A Combinatorial Upper Bound on the Length of Twang Cascades *

Leon Sering

Abstract

Damian et al. [2] introduced an intuitive type of transformations between simple polygons on a finite set of points in the plane. Each transformation consists of a sequence of atomic modifications of two types, called *stretches* and *twangs*. They proved that, for a given set of n points, the space of these simple polygons is connected by $O(n^2)$ such transformations.

We solve an open question of Damian et al. concerning a combinatorial upper bound on the length of these sequences. To this end, we show that the length of a twang cascade is bounded by $n^3/2$.

1 Introduction

This paper studies simple polygons on a fixed set S of n points in the plane. A *simple polygon* on S is a crossing-free cycle of straight line segments, each connecting two points of S such that every point of S is visited exactly once. In the following we abbreviate "simple polygon" by "polygon".

Sampling Polygons We consider the following challenging open problem:

Given a point set S as input, the task is to generate a random polygon on S with *uniform distribution*, that is, if r is the number of polygons on S, we want to choose one with a probability of $\frac{1}{r}$. So far no algorithm is known to do this efficiently.

There are, however, various partial solutions for this problem. On the one hand, there are efficient generators for specific subsets of simple polygons such as x-monotone polygons [5] or star-shaped polygons [1]. On the other hand, there are algorithms that produce all possible polygons with positive probability but they are not generated uniformly at random. For example, the *2-Opt-Moves* algorithm [1] produces a random permutation of S and then removes all self-intersections step by step.

Random Walk Approach A different approach is to repeatedly apply random modifications to a polygon, a so-called *Markov chain Monte Carlo sampling*. We define a class of transformations, each turning one polygon into another, such that for every polygon the number of applicable transformations is bounded polynomially. Then the algorithm do a *random walk* on the *transformation graph*, which is the directed graph whose vertices are the polygons on S and where two vertices are connected by an directed edge whenever there is a transformation that turns one polygon into the other. If the transformation graph is connected, one can choose transition probabilites such that the distribution of the random walk converges to the uniform distribution.

The simplest transformation is the 2-flip. It removes two edges of the polygon and adds two other edges in order to get a new polygon. Unfortunately, there are point sets where the flip graph is not connected [4], which is shown by the existence of an isolated vertex.

For the k-flip, $k \geq 3$, it is unknown whether the corresponding flip graph is always connected. Hernando et al. [3] give a good overview of different types of flips and of subclasses of polygons that are connected by them.

2 Stretches and Twangs

Damian et al. [2] imagine the polygon as an elastic band that is attached to the points in S at its vertices. They introduce a new transformation, called *forward move*, based on the idea of deforming this elastic band. It consists of two different types of atomic operations, namely *stretches* and *twangs*.

Polygonal Wrap Neither stretches nor twangs are simplicity preserving. They produce an object called *polygonal wrap*, which does not have any proper crossings but can be self-touching.

Definition 1 ([2]) *A* polygonal wrap *W on S of length $m \geq n$ is a cyclic polygonal chain such that*

(W1) The wrap only bends at points of S.

(W2) Every point in S is visited at least once.

(W3) The wrap does not contain any proper crossings, i.e., there exists an arbitrarily small perturbation of the vertices of W that makes the cyclic polygonal chain non-self-intersecting.

*This is a short version of a master thesis with the same title submitted to the Institute of Computer Science at the University of Würzburg in 2016. The full version is available at http://www1.pub.informatik.uni-wuerzburg.de/pub/theses/2016-sering-master.pdf

If a point is visited more than once, we call it a point in multiple contact*, and a subsequence* (a, b, a) *is a* needle-pin *at* b*. Furthermore, we say that a line segment* ab does not properly cross W *if there is an arbitrarily small perturbation of* W *and* ab *such that* W *and* ab *do not intersect. For a subsequence* (a, b, c)*, we call the cone in the minor arc the* convex side *and the complement the* reflex side*.*

Twang A twang is an operation that transforms one polygonal wrap into another. Informally, we choose a point in multiple contact that is not a needle-pin, detach one contact from the point and let the band snap back. The snapping band does not cross other vertices but attaches to them instead; see Figure 1a.

