Technische Universität München

Department of Mathematics

Bachelor's Thesis

Extension Complexity of Convex n-Gons

Josef Wittmann

Supervisor: Prof. Dr. Stefan Weltge

Advisor: Prof. Dr. Stefan Weltge

Submission Date: March 28, 2021

I assure the single handed composition of this bachelor's thesis only supported by declared resources.

Munich, March 28, 2021

Zusammenfassung

Das Thema dieser Arbeit ist die *Extension Complexity* konvexer Polygone, was auf Deutsch so viel wie *Erweiterungskomplexität* bedeutet. Die Idee dahinter ist, ein Polytop als lineare Abbildung eines höherdimensionalen Polytops darzustellen. Diese *erweiterte Darstellung* nennt man *Extended Formulation* und dessen *Größe* ist die Anzahl der Facetten des höherdimensionalen Polytops.

Die *Erweiterungskomplexität* eines Polytops P ist nun die kleinste Größe aller erweiterten Darstellungen dieses Polytops und man bezeichnet diese mit $\mathrm{xc}(P)$.

Wir behandeln in dieser Arbeit die besten bekannten Schranken dieser Erweiterungskomplexität $\mathrm{xc}(P)$ für konvexe Polygone.

Bisher ist bekannt, dass für ein konvexes Polygon P mit n Ecken, hier als n-Eck bezeichnet,

$$\mathrm{xc}(P) \in \Omega(n^{1/2}) \cap O(n^{2/3})$$

gilt. Liegen alle Punkte von P auf einem gemeinsamen Kreis, ist P also *zyklisch*, lässt sich die Abschätzung asymptotisch präzise zu

$$\mathrm{xc}(P) \in \Theta(n^{1/2})$$

verbessern.

Wir geben in dieser Arbeit einen Überblick über die Beweise, die zu den oben genannten Schranken führen, und konzentrieren uns vor allem auf Zusammenhänge der Aussagen und Grenzen der Vorgehensweisen.

Wir beginnen damit, den Beweis für $\mathrm{xc}(P) \in O(n^{2/3})$ für allgemeine n-Ecke P darzustellen und dessen Grenzen aufzuzeigen.
Das Vorgehen ist dort rein geometrischer Natur und versucht für jedes n-Eck eine Teilfolge u von Ecken zu finden, die groß genug ist und eine möglichst kleine Erweiterungskomplexität besitzt (für $m \in \Omega(n^{2/3})$ Ecken $\mathrm{xc}(u) \in O(m^{1/2})$). Durch diese Aussage kann man dann induktiv zeigen, dass $\mathrm{xc}(P) \in O(n^{2/3})$ gilt.
Ein zentraler Satz in diesem Vorgehen erlaubt uns einfache, dreidimensionale erweiterte Darstellungen eines vereinfachten Polygons zusammenzufügen, um eine erweiterte Darstellung unseres gewünschten Polygons zu erhalten. Wir zeigen, dass dieser Satz für ein n-Eck nur erweiterte Darstellungen der Größe $\Omega(n^{1/2})$ erzeugen kann.

Als nächstes geben wir einen Überblick zum Beweis von $\mathrm{xc}(P) \in O(n^{1/2})$ für zyklische n-Ecke P, wobei diese obere Schranke asymptotisch optimal ist.
Das Vorgehen benutzt hier im Kern den linear-algebraischen Ansatz, der Eigenschaften gewisser Matrizen mit der Erweiterungskomplexität eines Polygons verbindet.

Anschließend vergleichen wir beide Vorgehensweisen und versuchen Ähnlichkeiten und Unterschiede auszuarbeiten, auch wenn die beiden Ansätze grundlegend verschieden sind. Wir wiederholen auch in asymptotischer Betrachtung, wie die beiden Ansätze die Schranken hergeleitet haben.

Im letzten Abschnitt geben wir noch einen weiteren Beweis für $\mathrm{xc}(P) \in \Omega(n^{1/2})$ für zyklische n-Ecke, damit auch für allgemeine n-Ecke. Wir verwenden dafür einen Satz, der uns erlaubt diese untere Schranke für eine Familie von Polygonen abhängig von ihrem gegenseitigen Abstand zu formulieren.

Wir schließen die Arbeit mit einer Übersicht über aktuelle Vermutungen zur Erweiterungskomplexität von Polygonen.

Contents

Abstract

We examine the state of extended formulations for convex n-gons, focusing on bounds for the extension complexity in particular.

First, we analyze the geometric proof of the best known upper bound $\text{xc}(P) \in O(n^{2/3})$ and present the limits of some of its theorems.

Then, we outline the linear-algebraic approach, which results in $\text{xc}(P) \in O(n^{1/2})$ for cyclic polygons.

Finally, we provide another proof for $\text{xc}(P) \in \Omega(n^{1/2})$ for cyclic polygons, and therefore for all polygons.

1 Introduction

1.1 Extension Complexity

Six linear inequalities are required to describe a regular hexagon. However, the same polygon[1] can be described by a polytope[2] in $\mathbb{R}^3$ with five linear inequalities and a linear projection, see Figure 1.

