Local Complexity of Polygons

Fabian Klute[*1], Meghana M. Reddy[†‡2], and Tillmann Miltzow[§1]

[1]ETH Zürich, Department of Computer Science
[2]Utrecht University, Information and Computing Science
Department

January 20, 2021

Abstract

Many problems in Discrete and Computational Geometry deal with
simple polygons or polygonal regions. Many algorithms and data-structures
perform considerably faster, if the underlying polygonal region has low
local complexity. One obstacle to make this intuition rigorous, is the
lack of a formal definition of local complexity. Here, we give two possible
definitions and show how they are related in a combinatorial sense. We say
that a polygon P has *point visibility width* $w = [\![\mathtt{pvw}]\!]$, if there is no point
$q \in P$ that sees more than w reflex vertices. We say that a polygon P has
chord visibility width $w = [\![\mathtt{cvw}]\!]$, if there is no chord $c = \mathrm{seg}(a, b) \subset P$ that
sees more than w reflex vertices. We show that

$$[\![\mathtt{cvw}]\!] \leq [\![\mathtt{pvw}]\!]^{O([\![\mathtt{pvw}]\!])},$$

for any simple polygon. Furthermore, we show that there exists a simple
polygon with
$$[\![\mathtt{cvw}]\!] \geq 2^{\Omega([\![\mathtt{pvw}]\!])}.$$

1 Introduction

In Discrete and Computational Geometry we study many problems with respect
to the input size n and other natural parameters. One famous example is the
computation of the convex hull of a set of points in the plane. While $\Theta(n \log n)$
time is worst case possible, this can be improved to $\Theta(n \log h)$, where h is the
number of vertices on the convex hull [4]. Here, the number of vertices on
the convex hull is a natural parameter to study this problem. We also say
sometimes that the algorithm is output-sensitive. Another famous example, is
the spread Δ of a set of points in the plane. That is the ratio between the largest
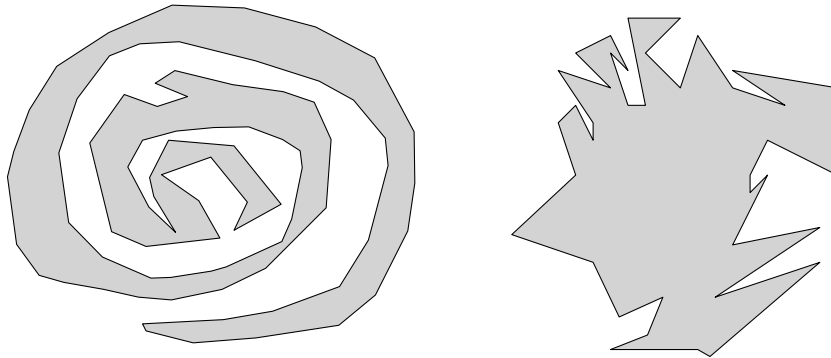
Figure 1: The polygon on the left has intuitively lower local complexity than on the right.

and the smallest distance, between any two points. Efrat and Har-Peled were the first to find an approximation algorithm for the art gallery problem under the assumption that the underlying set of vertices has bounded spread [2]. A third example is the number of reflex vertices of a polygon. This parameter gave raise to some FPT algorithms for the art gallery problem [1].

In this work, we introduce two new parameters that are meant to capture rigorously the idea of local complexity. Consider the polygons shown in Figure 1, most researchers would probably agree that the polygon on the left has lower local complexity than the polygon on the right. Yet, it is not straightforward how to define this rigorously in a mathematical sense.

Here, we give two possible definitions and show how they are related in a combinatorial sense. We say that a polygon P has *point visibility width* $w = [\![\texttt{pvw}]\!]$, if w is the smallest number such that there is no point $q \in P$ that sees more than w reflex vertices. We say that a polygon P has *chord visibility width* $w = [\![\texttt{cvw}]\!]$, if w is the smallest number such that there is no chord $c = \text{seg}(a, b) \subset P$ that sees more than w reflex vertices.

We show the following theorem.

