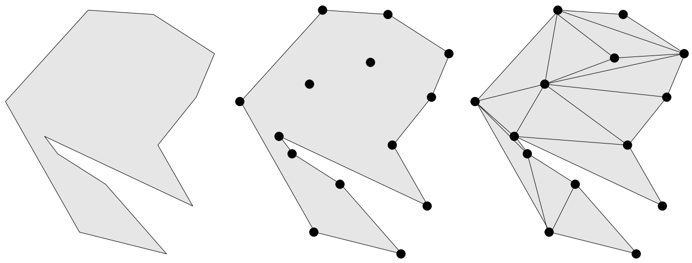

Fig. 1. A polygon P, a points-in-polygon-pair (S, P), and one of the triangulations of (S, P).

P with vertex set S is called *triangulation of* (S, P); $\mathcal{T}(S, P)$ denotes the set of all triangulations of (S, P).

Triangulations are basic building blocks in a number of applications, e.g. for the finite-element method and for the representation of terrains. Depending on their purpose, some of the triangulations in $\mathcal{T}(S, P)$ are considered (or, are in fact,) better than others. The *constrained Delaunay triangulation* is known to avoid small angles as much as possible and it can be constructed in $O(n \log n)$, [4]. A triangulation that minimizes the sum of edge lengths, a so-called *min-weight triangulation*, has so far refused an algorithm for efficient construction[1] . It is a challenging open problem to decide whether this problem is NP-hard. Similarly, the complexity status of counting the number of triangulations of (S, P) is open, with the currently best known bound of $O(n \cdot |\mathcal{T}(S, P)|)$, [2]. See [3] for a survey of algorithms for constructing triangulations of various kinds.

In view of this situation, Aichholzer introduced the concept of a triangulation path as a tool that makes dynamic programming approaches feasible for counting triangulations and for constructing optimal triangulations according to decomposable optimality criteria (e.g. min-weight), hopefully significantly faster than the number of such triangulations. Experiments support such evidence [5]. In the rest of this section we provide a definition of this concept which slightly deviates from Aichholzer's in notation only, and we describe our results.

Notation. For two points p and q, we let $\overline{pq}$ denote the straight line segment connecting p and q (i.e. the convex hull of p and q). We also consider directed segments $\overrightarrow{pq}$.

For a polygon P we use $E(P)$ for the set of edges of P. Similarly, $E(T)$ denotes the set of edges of a triangulation T. In $E(P)$ and $E(T)$ edges are segments.

An *edge on* S is a segment with both endpoints in S. A set A of edges on S is called *plane in* (S, P), if all edges in A are contained in P, and these edges intersect at

[1] On the experimental side, in recent years a significant break-through has been obtained for the computation of min-weight triangulations, [6].

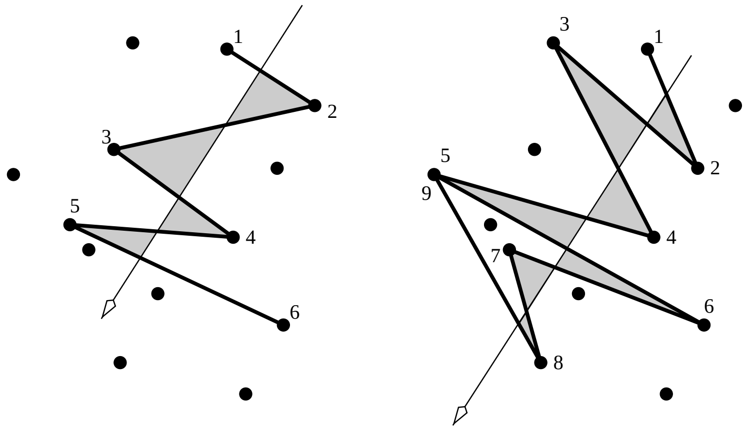

Fig. 2. Plane paths twining around a directed segment. The triangles that are required to be disjoint from S are shaded.

common endpoints only.

Plane paths on S twining around a segment. A *plane path* π *on* S is a sequence $(p_0, p_1, \ldots, p_k)$ of points in S such that no two consecutive points are the same, and the *edges* $\overline{p_{i-1}p_i}$, $1 \le i \le k$, *of* π are pairwise disjoint except for overlap in common endpoints. Points may repeat in the sequence! $E(\pi)$ denotes the set of edges $\{\overline{p_{i-1}p_i} : 1 \le i \le k\}$ of π.