Definition 2 ([2]) *The operation* $\mathbf{Tw}(abc)$ *is defined for a subsequence* (a, b, c) *of a polygonal wrap* W *whenever the following three conditions hold:*

(T1) b *is in multiple contact.*

(T2) b *is not a hairpin.*

(T3) (a, b, c) *does not surround any other visits of* b*.*

Then $\mathbf{Tw}(abc)$ *replaces the subsequence* (a, b, c) *in* W *by* $\mathbf{sp}(abc)$*, where* $\mathbf{sp}(abc)$ *is the shortest path from* a *to* c *inside of the triangle* $\triangle abc$ *that does not properly cross* W*.*

As long as there is at least one point in multiple contact, we can apply a twang to the wrap. Furthermore, a polygonal wrap without any points in multiple contact is a polygon. Although a twang might produce more multiple contacts in the process, but we will show that repeated twanging will in fact terminate and restore a polygon.

Stretch Informally, a stretch is the operation of taking an edge e of the elastic band and attaching it to a point p. Similar to a twang, the band does not cross other points but instead wraps around them; see Figure 1b.

Definition 3 ([2]) *Given a polygonal wrap* W*, we say an edge* e *of* W is visible *to a point* p *if there is a point* x *in the interior of* e *such that the line segment* px *does not properly cross* W*. The point* x *is called the* spotted point*. The operation* $\mathbf{St}(e, p)$ *is defined for any edge* $e = (a, b)$ *of* W *and any vertex* $p \in S$ *if* e *is visible to* p*. To execute* $\mathbf{St}(e, p)$*, we replace* (a, b) *by* $(\mathbf{sp}(axp), \mathbf{sp}(pxb))$*, where* x *is the spotted point.*

In other words, we first add the spotted point x as a pseudo-vertex to the polygonal wrap and replace (a, b) by (a, x, p, x, b). Afterwards, we twang at x twice such that x can be removed again.

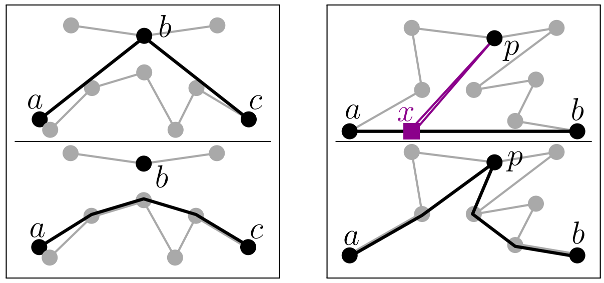

(a) Twanging at b: (a, b, c) is replaced by $sp(abc)$.

(b) Stretching (a, b) to p. Add (x, p, x) temporarily and twang at x twice.

Figure 1: Twang and Stretch.

A stretch is the main tool to modify a polygon and turn it into another polygon. But since it always creates at least one multiple contact, we need to apply a sequence of twangs in order to obtain a polygon.

Twang Cascade and Forward Move Next we define the transformation for the random walk.

Definition 4 ([2]) *Given a polygon* P *with an edge* e *and a point* p *such that* p *can see* e *and the spotted point* x *lies on the reflex side of the visit* (u, p, v)*, a* forward move *consists of a stretch* $\mathbf{St}(e, p)$ *followed by a so-called* twang cascade*, which starts with the twang* $\mathbf{Tw}(vpw)$ *and repeatedly twangs as long as there are vertices in multiple contact.*

In order to guarantee reversibility of the transformation, we also consider the time-reversal of a forward move and call it *reverse move.*

Pocket Reduction The main result of Damian et al. states the following:

Theorem 1 *The transformation graph combining forward and reverse moves is connected and has a diameter of* $O(n^2)$*. Each node has degree* $\omega(n)$*.*

The key idea is to transform every polygon to a canonical polygon (see Figure 2a) by reducing one pocket after the other until there is only one non-reduced pocket left. This can be done by $O(n^2)$ forward moves.

3 Upper Bound on the Length of Twang Cascades

Damian et al. did not give a combinatorial upper bound on the length of a twang cascade and, therefore, no polynomial bound on the running time of a forward move is known.

They showed, however, that there are forward moves that have twang cascades of length $\Theta(n^2)$ as shown in Figure 2b. In the following we will solve this

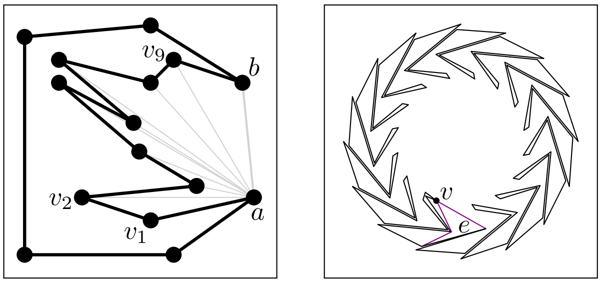

(a) Canonical polygon P_a with non-reduced pocket ab and vertices occurring clockwise around a.