Theorem 1. *For every polygon with chord visibility width $[\![\texttt{cvw}]\!]$ and point visibility width $[\![\texttt{pvw}]\!]$, it holds that*

$$[\![\texttt{pvw}]\!] \leq [\![\texttt{cvw}]\!] \leq [\![\texttt{pvw}]\!]^{O([\![\texttt{pvw}]\!])}.$$

Moreover, there are polygons such that

$$[\![\texttt{cvw}]\!] \geq 2^{\Omega([\![\texttt{pvw}]\!])}.$$

Note that Hengeveld and Miltzow already defined the notion of chord visibility width in a very similar way [3]. Specifically, they showed that the art gallery problem admits an FPT algorithm with respect to chord visibility width. For a parameter to be interesting to study, we usually have three criteria.

naturalness: Although there is no definition of what it means to be mathematically natural, many researchers seem to have a common understanding of this notion.

relevance: The parameter is at least for some fraction of instances reasonably low.

profitable: Using the parameter, we should be able to design better algorithms and prove useful run time upper bounds.

We believe that both parameters are mathematically natural. Theorem 1 indicates that the chord visibility width can be exponentially larger than the point visibility width. Thus we would expect that chord visibility width is potentially more profitable. We would expect that both parameters are equally relevant as the example that we give is fairly contrived. The remainder of this paper is dedicated to proving Theorem 1.

2 Chord visibility width vs Point visibility width

We prove Theorem 1 in two parts. First, we show the second half of the theorem in Section 2.1 by constructing a polygon for which it holds that $[\![\texttt{cvw}]\!] \geq 2^{\Omega([\![\texttt{pvw}]\!])}$. Second, we show the first half of Theorem 1 in Section 2.2 by analysing how the reflex vertices visible from a chord in a simple polygon P restrict each others vision and relating this to the point visibility width of the polygon.
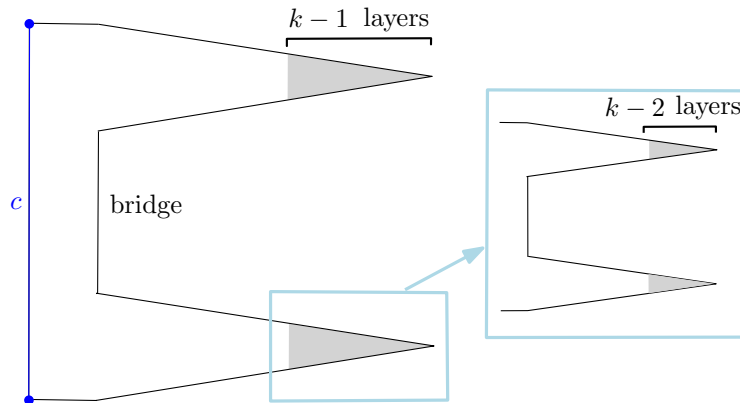
2.1 Lower bound

Figure 2: Construction of the Iterated Comb.

We construct a polygon P, called the *Iterated Comb*, see Figure 2. In the following, let $k \in \mathbb{N}$. The Iterated Comb consists of k *layers*, each layer consists of two *spikes* and each spike further splits into two more spikes in the subsequent layer. Observe, that the entire polygon is visible from the chord connecting the two left-most points of the polygon. The distance between consecutive spikes in a layer, referred to as the *bridge*, is adjusted such that if at least one vertex in the interior of a spike is visible from a point p on c, then p cannot see any interior vertex of any other spike. This property is achieved by stretching the bridges vertically. More specifically, for $1 < i \leq k$, the length of the bridge of the i^{th} layer is increased such that the property holds for layer i and then the bridge of the previous layer is adjusted accordingly. By iteratively stretching the

3

bridges from the last layer to the first layer, it can be ensured that the property holds for every layer. This property is illustrated in Figure 3 for $k = 2$. In the first layer, the point p sees an interior vertex of the first spike and no interior vertex of the second spike. Similarly, in the second layer point p sees an interior vertex of the second spike and no interior vertex of the first spike.
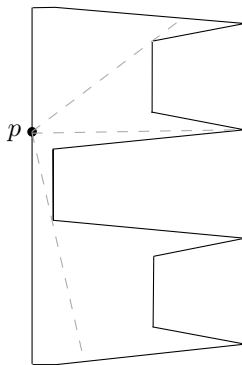