Definition 1 *Let σ stand for a directed segment $\overrightarrow{pq}$ disjoint from S. We say that a plane path π* twines around *σ if (i) all edges of π intersect σ, and (ii) for each i, $1 \le i \le k-1$, the open triangle bounded by the consecutive pair of segments $\overline{p_{i-1}p_i}$ and $\overline{p_ip_{i+1}}$ and by σ is disjoint from S; see Figure 2.*

(i) and (ii) in Defi nition 1 imply that the points q_i of intersection of $\overline{p_{i-1}p_i}$ with σ form a monotone increasing sequence $(q_1, q_2, \ldots, q_k)$ on σ (i.e. $q_i \in \overline{pq_{i+1}}$ for all $1 \le i \le k-1$).x

Triangulation paths. We call a directed segment σ *generic w.r.t. a points-in-polygon-pair* (S, P), if $\sigma \subseteq P$ and $\sigma \cap S = \emptyset$.

Given a directed segment σ generic w.r.t. (S, P), we call a plane path π *twining around σ in* (S, P), if π twines around σ, and $E(\pi)$ is plane in (S, P). Note that in this case, $E(\pi)$ can always be completed to a triangulation of (S, P). Given (S, P) and σ there may be many paths twining around σ in (S, P).

Here comes the main result of Aichholzer [1] that makes these concepts remarkable and useful.

Lemma 2 *Let T be a triangulation of the points-in-polygon-pair (S, P), and let*

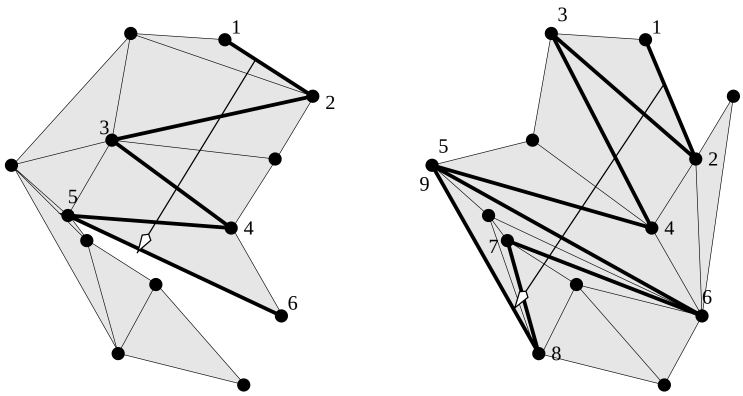

Fig. 3. Two triangulation paths.

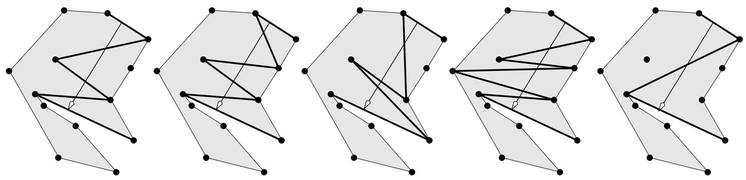

Fig. 4. A sample of paths in $\Pi_\sigma(S, P)$. Note that the union of no two such paths is plane – a consequence of Lemma 2.

σ be a generic directed segment connecting two points on distinct edges in $E(P)$. Then there is a unique path π twining around σ with $\{e', e''\} \subseteq E(\pi) \subseteq E(T)$, where e' and e'' are the edges in $E(P)$ carrying the endpoints of σ.

Hence, given (S, P), T, and σ as in the previous lemma, we can define the σ-*triangulation path*, or σ-*path* for short, $tp_\sigma = tp_\sigma(T) = tp_\sigma(S, P, T)$, as the unique plane path twining around σ in (S, P) with $\{e', e''\} \subseteq E(\pi) \subseteq E(T)$ (e' and e'' as in the lemma); see Figure 3