(b) $\mathbf{St}(e, v)$ initiates a forward move twanging multiple times around the center.

Figure 2: Canonical polygon and quadratic twang cascade.

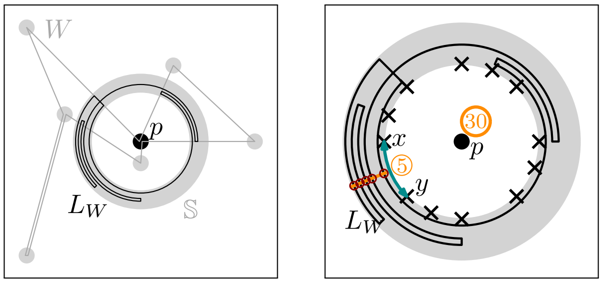

(a) The radial loop L_W is created by projecting the polygonal wrap W to the unit circle.

(b) $[xy]$ appears five times in L_W.
Hence $\mathbf{c}_W([xy]) = 5$.

Figure 3: Radial loop and coins.

open question and show that the length of a twang cascade is bounded by $O(n^3)$.

Given a set S of $n > 2$ points in the plane in general position and a polygonal wrap W that is created by a stretch on a polygon, we fix a point $p \in S$ and give an upper bound on the number of twangs that can occur at p in a twang cascade. We will show that this number is bounded by $O(n^2)$.

Markers, Elementary Arcs and Radial Loop For simplicity, consider p to be at the origin, that is, $p = (0, 0)$. First, we radially project every point in $S \backslash \{p\}$ to the unit circle.

Definition 5 *Consider the* radial projection*:*

$$\Pi \colon \mathbb{R}^2 \backslash \{(0,0)\} \to \mathbb{S}, \quad q \mapsto \frac{q}{\|q\|}$$

Here $\mathbb{S}$ is the unit circle. We define

$$S^+ := \Pi(S \backslash \{p\}), \quad S^- := -S^+, \quad S^\pm := S^+ \cup S^-.$$

We call the elements of $S^\pm$ markers*. For each point $q \in S \backslash \{p\}$, we write q^+ for $\Pi(q)$ and q^- for $-\Pi(q)$. The bijective map $\iota : \mathbb{S} \to \mathbb{S}$ with $x \mapsto -x$ is called* central inversion*.*

The arcs between two neighbouring markers play an important role throughout this section.

Definition 6 *Let $x, y \in S^\pm$ be two markers. If the minor arc between x and y does not contain any other marker of $S^\pm$, we call it an* elementary arc *and write $[xy]$. Let Λ be the set of all elementary arcs.*

Since we assume that S is in general position, S^+ and S^- are disjoint, and therefore $|S^\pm| = 2n-2$. This is also the number of elementary arcs. The central inversion induces a bijection on the markers $q^+ \mapsto q^-$ and $q^- \mapsto q^+$ for all $q \in S \backslash \{p\}$ and, therefore, a bijection on the set of elementary arcs Λ as well:

$$[xy] \qquad \mapsto \qquad [\iota(x)\iota(y)]$$

Next, we want to count how often the polygonal wrap goes around p in specific directions. This is well defined except for the spots where the wrap goes through p. For our purpose it is important to handle these cases as if the wrap went around the reflex side of p. We radially project the polygonal wrap to the unit circle in the following way.

Definition 7 *The* radial projection *of an edge (q, r) of W with $q \neq p$ and $r \neq p$ is the sequence of all elementary arcs between q^+ and r^+ on the minor arc. For a subsequence (q, p, r) in W, the* radial projection *is the sequence of the elementary arcs between q^+ and r^+ on the major arc. If we incrementally replace every edge of W by its radial projection, we get the* radial loop *L_W. This is visualized in Figure 3a.*

Coin System Next, we introduce an integral potential, called coins. We assign as many coins to an elementary arc as often as the radial loop winds around p in the direction of this elementary arc; see Figure 3b. Whenever we twang at p, we remove at least two coins. Since the current number of coins is always non-negative, the number of twangs at p is bounded by the initial number of coins divided by two.