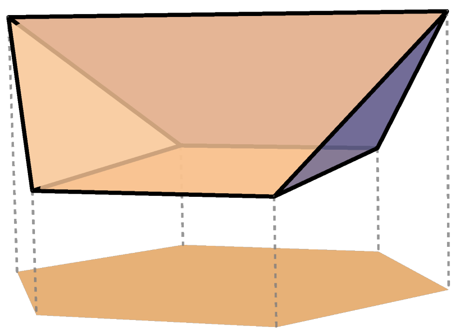

Figure 1. A regular hexagon as a projection of a three-dimensional polytope with five facets. [8, Figure 1]

This is the key idea behind extended formulations. Simplifying the representation of polygons by projecting simpler, higher-dimensional polytopes "down" to the desired polytope. One can think of it as "compressing" the representation. This mental image is inspired by the fact, that regular n-gons (i.e. polygons with n vertices), only require $O(\log(n))$ inequalities in their compressed form [6].

This intuition is formalized as follows:

Definition 1 (Extended Formulation)**.** Let P be a d-dimensional polytope, Q an m-dimensional polytope and $\pi : \mathbb{R}^m \to \mathbb{R}^d$ a linear projection. The pair (Q, π) is called an *extended formulation* of P, if $\pi(Q) = P$. The *size* of (Q, π) is defined as the number of facets of Q.

[1] We define a polygon to be a two-dimensional polytope, i.e. it is convex.
[2] A polytope is the convex hull of a finite set of points.

Definition 2 (Extension Complexity)**.** The *extension complexity* of a polytope P is the smallest possible size of an extended formulation of P. It is usually denoted with $\mathrm{xc}(P)$.

Polytopes with a small extension complexity can be "compressed" very well with extended formulations. Such simplified formulations can be used for building faster algorithms for hard linear programs [17], which seems to be a driving reason for the popularity of research in extended formulations.

1.2 Overview of Current Research

The application of extended formulations for solving hard linear problems seems promising. However, research in the past years shows unpromising results:

- The *traveling salesman polytope* can't have polynomial extension complexity [4]. For this problem, extended formulations don't provide an improvement over traditional methods.
- The results for the *matching polytope* are even worse: In contrast to the polynomial-time algorithm for solving this problem [5], there is no polynomial-size extended formulation [12].

That's why the study of polygons has importance as a benchmark for extended formulations in general (Braun and Pokutta called it "prototypical importance" [2]). Current research focuses on finding good bounds for the extension complexity, rather than developing algorithms for constructing extended formulations.

There are currently two main approaches for finding bounds for the extension complexity: The first one is the natural geometric approach. The second one is a linear-algebraic approach based on a remarkable result of Yannakakis [17], which we recall briefly:

Definition 3 (Slack Matrix)**.** Let P be a polytope. Then a *slack matrix*[3] M of P is a nonnegative matrix, whose rows are indexed by the vertices of P and whose columns are indexed by the linear constraints of some representation of $P = \{x \in \mathbb{R}^d \mid Ax \leq b\}$. The entries of M are defined as $m_{ij} = b_j - a_j v_i$, where a_j is a row of A and v_i is a vertex of P.

Definition 4 (Nonnegative Rank)**.** The *nonnegative rank* of a nonnegative matrix $M \in \mathbb{R}^{n\times m}_{\geq 0}$ is the smallest integer k, for which $M = TU$, where $T \in \mathbb{R}^{n\times k}_{\geq 0}$ and $U \in \mathbb{R}^{k\times m}_{\geq 0}$ are nonnegative matrices. We define $\mathrm{rank}_+(M) := k$.

The nonnegative rank of M can equivalently be defined as the minimum r for which the matrix M can be written as the sum of r nonnegative matrices of (ordinary) rank 1.

Theorem 5 (see [17])**.** *Let P be a polytope and M a slack matrix of P. Then*

$$\mathrm{xc}(P) = \mathrm{rank}_+(M).$$

[3]Note that the slack matrix depends on the representation of P. So there is no unique slack matrix to a given polytope.

Because this document is focused on polygons, we give a concise overview of the history of bounds and conjectures for the extension complexity of those:

Define $\mathcal{P}_n$ as the set of all convex polygons with n vertices. Then

$$\text{pc} : \mathbb{N} \to \mathbb{N},\ n \mapsto \max\{\text{xc}(P) \mid P \in \mathcal{P}_n\}$$

is called *polygon complexity*. It describes the largest extension complexity for a polygon with n vertices.

The lower bound $\text{pc}(n) \in \Omega(n^{1/2})$ is quite easy to prove, which was done many times, for example by Fiorini, Rothvoß, and Tiwary [3] using a counting argument. In this paper we will provide another proof for this lower bound.

The upper bound is more challenging. Only few improvements have been made for a long time and the trivial upper bound $\text{pc}(n) \leq n$ could only be improved by constant factors. One such improvement was done independently by Shitov [13] and Padrol and Pfeifle [10] and proved that $\text{pc}(n) \leq (6n + 6)/7$. The basis for their proofs was $\text{pc}(7) = 6$, i.e. each heptagon has extension complexity at most 6.

Because of missing improvements, $\text{pc}(n) \in \Theta(n)$ was conjectured [3]. This was refuted by Shitov [14] proving $\text{pc}(n) \in o(n)$. He later improved the upper bound to $\text{pc}(n) \in O(n^{2/3})$ [15], which is the best known bound so far.

In summary, this is currently known about the polygon complexity:

$$\text{pc}(n) \in \Omega(n^{1/2}) \cap O(n^{2/3})$$

There is another recent paper by Kwan, Sauermann, and Zhao [8], which focuses on cyclic polygons[4]. The authors proved that $\text{xc}(P) \in O(n^{1/2})$ for all cyclic n-gons. This led them to propose that this upper bound could also hold for general n-gons since "cyclic polygons seem to represent quite a diverse cross-section of the space of all polygons" [8, p. 3].