Figure 3: Point p sees interior points of at most one spike of any layer

Chord visibility width of the Iterated Comb   Clearly, the chord which sees the highest number of reflex vertices is the chord defined by the two left-most vertices. Let this chord be c. The number of reflex vertices of the first layer visible from c is two. Similarly, the number of reflex vertices of the i^{th} layer visible from c is 2^i. Summing up over all k layers, the number of reflex vertices visible from c is $\Theta(2^{k+1})$, and hence $[\![\mathtt{cvw}]\!] = \Theta(2^{k+1})$.

Point visibility width of the Iterated Comb

Claim 1. *Chord c contains at least one of the points in P which see the highest number of reflex vertices of P.*

Proof. Let q be a point in polygon P which sees the highest number of reflex vertices of P. Let p be a point on chord c which has the same y-coordinate as q. Assume $p \neq q$. Let r be a reflex vertex visible from q. Since P is monotone with respect to y-axis, the triangle pqr must be empty. This implies that r is visible from p as well. Hence, the point p also sees the highest number of reflex vertices in P since p sees at least as much as q. Refer to Figure 4 for an illustration. $\square$

Without loss of generality, assume the point with highest visibility is the topmost point on c, denoted by p. Both the reflex vertices in layer one are visible from p. In each subsequent layer, p can see the reflex vertices that are in the interior of the first spike, which is two reflex vertices, p cannot see any of the other reflex vertices in the other spikes by construction. Summing it up, we can conclude that $2k$ reflex vertices are visible from p, and thus $[\![\mathtt{pvw}]\!] = 2k$. Hence the Iterated Comb has $[\![\mathtt{cvw}]\!] \geq 2^{\Omega([\![\mathtt{pvw}]\!])}$.
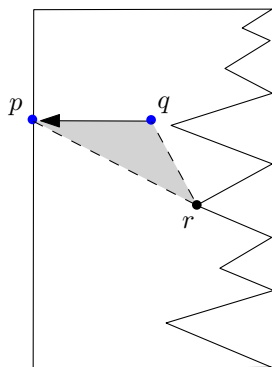
4

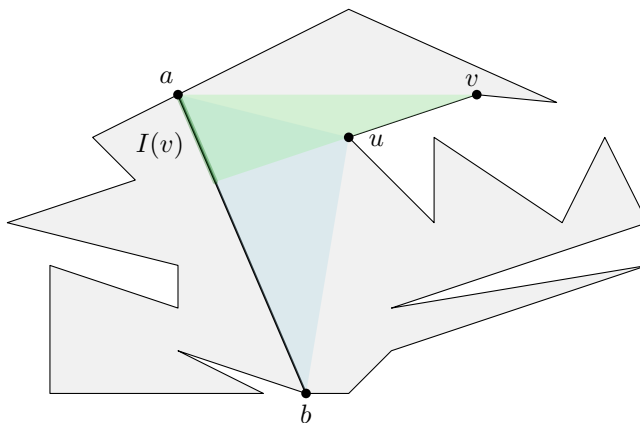Figure 4: Point p sees all the reflex vertices visible from q

Figure 5: The vertex v sees a subinterval $I(v) \subseteq s$ which is restricted by a and u.

2.2 Upper bound

Next, we show that we can upper bound the chord visibility width in terms of the point visibility width.

To this end, we prove the following lemma.

Lemma 2. *For every simple polygon, it holds that*

$$\llbracket \mathtt{cvw} \rrbracket \leq \llbracket \mathtt{pvw} \rrbracket^{O(\llbracket \mathtt{pvw} \rrbracket)}.$$

The rest of this paragraph is dedicated to the proof of Lemma 2.

For that purpose assume, we are given a simple polygon P together with a chord $s \subset P$. Furthermore, we assume that no point in P sees more than $k = \llbracket \mathtt{pvw} \rrbracket$ reflex vertices of P. Let us denote by R the set of all reflex vertices that see at least one point of $s = \mathrm{seg}(a, b)$. Furthermore, we also include the two endpoints of s in the set R. As P is a simple polygon it holds that every reflex vertex $r \in R \setminus \{a, b\}$ sees a subsegment $I(r) \subseteq s$. For convenience, we also call $I(r)$ an *interval*.