for examples. Given (S, P) and σ as before (without T), we consider now all σ-paths that can appear for triangulations. We denote by $\Pi_\sigma = \Pi_\sigma(S, P)$ the set of all plane paths π twining around σ in (S, P) with $\{e', e''\} \subseteq E(\pi)$. Note that *every* such path π can be extended to (possibly many) triangulations T of (S, P) with $E(\pi) \subseteq E(T)$, and for all these triangulations $\pi = tp_\sigma(S, P, T)$. In fact, Π_σ induces a partition of $\mathcal{T}(S, P)$ into $|\Pi_\sigma|$ classes of the form $\{T \in \mathcal{T}(S, P) : tp_\sigma(S, P, T) = \pi\}$, $\pi \in \Pi_\sigma$ (see Lemma 2). At this point it is maybe clear how the paths in Π_σ can be used to enumerate or count the triangulations in $\mathcal{T}(S, P)$ in a recursive manner. We refer to Aichholzer's paper for details [1].

Results. Our contribution is an algorithm for enumerating all triangulation paths in $\Pi_\sigma(S, P)$ – a crucial step in the efficient employment of σ-paths. The algorithm runs in time $O(|\Pi_\sigma(S, P)| \cdot n^3 \log n)$ using $O(n)$ working space. For the algorithm we introduce the notion of flips between such paths, and define a structure on all paths such that the reverse search approach can be applied, [2]. This result will be described in Section 2.

In Section 3 we refute Aichholzer's conjecture in [1] that points sets S in convex position (with P the boundary of the convex hull) maximize $|\Pi_\sigma(S,P)|$ in terms of n. To this end we describe a point set S, such that $(S, \mathit{conv}(S))$ exhibits $\Omega(2^{2n-\Theta(\log n)})$ triangulation paths (compared to $\binom{n-2}{\lfloor \frac{n}{2} \rfloor -1} = \Theta(2^{n-\Theta(\log n)})$ for S in convex position).

2 Flips on Paths and Reverse Search

We first recapitulate Lawson-flips in triangulations – an operation that allows simple, though not the most efficient algorithms for constructing (constrained) Delaunay triangulations. We will then derive a corresponding operation for triangulation paths which will lead us to the enumeration procedure.

Flips in triangulations. Consider a triangulation T of a points-in-polygon-pair (S,P). Let $e = \overline{pq}$ be an edge in $E(T) \setminus E(P)$. Edge e is adjacent to two empty triangles *conv*$\{p,q,x\}$ and *conv*$\{p,q,y\}$ in T. We call such an edge e a *flip edge in* T, if *conv*$\{x,p,y,q\}$ is a quadrilateral. If, moreover, the circle through the points p, q, and x contains the point[2] y, we call e a *Lawson edge in* T; otherwise, it is called *anti-Lawson edge*. The set of flip edges, of Lawson edges, and of anti-Lawson edges of T will be denoted by $F(T)$, $L(T)$ and $A(T)$, respectively; $F(T) = L(T) \dot\cup A(T)$.

For a flip edge e in T, we can obtain a new triangulation, *flip*(T,e), by substituting $\overline{xy}$ for e, (x and y as above); if e is a Lawson edge, this operation is called *Lawson-flip*.

Let C be a set of edges on S that is plane in (S,P). Then it is known that there is a unique triangulation T of (S,P) such that $L(T) \subseteq C \subseteq E(T)$, [4]; it is called the C*-constrained Delaunay triangulation of* (S,P), denoted by $\mathit{CDT}_C = \mathit{CDT}_C(S,P)$. This triangulation can be obtained by starting with an arbitrary triangulation containing all edges in C, and then successively applying Lawson-flips to edges not in C as long as possible.

We want to make this process more deterministic (obviously we may have many edges for Lawson-flips to choose from). To this end we agree on some total order $<$ on all edges on S. Given a set C of edges on S that are plane in (S,P), we define a

[2] This is equivalent to the fact that the circle through the points p, q, and y contains x. And it implies that *conv*$\{x,p,y,q\}$ is a quadrilateral.

function f_C from the set of all triangulations containing C to all such triangulations plus a special symbol $\perp$ by

$$T \mapsto \begin{cases} \perp & \text{if } L(T) \subseteq C \text{, and} \\ \mathit{flip}(T, e) & \text{otherwise,} \end{cases}$$
$$\text{where } e = \min(L(T) \setminus C).$$