1.3 Overview of this Document

In this paper we focus on the key ideas behind the discussed papers *Sublinear extensions of polygons*, Shitov [15] and *Extension complexity of low-dimensional polytopes*, Kwan, Sauermann, and Zhao [8]. We will make use of examples and skip most technical detail.

In the first part we examine the results which led to the best known upper bounds. We will start with Shitov [15], which provides the upper bound of $O(n^{2/3})$ for arbitrary n-gons by a purely geometric approach. From there we have a look at cyclic n-gons and the upper bound of $O(n^{1/2})$, which was provided by a semi-geometric/semi-algebraic approach by Kwan, Sauermann, and Zhao [8]. Next, we try to compare these two approaches, as far as they are comparable. First on a high level comparing key ideas and later on a more technical level.

In the second part we give another proof of the lower bound and present the applied theorem [1, Theorem 1]. We also give a quick overview of the key ideas behind it.

[4]A polygon is called cyclic, if all its vertices lie on a circle.

2 Upper Bounds for the Extension Complexity

2.1 Arbitrary Convex n-Gons

In this section we give a detailed overview of *Sublinear extensions of polygons*, Shitov [15], which proves $\mathrm{xc}(P) \in O(n^{2/3})$ for any n-gon P. We focus on the underlying ideas and present only important arguments in more detail. In contrast to the source, we present the results first, because, in doing so, we can highlight the reasoning more clearly.

The gist of the approach is as follows:

1. We cut the polygon into smaller slices, for which we can prove small extension complexity. Joining these slices estimates the maximum extension complexity of the polygon from above.
2. We prove this small extension complexity by inductively finding a large enough subset of vertices in the slice, for which we can build a small extended formulation. Joining these subsets will again estimate the maximum extension complexity of the slice from above.
3. We build the small extended formulation by building specific three-dimensional extended formulations for a surrounding polygon. These are "glued" together resulting in an extended formulation for the original set of vertices.

In Appendix A we added a graphical representation of the dependencies in the proofs of the source. It may be of help in understanding how the statements lead to the main result or how some statements are connected.

2.1.1 Main Result

Theorem 6 ([15, Theorem 5]). *Every convex n-gon P has* $\mathrm{xc}(P) \leq 147n^{2/3}$.

The starting point for proving this theorem is the following lemma, adopted from Weltge [16, Proposition 3.1.1], which states that the extension complexity behaves well for unions of polytopes:

Lemma 7 ([15, Lemma 8]). *Let P and Q be polytopes*[5] *in $\mathbb{R}^d$, each different from a point. Then*

$$\mathrm{xc}(\mathrm{conv}(P \cup Q)) \leq \mathrm{xc}(P) + \mathrm{xc}(Q).$$

This statement enables us to split a polygon into smaller parts when estimating its extension complexity.

We will go on to define those smaller parts, which we can handle well.

Definition 8 (Turning angle). If $P \subset \mathbb{R}^2$ is a polygon, then the *turning angle* at a vertex v is $\pi - \angle v_- v v_+$, where v_- and v_+ are two vertices adjacent to v in P.
In other words, it is the amount the angle at v deviates from a straight line.
The *turning angle* of an edge e of a polygon is the sum of the turning angles at the two endpoints of e.

[5] One is tempted to think of these polytopes as being disjunct. But they can be arranged arbitrarily.

Definition 9 (Correct sequence)**.** A sequence $v = (v_1, \ldots, v_n)$ of distinct points on a plane is called *correct* if these points are the vertices of their convex hull P and the segment between any pair of consecutive points in v is an edge of P.

Unless stated otherwise, we assume that the vertices of a correct sequence are in *clockwise* order.

Definition 10 (Thin sequence)**.** Let $\alpha \in (0, \pi)$ and $n \geq 3$ be an integer. A correct sequence $v = (v_1, \ldots, v_n)$ is called *α-thin*, if the turning angle of the edge $\operatorname{conv}\{v_1, v_n\}$ in the polygon $\operatorname{conv} v$ is greater than $2\pi - \alpha$, that is,

$$\angle v_n v_1 v_2 + \angle v_1 v_n v_{n-1} < \alpha.$$

We say that v is *thin*, if it is α-thin for some $\alpha \in (0, \pi)$.

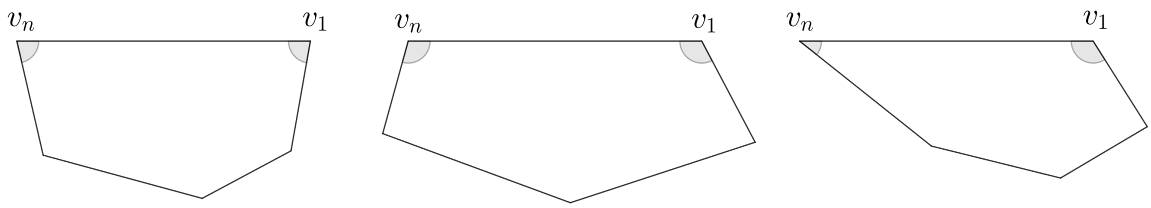

Figure 2. Examples of thin and not-thin sequences: The sequence in the middle is the only one, which is *not* thin.

For thin sequences, all lines $v_i \wedge v_j$ with $\{i, j\} \neq \{1, n\}$ meet on the same side of $v_1 \wedge v_n$, where $\cdot \wedge \cdot$ is used to describe the line joining two points or the intersection of two lines.