Note that every interval is *restricted* by exactly two points in R, see Figure 5. In case of ambiguity, due to collinearities, we say the point in R closer to s is

5

the restricting point. Those vertices can be either the endpoints of s (a and b) or a different reflex vertex in R. We show the following claim.

Claim 2. *If u is a reflex vertex that restricts the reflex vertex v then it holds that*

$$I(v) \subseteq I(u).$$

Proof. The triangle T formed by $I(v)$ and v is trivially convex and fully contained inside P. The reflex vertex u is on the boundary of the triangle and thus sees every point of T. In particular also $I(v)$. □

Given the previous claim, we construct the visibility restriction graph G as follows. The vertices are formed by the points in R. We say that uv forms a directed edge, if u is restricted by v. We summarize a few useful properties of G in the following claim.

Claim 3. *The visibility restriction graph of a polygon with point visibility width at most k has the following properties.*

1. *The segment endpoints a, b are the only two sinks.*

2. *The out-degree is two for every vertex $v \in R \setminus \{a, b\}$.*

3. *The in-degree is at most $k - 1$ for every vertex $v \in R$.*

4. *The longest path has at most $k + 1$ vertices.*

Proof. By definition, every reflex vertex is restricted by exactly two vertices in R. This implies Item 1 and 2.

Any reflex vertex v can see itself and all its neighbors. Its in-degree neighbors are also reflex vertices. As no point can see more than k reflex vertices v has at most $k - 1$ in-degree neighbors. This concludes the proof of Item 3.

Finally, to prove Item 4, let $p = u_1 u_2 \ldots u_l$ be a directed path. Then it holds that there is a point

$$q \in I(u_l) \subseteq \ldots \subseteq I(u_2) \subseteq I(u_1) = s.$$

The point q sees all reflex vertices of the path p. As no point sees more than k reflex vertices, it holds that p has at most k reflex vertices. As all but potentially the first vertex is a reflex vertex, we have $l \leq k + 1$. □

The properties of the last claim enable us to give an upper bound on the size of G and thus also on the size of R.

Claim 4. *The visibility restriction graph G of a polygon with point visibility width $[\![\mathtt{pvw}]\!] = k$ has at most $k^{O(k)}$ vertices.*

Proof. We organize G into layers depending on the distance from a and b. Note that if layer i has t vertices then layer $(i + 1)$ has at most $t \cdot k$ vertices. As there are at most $k + 1$ layers and the first layer has size two we get that G has at most

$$2 + 2k + 2k^2 + 2k^3 + \ldots + 2k^k = k^{O(k)}$$

vertices. □

3 Conclusion

We believe that local complexity has the potential to be a useful parameter. We gave two ways to define local complexity in a rigorous way and showed how those two ways relate to one another. We want to end with a few open questions.

1. Can we find algorithms and data structures that can make use of low local complexity?

2. Can we compute or approximate the point visibility width and chord visibility width in an efficient manner? Note that this is more a theoretical question. We do not necessarily need to know the chord visibility width of a polygon to use the concept in the design and analysis of an algorithm.

3. Are there other ways to formalize the idea of low local complexity within a polygonal region?

References

[1] Akanksha Agrawal and Meirav Zehavi. Parameterized analysis of art gallery and terrain guarding. In *International Computer Science Symposium in Russia*, volume 12159 of *LNCS*, pages 16–29. Springer, 2020. `doi:10.1007/978-3-030-50026-9_2`.

[2] Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Information Processing Letters*, 100(6):238–245, 2006. `doi:https://doi.org/10.1016/j.ipl.2006.05.014`.

[3] Simon Hengeveld and Tillmann Miltzow. A practical algorithm with performance guarantees for the art~gallery problem. *CoRR*, abs/2007.06920, 2020. `arXiv:2007.06920`.

[4] David G. Kirkpatrick and Raimund Seidel. The ultimate planar convex hull algorithm? *SIAM Journal on Computing*, 15(1):287–299, 1986. `doi:10.1137/0215021`.