Here, the minimum refers to the total order on edges we had agreed on. The C-constrained Delaunay triangulation is the only triangulation that maps to $\perp$, and it can be shown that f_C is acyclic (successive application of the function will not yield the initial argument). We can 'view' f_C as a tree with root CDT_C, and we can use a depth-first procedure for enumerating all triangulations that contain C, starting at the root CDT_C. Note that we can easily determine $f_C(T)$ for a triangulation T, and, moreover, we can enumerate the elements of $f_C^{-1}(T)$ in a canonical order. To this end we apply flips to anti-Lawson edges in T wherever possible, and check whether the resulting triangulation does indeed map to T. This can be done in time $O(|E(T)|) = O(n)$, provided appropriate data-structures have been prepared, [2]. This enumeration method is called *reverse search* [2], since successive application of f_C constructs CDT_C, and the enumeration turns this process around in order to search for all possible triangulations.

Searching for the σ-path of $\mathit{CDT}_\emptyset$. Let (S, P) and σ be fixed. We describe now a search for the σ-path of $\mathit{CDT}_\emptyset$, with the goal of turning around this process in order to enumerate Π_σ afterwards.

For a path $\pi \in \Pi_\sigma$ consider the triangulation $T = \mathit{CDT}_{E(\pi)}$. If $L(T) = \emptyset$, then $T = \mathit{CDT}_\emptyset$ and π equals $\mathit{tp}_\sigma(\mathit{CDT}_\emptyset)$, the σ-path we were searching for. Otherwise, we let $e = \min L(T)$, which has to be an edge of π, and we flip this edge e. The resulting triangulation $T' = \mathit{flip}(T, e)$ has a σ-path π' different from π, and we proceed with π'. We have just described the function $g : \Pi_\sigma \to \Pi_\sigma \cup \{\perp\}$

$$\pi \mapsto \begin{cases} \perp, & \text{if } L(T) = \emptyset \text{, and} \\ \mathit{tp}_\sigma(\mathit{flip}(T, \min L(T))), & \text{otherwise,} \end{cases}$$
$$\text{where } T = \mathit{CDT}_{E(\pi)}.$$

Lemma 3 (i) $g(\pi) = \perp$ *iff* $\pi = \mathit{tp}_\sigma(\mathit{CDT}_\emptyset)$.
(ii) g *is acyclic, in the sense that iterated application of* g *to* π *cannot result in* π.

PROOF. (i) Consider some $\pi \in \Pi_\sigma$ and let $T = \mathit{CDT}_{E(\pi)}$.
Note that $\pi = \mathit{tp}_\sigma(\mathit{CDT}_{E(\pi)})$, i.e. π is the unique σ-path of T. If $g(\pi) = \perp$, i.e. $L(T) = \emptyset$, then $T = \mathit{CDT}_\emptyset$, and thus $\pi = \mathit{tp}_\sigma(\mathit{CDT}_\emptyset)$. On the other hand, if $\pi = \mathit{tp}_\sigma(\mathit{CDT}_\emptyset)$, then $E(\pi) \subseteq E(\mathit{CDT}_\emptyset)$ and thus $T = \mathit{CDT}_{E(\pi)} = \mathit{CDT}_\emptyset$. Hence,

$L(T) = \emptyset$ and $g(\pi) = \perp$.
(ii) It is known that Lawson-flips are acyclic operations on the set of all triangulations. It remains to observe that, for π with $g(\pi) \neq \perp$, $CDT_{E(g(\pi))}$ can be obtained from $CDT_{E(\pi)}$ by a sequence of Lawson-flips. □

In other words, the lemma guarantees that the directed graph induced by g on Π_σ is a tree with root $tp_\sigma(CDT_\emptyset)$. Now we have to establish an effective computation of g and of its inverse g^{-1} defined by $g^{-1}(x) := \{\pi : g(\pi) = x\}$ for $x \in \Pi_\sigma \cup \{\perp\}$.

Given $\pi \in \Pi_\sigma$, we can determine $CDT_{E(\pi)}$ in time $O(n \log n)$ by known algorithms, [4]. We take a careful look at the changes from $tp_\sigma(T)$ to $tp_\sigma(flip(T, e))$; see Figure 5 for an illustration of the possible cases.