Definition 8 *For an elementary arc $[xy]$ let $\mathbf{c}_W([xy])$ be the number of times $[xy]$ appears in the radial loop L_W. We assign $\mathbf{c}_W([xy])$ many coins to $[xy]$. Furthermore, let $\mathbf{c}_W(p)$ be the total number of coins on the unit circle, that is,*

$$\mathbf{c}_W(p) := \sum_{[xy] \in \Lambda} \mathbf{c}_W([xy])$$

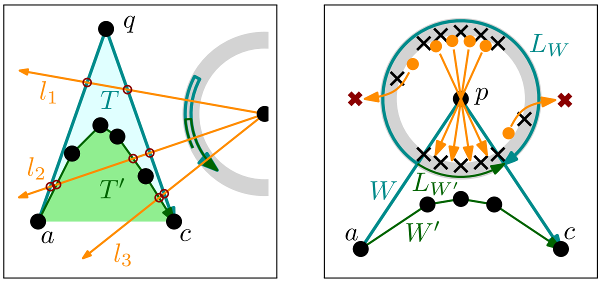

(a) l_1 crosses T but not T'. l_2 crosses both twice. l_3 crosses both once.

(b) The coin movement if we twang at p.

Figure 4: A twang only decreases the number of coins.

The Combinatorial Upper Bound In this section we want to prove the combinatorial upper bound on the length of twang cascades by bounding the number of twangs at each point p.

The following theorem states that the number of coins does not increase when we execute a twang.

Theorem 2 *A twang $W \to W'$ at some point $q \in S \backslash \{p\}$ never increases the number of coins at p. In other words, $\mathbf{c}_{W'}(p) \leq \mathbf{c}_W(p)$.*

The proof uses the fact that the twanged subsequence T' is in convex position and lies inside the original subsequence T; see Figure 4a. If T' goes around p in the direction of an elementary arc then so did T. Therefore, the number of coins on each elementary arc can only decrease. Note that a special case occurs when p is a neighbour of q. We omit the details.

The next theorem considers the case of twanging at p. In order to bound the number of twangs by the number of coins, we need to make sure that the number of coins decreases each time we twang.

Theorem 3 *Every twang $W \to W'$ at p decreases the number of coins at p by at least two. In other words, $\mathbf{c}_{W'}(p) \leq \mathbf{c}_W(p) - 2$.*

One can imagine that all elementary arcs $[xy]$ on the reflex side give one of their coins to $\iota([xy])$, the elementary arc on the opposite side, if $\iota([xy])$ lies on the convex side. But there are at least two elementary arcs whose opposite arcs are on the reflex side as well. Therefore, one coin from each of them is removed from the system. How the coins move is illustrated in Figure 4b.

In the next step we give an upper bound on the number of coins on each point p at the beginning of a twang cascade.

Lemma 4 *Let P be a polygon on S, and $p \in S$. Then $\mathbf{c}_P(p) \leq n^2$.*

This is true due to the fact that each of the n edges can at most contribute one coin to at most n elementary arcs.

Lemma 5 *A stretch does not add more than $O(n)$ coins to p.*

This holds because a stretch means basically to add two edges followed by some twangs.

In summary, there are no more than $O(n^2)$ coins on p at the beginning of a twang cascade. Therefore, only $O(n^2)$ twangs can occur at p. Summing over all n points we receive our final result:

Theorem 6 *In every forward move, the length of a twang cascade is bounded by $O(n^3)$.*

If we count more carefully and consider a lower bound of coins for polygons, we can show that the length of a twang cascade is in fact bounded by $n^3/2$.

4 Open Problems

We still need to solve several problems to obtain an efficient random generator for polygons.

1. Can reverse moves be computed efficiently?
2. Is the transformation graph still connected if we allow only forward moves?
3. Is the random walk rapidly mixing?

We assume the first question to be answered with no and the last with yes.

References

[1] T. Auer and M. Held. RPG – heuristics for the generation of random polygons. *Proc. 8th Canad. Conf. Comput. Geom (CCCG'98)*, pages 38–44, 1998.

[2] M. Damian, R. Flatland, J. O'Rourke, and S. Ramaswami. Connecting polygonizations via stretches and twangs. *Theory of Computing Systems*, 47(3):674–695, 2010.

[3] C. Hernando, M. E. Houle, and F. Hurtado. On local transformation of polygons with visibility properties. *Theoretical Computer Science*, 289(2):919–937, 2002.

[4] M. E. Houle. On local transformation of polygons. In *Proc. Computing: 2nd Australasian Theory Symp. (CATS'96)*, volume 18 of *Australian Comput. Sci. Commun.*, pages 64–71, 1996.

[5] C. Zhu, G. Sundaram, J. Snoeyink, and J. S. Mitchell. Generating random polygons with given vertices. *Computational Geometry*, 6(5):277–290, 1996.