Observation 11 (Splitting into thin sequences, [15, Observation 30])**.** *Let P be a polygon with n vertices, and let $q \geq 3$ be an integer. Then the vertices of P can be partitioned into at most q sets each of which is either a point or a pair of points, or a set that forms a $2\pi/q$-thin sequence.*

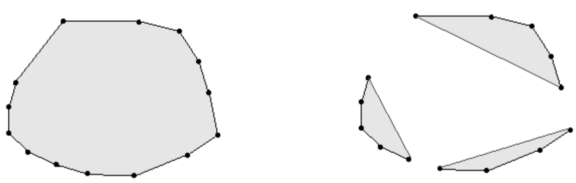

Figure 3. Splitting a polygon into three thin sequences. [15, Figure 4]

Proof outline. For each set, take as many consecutive vertices until the $2\pi/q$-thinness would be violated. Then begin a new set, which starts with at least $2\pi/q$ higher turning angle measured from the start. Therefore, at most q sets result in the full turning angle of 2π covering the polygon. □

Based on Observation 11 we choose to split the original n-gon into 12 distinct $\pi/6$-thin sequences. Joining those (constantly many) sequences will not change the asymptotical bound for the extension complexity of the original polygon in light of Lemma 7.

Now we will go on to show that each $\pi/6$-thin sequence has extension complexity $O(n^{2/3})$. We do this by inductively extracting a subsequence with a small extended formulation.

Theorem 12 ([15, Theorem 58]). *Let v be a $\pi/6$-thin sequence with $n = 1024\tau^3 + 8\tau$ vertices, where $\tau \in \mathbb{N}$. Then v contains a subsequence u with $|u| \geq 4\tau^2$ and $\mathrm{xc}(u) \leq 12\tau$.*

Proving this theorem is the objective of the rest of this section. We preempted it to provide a better understanding of the results.

The following corollary is a direct result of it:

Corollary 13 ([15, Corollary 59]). *Let v be a $\pi/6$-thin sequence with $n > 263\,000$ vertices. Then v contains a subsequence u with $|u| \geq \frac{1}{36} n^{2/3}$ and $\mathrm{xc}(s) \leq \left(\frac{72}{43} n\right)^{1/3}$.*

Proof outline. Apply Theorem 12 with $\tau = \left\lfloor (n/1032)^{1/3} \right\rfloor$, which gives the desired result for $n > 263\,000$. □

Corollary 14 ([15, Corollary 60]). *Let v be a $\pi/6$-thin sequence. Then*

$$\mathrm{xc}(v) \leq \frac{324}{\sqrt[3]{129}} n^{2/3}.$$

Proof outline. Use induction for $n > 263\,000$ ($n \leq 263\,000$ is trivial) and apply Corollary 13 and Lemma 7 for the induction step:

$$\mathrm{xc}(\mathrm{conv}\, v) \leq \frac{324}{\sqrt[3]{129}} \left(n - \frac{n^{2/3}}{36}\right)^{2/3} + \left(\frac{72n}{43}\right)^{1/3} < \frac{324}{\sqrt[3]{129}} n^{2/3}$$

□

We can now prove the main result.

Proof outline of Theorem 6. Using Observation 11 we split the n-gon P into twelve disjoint $\pi/6$-thin sequences[6] with sizes $n_1, \ldots, n_{12}$.

We apply Lemma 7 and Corollary 14 and get

$$\mathrm{xc}(P) \leq \frac{324}{\sqrt[3]{129}} \left(\sum_{i=1}^{12} n_i^{2/3}\right) \leq \frac{324}{\sqrt[3]{129}} \left(12 \left(\frac{n}{12}\right)^{2/3}\right) < 147\, n^{2/3}.$$

□

2.1.2 Building Small Extended Formulations

In this subsection, we build extended formulations of small size for polygons with special properties. The key idea is to omit some vertices to obtain a simpler surrounding polygon. For this polygon we build specific three-dimensional extended formulations with respect to the omitted vertices. If we now combine these extended formulations, we get a higher-dimensional extended formulation for the original polygon with small extension complexity.

Before we can formulate this theorem, we have to introduce *acute polyhedra* and *acute diagrams*, which help us build three-dimensional extensions for the surrounding polygon.

[6] Technically, we have to consider the case, when a sequence has less than three points. We omit it, since it does not provide further insight.

Definition 15 (Acute Polyhedron)**.** Assume $P \subset \mathbb{R}^3$ is a polyhedron and B is one of its facets. If all other facets $F \neq B$ of P share an edge with B and the angle[7] between B and F is acute, then P is called an *acute polyhedron* with *base* B.[8]

Definition 16 (Main edge)**.** An edge e of an acute polyhedron with base B is called *main*, if exactly one endpoint of e lies on B.

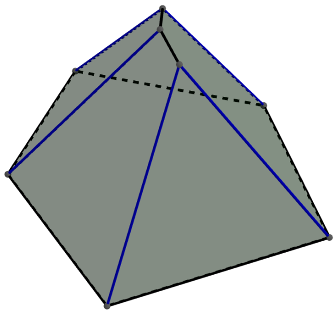

Figure 4. An acute polyhedron with main edges colored.

Lemma 17. *Any base vertex of an acute polyhedron belongs to a unique main edge.*

Proof outline. A base vertex v with two main edges would be contained in three non-base faces. Since each face has to include a base edge by Definition 15, two of them coincide on the base, which is impossible for acute polyhedra. □

Now follows an important lemma, which allows us to construct an acute polyhedron for a given base, where all main edges but one have fixed direction with respect to the base.