Lemma 4 *Let π be the σ-path of triangulation T, and let e be an edge in $F(T) \cap E(\pi)$. Let $\overline{xy}$ be the edge that replaces e in flip(T, e).*
(i) *$\overline{xy}$ intersects σ.*
(ii) *Let p', p, q, q' be the consecutive subsequence of π such that $e = \overline{pq}$ (e cannot be the first or last edge on π). W.l.o.g. let p and y be on the same side of σ (and, thus, q and x are on the same side according to (i)).*

The σ-path of flip(T, e) can be obtained by substituting p, q in π by one of the following sequences:

(I)	$p, x, y, q,$	*if* $x \neq p'$ *and* $y \neq q'$,
(II)	$y, q,$	*if* $x = p'$ *and* $y \neq q'$,
(III)	$p, x,$	*if* $x \neq p'$ *and* $y = q'$,
(IV)	*the empty sequence*,	*if* $x = p'$ *and* $y = q'$.

Special cases occur if pp' was a convex hull segment in the cases (II,IV) or qq' was a convex hull segment in the cases (III,IV). In this case we prefix the the obtained sequence with p, or respectively we suffix it with q (rule CHR).

PROOF. (i) Let p', p, q, q' be the consecutive subsequence of π such that $e = \overline{pq}$. One of the edges $\overline{px}$ or $\overline{py}$ must either be equal to $\overline{pp'}$ or it must lie in the angle formed by $\overline{pp'}$ and $\overline{pq}$ (since $conv\{p, q, x\}$ and $conv\{p, q, y\}$ are the triangles incident to $\overline{pp'}$ in T); say it is $\overline{px}$. If $\overline{px} = \overline{pp'}$ then x lies on the side of σ opposite to p; if $\overline{px} \neq \overline{pp'}$, then x must lie on the side of σ opposite to p, as well, since otherwise x is in the forbidden triangle bounded by $\overline{pp'}, \overline{pq}$ and σ (forbidden in the sense that x in the triangle would contradict π being a σ-path). By symmetry, y must lie on the side of σ opposite to q. Therefore x and y lie on opposite sides of σ and we are done.
(ii) If the sequences obtained by the transformations (by I-IV,CHR) define paths

that twine around σ, and whose segments are contained in $E(\mathit{flip}(T,e))$, we are done, because of the uniqueness property for σ-paths, see Lemma 2.

The final rule (CHR) guarantees that the obtained sequences start and end with the needed convex hull segments, in (I) we do not change in any case the two first and last points in the sequence, in (II,IV), if pp' is not on the convex hull nothing happens, otherwise by the rule the sequence starts with p, p' as required; the same argument applies for (III,IV) with q and qq' by symmetry. Note that p, y and q' lie on one side of σ while q, x and p' lie on the opposite side of σ. Hence in all sequences in (I-IV) consecutive points are distinct and lie on opposite sides of σ. In all four cases (I-IV), the segments defined by the new sequence are in $E(\pi) \cup \{\overline{px}, \overline{xy}, \overline{yq}\} \setminus \{\overline{pq}\} \subseteq E(\mathit{flip}(T,e))$, implying that there are no intersections in the interiors of these segments. We have proved that the new sequences define plane paths, whose segments are contained in $E(\mathit{flip}(T,e))$.

It remains to show that the defined paths twine around σ. First clearly $\overline{px}$, $\overline{yq}$ and $\overline{xy}$ intersect σ – we have already noticed that their endpoints lie on opposite sides of σ.

Thus we are left with the proof of the required emptiness property (Definition 1 (ii)). Let p'' be the predecessor of p' or be p if pp' is on the convex hull, and q'' be the successor of q' in π or q if qq' is on the convex hull. Let F_{abc} denote the open triangle spanned by $\overline{ab}$, $\overline{cb}$ and σ. In the case two of a, b, c are the same F_{abc} is defined to be the empty-set.

- $F_{p'px}$ should be empty in (I,III) and this is the case, because $F_{p'px} \subseteq F_{p'pq}$, since $\overline{px}$ lies in the angle formed by $\overline{pp'}$ and $\overline{pq}$;
- F_{pxy} should be empty in (I,III) and this is the case, since F_{pxy} is contained in the relative interior of $\mathit{conv}\{p, x, q, y\}$, which is empty by definition of flip;
- In (II,IV) $x = p'$ and $F_{p''xy}$ is contained in the union of $F_{p''p'p}$ and the relative interior of $\mathit{conv}\{p, x, q, y\}$, implying emptiness as required.

These are all the forbidden triangles that appear in the four cases, plus $F_{yqq'}$, F_{xyq}, and $F_{xyq''}$ that can be handled analogously by symmetry. We have used that the forbidden triangles $F_{p'pq}$ and $F_{p''p'p}$ of π were empty by definition of σ-path. □

Hence, $g(\pi)$ can be computed by first computing $T = \mathit{CDT}_{E(\pi)}$, then determining the minimal edge e in $L(T)$ (or detecting that $L(T) = \emptyset$), and then performing the adaptations described in Lemma 4. Altogether, these operations can be performed in $O(n \log n)$.

The crucial property that follows from Lemma 4 is that the σ-path of $\mathit{flip}(T,e)$ is determined by $\mathit{tp}_\sigma(T)$ and $\overline{xy}$, so it is independent from the remaining triangulation T. This will be essential, when we want to turn around the process for reverse search.

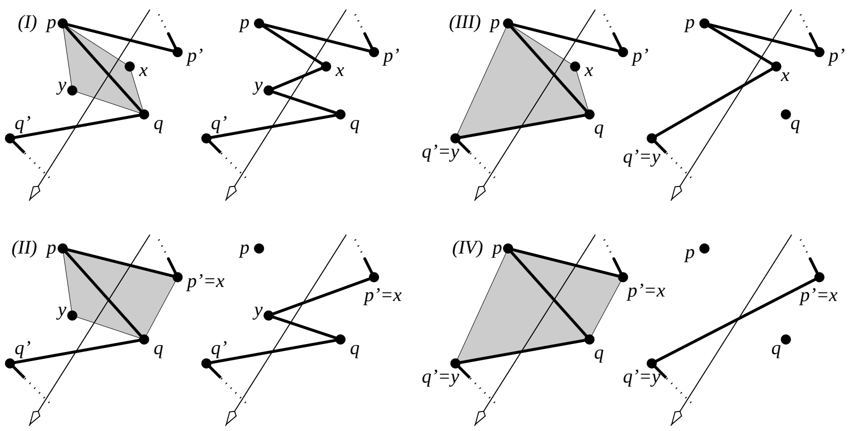

Fig. 5. The effect of flips to triangulation paths.

Turning it around for reverse search. Given a σ-path π, a *flip-pair for* π is an ordered pair $(\overline{pq}, \overline{xy})$ of edges on S such that (i) $\overline{pq} \in E(\pi) \setminus E(P)$, (ii) $E(\pi) \not\ni \overline{xy}$ intersects σ and $\overline{pq}$ in its relative interior, and it intersects no other edge of π in its relative interior, (iii) *conv*$\{p, x, q, y\}$ is a quadrilateral whose interior is disjoint from S, and (iv) $E(\pi) \cup \{\overline{px}, \overline{xq}, \overline{qy}, \overline{yp}\}$ is plane in (S, P). These are simply the necessary conditions for π to be extended to a triangulation where $\overline{pq}$ and $\overline{xy}$ can be flipped. Given such a flip-pair, we can perform a flip according to the rules described in Lemma 4 (ii); the resulting σ-path is denoted by *flip*$(\pi, \overline{pq}, \overline{xy})$.

We can now enumerate $g^{-1}(\pi)$, $\pi \in \Pi_\sigma$, in canonical order by going through the list of all edges e' on S in increasing order (according to the total order we had agreed on), and testing whether there is an edge $e \in E(\pi)$ such that (e, e') is a flip-pair for π and, if so, whether $g(\textit{flip}(\pi, e, e')) = \pi$. Note that there can exist at most one such e for each e'. There are $O(n^2)$ edges e' on S, and each such edge can be checked in $O(n \log n)$.

For g we compute the *CDT* and read out its σ-path (for the latter, see [5]), to check for a flip-pair we mainly need to test for segment intersections between $E(\pi)$ ($O(n)$ edges, see [1]) and the few segments relevant for the definition; *flip*(π, e, e') is by the lemma simply an operation on the path sequence and has at most linear cost in n, thus all operations can be done in $O(n \log n)$ time.

Hence, we can enumerate $g^{-1}(\pi)$ in time $O(n^3 \log n)$. In particular, within this time bound we can determine the first path in $g^{-1}(\pi)$, and given some π' in $g^{-1}(\pi)$, we can determine its successor in the canonical order, or determine that it is the last element in this order.