Lemma 18 ([15, Lemma 15])**.** *Let V be a polygon with vertices $v_1, \dots, v_n$ and let $y_1, \dots, y_{n-1}$ be a set of inner points of V. Then there is an acute polyhedron P with base V such that, for any $i \in \{1, \dots, n-1\}$, the image of the main edge passing from v_i under the orthogonal projection of P onto V is collinear to $v_i \wedge y_i$.*

Proof outline (see Figure 5). Define the planes H_i orthogonal to V containing v_i and y_i. The half-planes F_i are defined recursively with $v_i \wedge v_{i-1}$ as their base[9] and containing $F_{i+1} \cap H_i$. Begin constructing F_n with base $v_n \wedge v_{n-1}$ having an arbitrary acute angle with V. Iteratively construct $F_{n-1}, \dots, F_1$ (all have an acute angle with V, since they contain the ray $F_{i+1} \cap H_i$, which heads towards the interior of V). Now V and $F_1, \dots, F_n$ define the acute polyhedron. □

Observation 19 ([15, Observation 16])**.** *Let P be an acute polyhedron with base B. The orthogonal projection π of P onto the plane containing B maps the non-base points of an acute polyhedron injectively into the interior of B.*

Proof outline. Follows by definition of acute polyhedra (acute angle) and convexity of P. □

[7]The angle between two faces A and B with a common edge e is defined as the angle between two oriented segments that lie on A and B respectively, have their origin on e and are orthogonal to e.

[8]Note that all facets lie on the same side of B. Since otherwise, B would be no facet of P.

[9]We define $v_0 := v_n$, since v_1 and v_n share an edge.

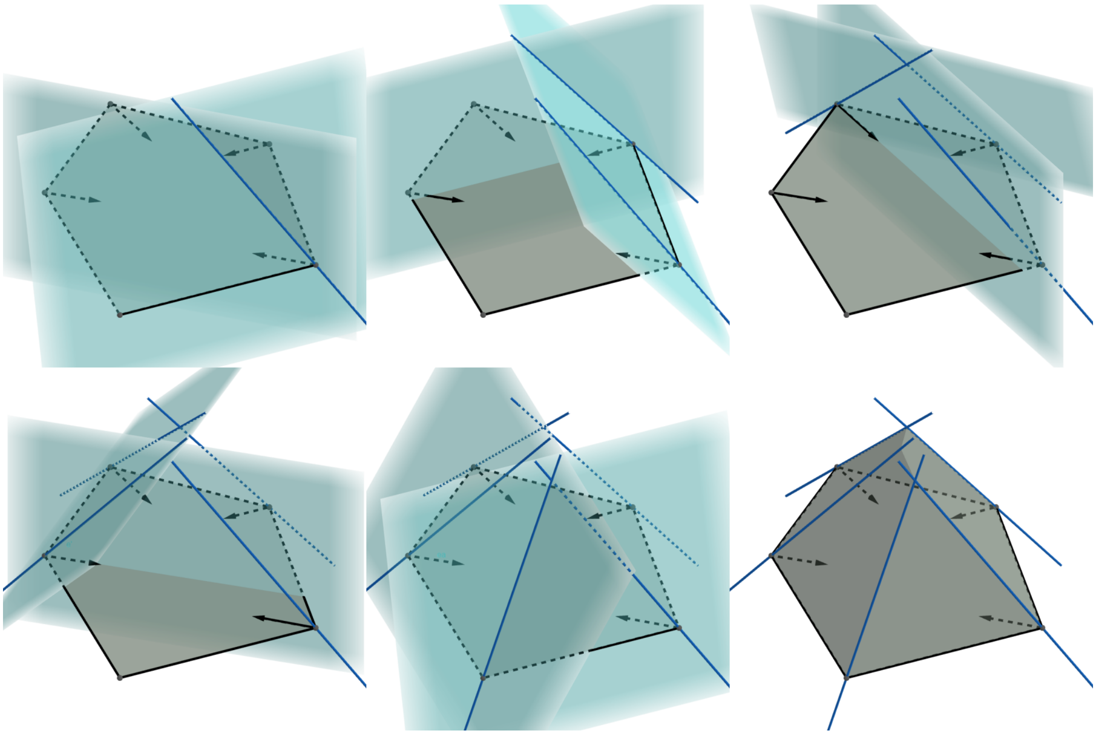

Figure 5. The construction of Figure 4 in the proof of Lemma 18.

Definition 20 (Acute diagram and lifting)**.** Let P be an acute polyhedron with base B and π like in Observation 19. The image of all edges of P under π gives us a diagram which we call the *acute diagram* of P relative to base B. We also say that P is an *acute lifting* of the corresponding diagram.

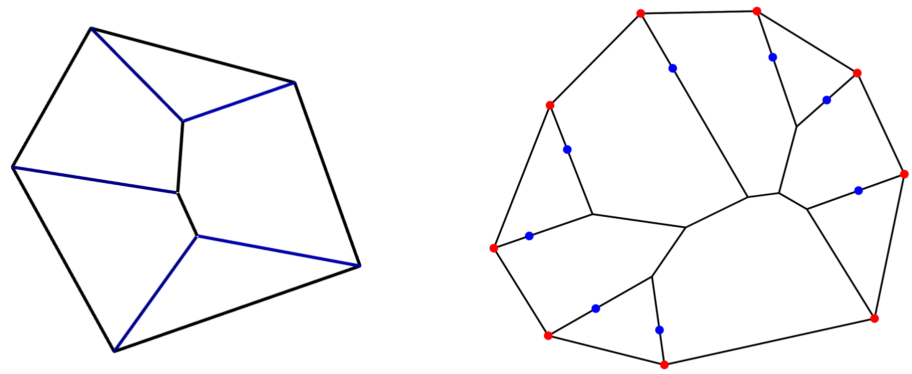

Figure 6. The acute diagram of Figure 4 with colored main edges and an example of an acute diagram with more vertices, where blue points fix the main arc directions as shown in Lemma 18.