The ingredients for reverse search are set, and we can conclude:

Theorem 5 *Given a points-in-polygon pair (S, P) and a generic directed segment σ connecting two points on distinct edges in $E(P)$, we can enumerate $\Pi_\sigma(S, P)$ in time $O(|\Pi_\sigma(S, P)| \cdot n^3 \log n)$ and space $O(n)$, $n = |S|$.*

3 Point Sets with many Paths

Surprisingly, although in [1] Aichholzer conjectured that the maximum number of triangulation paths is attained for a configuration of points in convex position, our best counterexample to his conjecture (with the largest number of paths) is to be found in his own aforementioned paper. While convex position allows $O(2^n)$ triangulation paths, the configuration to be exhibited has $\Omega(2^{2n-\Theta(\log n)})$ such paths.

We explain it here with the aid of Figure 6. For convenience, assume $n \geq 16$ is a multiple of 8, where $n = |S|$. There are $n/2$ points lying on a horizontal line (which we take as the x-axis) on the right side of l (here l is the supporting line of σ, with vertical orientation). There are two groups of $n/4$ points on two lines on the left side of l. The two left lines are arranged such that they 'split' the set on the right side in two parts of equal size. In order to move the points in general position, we perform a slight perturbation. Put $P = \mathit{conv}(S)$; $\mathit{conv}(S)$ has three vertices. We show that we count roughly 4^n paths for this choice of S, P, σ. Consider the points $\{a, b, c, d, e, f\}$ as in Figure 7 (i.e. a is the extreme right point of S, etc.). Let

- A and B be any subsets of size $n/8$ of the right group of points on the x-axis, with $a \in A$, $a \in B$;
- C and D be any subsets of size $n/8$ of the left group of points on the x-axis, with $b \in C$, $b \in D$;
- E and F be any subsets of size $n/8$ of the set of points on the line of negative slope, with $\{c, d\} \subseteq E$, $d \in F$;
- G and H be any subsets of size $n/8$ of the set of points on the line of positive slope, with $f \in G$, $\{e, f\} \subseteq H$.

For each such choice of A, B, C, D, E, F, G, H, consider the path which starts at a, then joins points from A and E (alternating), then from C and F, then from D and G, and finally from B and H and ends at a; the first edge of the path is $\overline{ac}$ and the last one $\overline{ea}$, and these two edges contain the two endpoints of σ respectively. The points are numbered (below or above) in Figure 6 according to their order in the path. The way the path goes is sketched in Figure 7.

Put $N = |\Pi_\sigma(S, P)|$. Denote by $H(q) = -q \log q - (1-q)\log(1-q)$ the binary entropy function, where $\log$ stands for the logarithm in base 2. From the well-known estimate

$$\binom{n}{\alpha n} = \Theta(n^{-\frac{1}{2}} 2^{H(\alpha) n}),$$

we get that for $0 < \beta \leq \alpha$,

$$\binom{\alpha n}{\beta n} = \Theta(n^{-\frac{1}{2}} 2^{\alpha H(\frac{\beta}{\alpha}) n}).$$

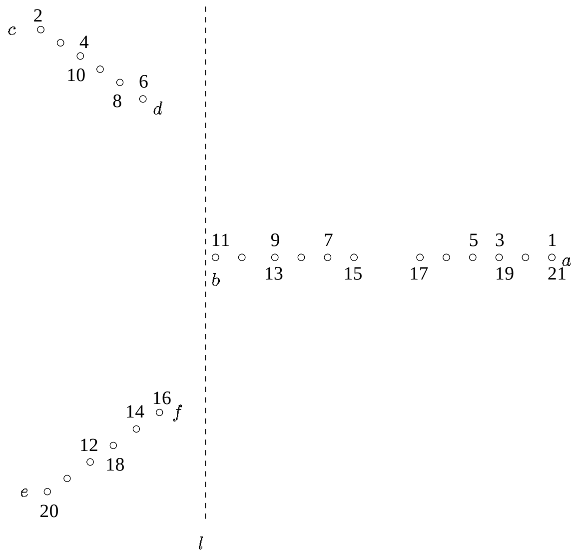

Fig. 6. A point configuration with many paths.