We can now reformulate Lemma 18 for acute diagrams:

Corollary 21 ([15, Corollary 26])**.** *Let V be a polygon with vertices $v_1, \ldots, v_n$ and let $y_1, \ldots, y_{n-1}$ be a set of inner points of V. Then there is an acute diagram with base V such that, for any $i \in \{1, \ldots, n-1\}$, the main edge from v_i lies on $v_i \wedge y_i$.*

We showed that every acute polyhedron has an acute diagram. Now we want to describe the properties of acute diagrams, which allow us to "lift" them into acute polyhedra.

Lemma 22 (Properties of acute diagrams, [15, Lemma 19])**.** *Let Δ be the acute diagram of an acute polyhedron P with base B. Then Δ is a planar straight-line graph such that*

(o) the base of Δ is B,

(i) every node of Δ has degree at least three,

(ii) the non-base nodes of Δ lie in the interior of the base,

(iii) every edge of the base is an arc of Δ,

(iv) every bounded face F of Δ contains exactly one arc e_F of the base,

(v) if a non-base arc e of Δ separates faces F, G, then e, e_F, e_G are concurrent.

We omit the proof, because we only lift polyhedra from diagrams in the following.

Lemma 23 (Lifting acute diagrams, [15, Lemma 25])**.** *A planar straight-line graph Δ satisfying (i)-(v) as in Lemma 22 is an acute diagram of some acute polyhedron P.*

Proof outline. We can show that each diagram Δ has a triangle formed by two main arcs and one base arc e with turning angle less than π (expect in the cases, where Δ is a triangle or trapezoid).[10]

Then we use induction on the number of nodes of the base of Δ. The induction basis are the two exceptions from above, where one can easily construct an acute polyhedron.

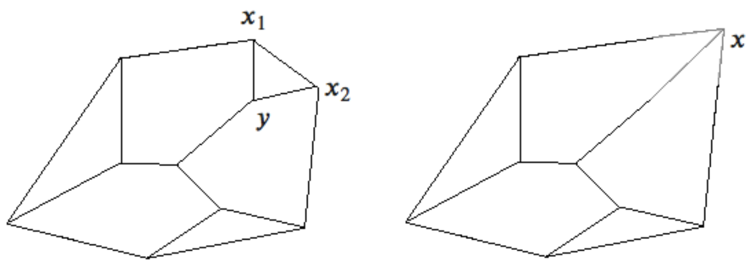

Figure 7. The inductive step in the proof of Lemma 23. [15, Figure 2]

For the induction step we find a triangle with base arc e like above and "remove" it by continuing the arcs next to it (see Figure 7). The continuation meets on the "outside" of e, since e has turning angle less than π.

From the induction hypothesis we can lift this diagram with one vertex less. We cut the resulting polyhedron by the plane defined by our triangle to get our desired polyhedron. □

We can now focus on the main result of the first part, which allows us to build small extended formulations by combining acute polyhedra for a surrounding polygon, where some vertices were omitted. These polyhedra are lifted from acute diagrams, which have to include the omitted vertices in their main edges.

[10] In the original paper it is shown by defining a flow on the inner arcs of Δ towards a base vertex s. For this flow we can find two consecutive "furthest" vertices, which therefore form a triangle. If the turning angle was too large, one can choose another s.

Theorem 24 (Gluing acute extensions together, [15, Theorem 28]). *Let P be a polygon with vertex set V. Let $\emptyset \neq S \subseteq V$ and $\delta \geq 1$ be an integer. Assume*

(i) for any $s \in S$ there are two vertices s', s'' on two edges of P adjacent to s,

(ii) there are δ points $\{s^1, \ldots, s^\delta\}$ in the interior of the triangle $T_s = \operatorname{conv}\{s, s', s''\}$,

(iii) the triangles T_s are disjoint for different s and

(iv) for any $i \in \{1, \ldots, \delta\}$ there is an acute diagram D^i with base P where, for any $s \in S$, the segment between s and s^i is a subset *of the main edge passing from s.*

$$\Rightarrow \operatorname{xc}\left(\operatorname{conv}\left((V \setminus S) \cup \bigcup_{s \in S} \{s', s'', s^1, \ldots, s^\delta\}\right)\right) \leq |V| + |S| + \delta$$

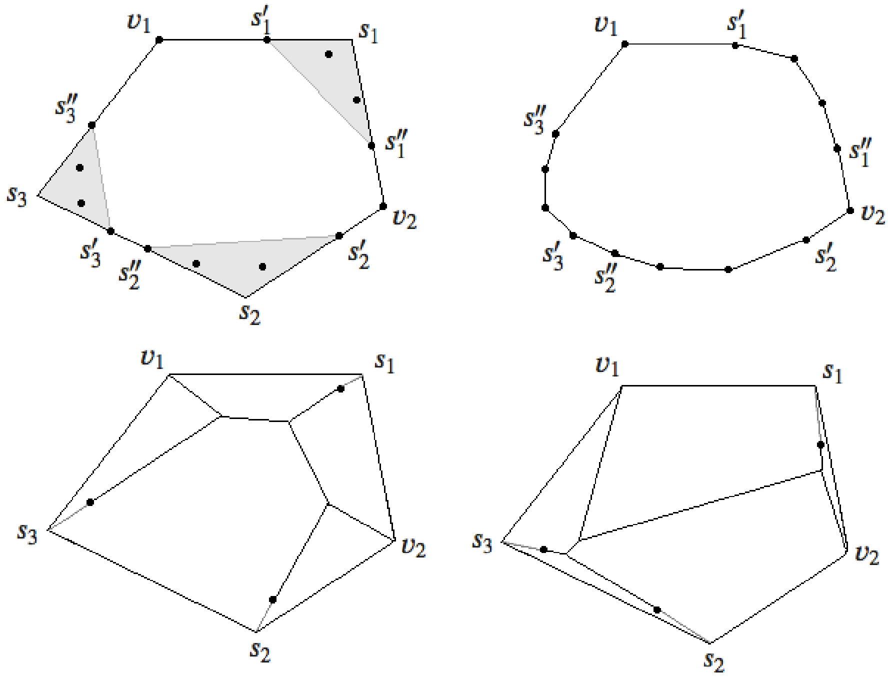

Figure 8. An application of Theorem 24: Two acute diagrams confirm that the 14-gon on the right has xc at most 10. [15, Figure 3]

Proof. We are going to construct a polytope $\mathcal{P}'$ in $\mathbb{R}^2 \times \mathbb{R}^\delta = \{(x, y, z^1, \ldots, z^\delta)\}$ with at most $|V| + |S| + \delta$ facets, which is an extended formulation of P.

Assume that $a_e x + b_e y + c_e \geq 0$ is the defining inequality of an edge e of P. Then the acute diagrams D^i are lifted into three-dimensional polytopes described by $z \geq 0$ and $a_e x + b_e y + c_e \geq \varepsilon_e^i z$ with $\varepsilon_e^i > 0$ for all edges e of P (we assume that all polygons are located in the upper half-space $z \geq 0$).

Then we define an auxiliary polytope $\mathcal{P} \in \mathbb{R}^2 \times \mathbb{R}^\delta$ with the following $|V| + \delta$ inequalities

$$a_e x + b_e y + c_e \geq \sum_{i=1}^{\delta} \varepsilon_e^i z,$$
$$z^i \geq 0,$$

where e runs over all edges of P and $i \in \{1, \ldots, \delta\}$.

We have $\dim \mathcal{P} = \delta + 2$, i.e. it has full dimension, since we can find a point (x, y) in the interior of P and $\varepsilon > 0$, such that $(x, y, \varepsilon, \ldots, \varepsilon)$ fulfills all inequalities strictly.

And from Observation 19 we know $\pi(\mathcal{P}) = P$, where $\pi(x, y, z^1, \ldots, z^\delta) = (x, y)$ is an orthogonal projection.

For any $s = (x_s, y_s) \in S$, $\sigma_s := (x_s, y_s, 0, \ldots, 0)$ is a vertex of $\mathcal{P}$, since it fulfills exactly $\delta + 2$ inequalities with equality, namely $z^i \geq 0$ and the two inequalities corresponding to the edges at s.

Therefore, there are $\delta + 2$ rays passing from σ_s: Two corresponding to the edges at s and δ of the form

$$(x_s + \alpha_s^i t, y_s + \beta_s^i t, 0, \ldots, 0, t, 0, \ldots, 0), t \geq 0, \tag{1}$$

where there are $i - 1$ zeros before t and (α_s^i, β_s^i) is pointing from s to s^i (and scaled such that we don't need a factor for t).

Now we define a half-space H_s, whose defining hyperplane is determined by the points $(x_{s'}, y_{s'}, 0, \ldots, 0)$, $(x_{s''}, y_{s''}, 0, \ldots, 0)$ and the δ points on (1), which project to s^i under π (this hyperplane is well-defined, as all those points are linearly independent). The orientation of H_s is set such that it does not include σ_s.

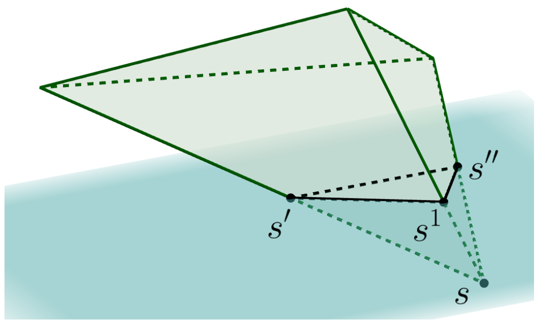

Figure 9. Cutting P with H_s in Theorem 24 for $\delta = 1$.

If we now intersect $\mathcal{P}$ with H_s, we cut the rays passing from σ_s at the defining points of H_s. Since these points are all contained in an edge of $\mathcal{P}$ (by definition of D^i), these $\delta + 2$ points become vertices by this cut. The different cuts themselves don't overlap on $\mathcal{P}$, since the triangles T_s are disjoint by definition and all s^i lie in the interior of T_s.

So we can define $\mathcal{P}' = \mathcal{P} \cap \bigcap_{s \in S} H_s$, which is given by $|V| + |S| + \delta$ inequalities, and conclude that $\pi(\mathcal{P}')$ is the desired polygon. $\square$

Remark 25. *There is a subtle difference between Theorem 24 and Corollary 21.*

In the theorem we assume that "the segment between s and s^i is a subset *of the main edge", but in the corollary we construct a diagram for which "the main edge from v_i lies on $v_i \wedge y_i$", implying that y_i can lie outside the main edge.*

To understand Theorem 24 better, we examine the limits it has in application. This observation is aside from the main argument.