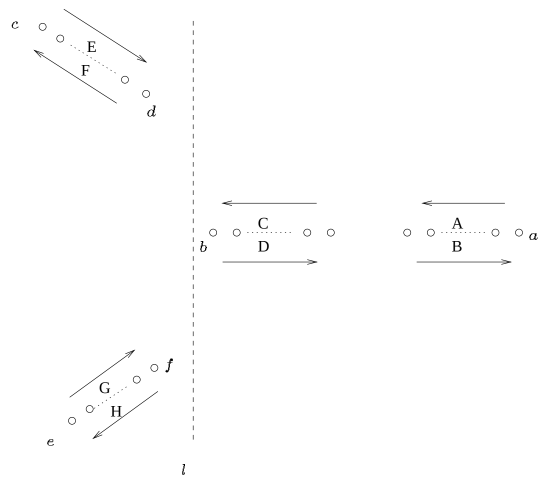

Fig. 7. Path description.

Let N_X be the number of possible choices of X for $X \in \{A, B, C, D, E, F, G, H\}$. We have

$$N_A = N_B = N_C = N_D = N_F = N_G = \binom{\frac{n}{4} - 1}{\frac{n}{8} - 1},$$

$$N_E = N_H = \binom{\frac{n}{4} - 2}{\frac{n}{8} - 2},$$

and

$$N \geq N_A \cdot N_B \cdot N_C \cdot N_D \cdot N_E \cdot N_F \cdot N_G \cdot N_H.$$

Ignoring the inverse polynomial factors that arise from the above estimates and from rewriting everything in terms of $\binom{\frac{n}{4}}{\frac{n}{8}}$, we get

$$N \gtrsim \binom{\frac{n}{4}}{\frac{n}{8}}^8 \approx (2^{H(\frac{1}{2})\frac{n}{4}})^8 = 2^{2n}.$$

Taking into account the factors we ignored, we obtain the following lower bound.

Theorem 6 *For infinitely many n, there exists a configuration S of n points, such that $(S, \text{conv}(S))$ has*

$$\Omega(2^{2n-\Theta(\log n)})$$

triangulation paths.

4 Open Problems

We have presented an algorithm that enumerates triangulation paths in time $O(n^3 \log n)$ per path reported. We make no claim that this procedure is more efficient than the back-tracking approach from [1,5] in practice. Still, our reverse search algorithm is the first method with a provable polynomial time bound per path. A run-time analysis for the back-tracking algorithm is still an open problem.

The 'ratio' between the number of paths and the number of triangulation is the critical parameter for the effectiveness of the path-based methods. Good upper bounds for the number of triangulation paths are still missing, but we have shown that points in convex position do not deliver the worst case examples. Such a bound should lie (ignoring polynomial factors) between our lower bound construction value of 2^{2n} and the actually best upper bound for the number of triangulations of 2^{8n} [7]. The latter comes from an encoding method that both shows and exploits the fact that for a given triangulation at least one third of all points have degree at most 6. Can this be used for bounding the number of paths? Is it possible to have configurations where there is a huge number of paths all extending to few triangulations? One open question in this direction is whether on average, the configuration in Section 3 has an exponential number of triangulations per path.

References

[1] O. Aichholzer, The path of a triangulation, *Proc. 15th Ann. ACM Symp. on Comput. Geom.* (1999) 14–23.

[2] D. Avis and K. Fukuda, Reverse search for enumeration, *Discrete Appl. Math.* **65** (1996) 21–46.

[3] M. Bern and D. Eppstein, Mesh generation and optimal triangulation, in *Computing in Euclidean Geometry* (D.-Z. Du and F.K. Hwang, Eds.), World Scientific (1992).

[4] L.P. Chew, Constrained Delaunay triangulations, *Algorithmica* **4** (1989) 97–108.

[5] B. Erkinger, Struktureigenschaften von Triangulierungen, Diplomarbeit, TU Graz (1998).

[6] R. Beirouti and J. Snoeyink, Implementations of the LMT heuristic for minimum weight triangulation, *Proc. 14th Ann. ACM Symp. on Comput. Geom.* (1998) 96–105.

[7] M.O. Denny and C.A. Sohler, Encoding a triangulation as a permutation of its point set, *Proc. 9th Canadian Conf. on Comput. Geom.* (1997), 39–43.