Observation 26 (Limits of Theorem 24)**.** *Let P be a polygon with n vertices. By applying Theorem 24 we obtain an extended formulation for P with size in $\Omega(n^{1/2})$.*

Proof. We set $P := \operatorname{conv}\left((V \setminus S) \cup \bigcup_{s \in S} \{s', s'', s^1, \ldots, s^\delta\}\right)$ as in Theorem 24. We define $v := |V|$ and $s := |S|$. Throughout this proof we assume that we can apply Theorem 24 for our choice of v, s and δ. Then

$$n \leq (v - s) + s(2 + \delta) = v + s(1 + \delta). \tag{2}$$

With (2) we can express δ dependent on n, v and s:

$$\delta \geq \frac{n - v}{s} - 1$$

We now define $\delta := (n - v)/s$ as the minimal possible value[11] for given v and s, since the size of the extended formulation increases with δ. This way we can define

$$f_n(v, s) := v + s + \delta = v + s + \frac{n - v}{s}$$

as the size of the extended formulation for given v and s.

$$\frac{\partial}{\partial v} f_n(v, s) = 1 - \frac{1}{s} > 0$$

shows that v should be chosen minimal for minimal extension size. So we set $v = s$ according to the assumptions of Theorem 24.

$$\begin{aligned} f_n(s, s) &= s + s + \frac{n - s}{s} \\ &= 2s + \frac{n}{s} - 1 \\ \frac{\partial}{\partial s} f_n(s, s) &= 2 - \frac{n}{s^2} \Rightarrow s_{min} := \sqrt{\frac{n}{2}} \\ \frac{\partial^2}{\partial^2 s} f_n(s, s) &= \frac{2n}{s^3} > 0 \end{aligned}$$

This shows that f_n has a minimum at $s_{min} = \sqrt{n/2}$ with $v_{min} = s_{min}$.

Finally, we can find the minimal value for f_n:

$$f_n(s_{min}, s_{min}) = 2\sqrt{2n} - 1$$

□

2.1.3 Finding Good Subsequences

After defining this central theorem, we have to make a way to apply it to general polygons.[12] Hence, we formalize the notion of "good" sequences in terms of Theorem 24.

Definition 27 (t-scattered)**.** Let $t, n \in \mathbb{N}$ and $G \subseteq \{1, \ldots, n\}$. G is called *t-scattered* if for $g, g' \in G, g \neq g'$ we have $|g - g'| \geq t$.

Definition 28 (G-envelope)**.** Assume $n \geq 5$ is an integer and $G \subseteq \{3, 4, \ldots, n - 2\}$ is a 3-scattered subset. If $v = (v_1, \ldots, v_n)$ is a thin sequence, then the *G-envelope* of v is the sequence v_G obtained from v by replacing the points v_{g-1}, v_g, v_{g+1} with $o_g := (v_{g-2} \wedge v_{g-1}) \wedge (v_{g+1} \wedge v_{g+2})$ for all $g \in G$.

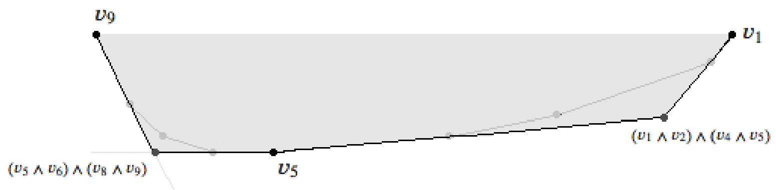

Figure 10. The G-envelope of the sequence in Figure 11 with $G = \{3, 7\}$. [15, Figure 6]

Definition 29 (G-good)**.** Assume $n \geq 5$ is an integer and $G \subseteq \{3, 4, \ldots, n-2\}$ is a 3-scattered subset. A thin sequence $v = (v_1, \ldots, v_n)$ is called *G-good* if there is an acute diagram with the base $\operatorname{conv} v_G$ such that, for any $g \in G$, the main edge passing from o_g *contains* v_g.

Now we find another way to determine which sequences are actually *G-good*. Therefore, we have to define the following:

Definition 30. Let $v = (v_1, \ldots, v_n)$ be a thin sequence. For indexes $i, \hat{\imath}, k, \hat{\jmath}, j$ satisfying $1 \leq i < \hat{\imath} < k < \hat{\jmath} < j \leq n$, we define $\rho(v, i, \hat{\imath}, k, \hat{\jmath}, j)$ as the ray passing from the point $(v_i \wedge v_{\hat{\imath}}) \wedge (v_{\hat{\jmath}} \wedge v_j)$ towards v_k.

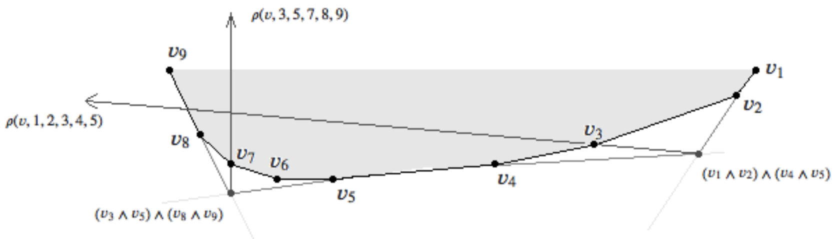

Figure 11. Rays $\rho(v, 1, 2, 3, 4, 5)$ and $\rho(v, 3, 5, 7, 8, 9)$. [15, Figure 5]

Lemma 31 ([15, Lemma 42])**.** *Assume $n \geq 5$ is an integer and $G \subseteq \{3, 4, \ldots, n-2\}$ is a 3-scattered subset. Assume $v = (v_1, \ldots, v_n)$ is a $\pi/2$-thin*[13] *sequence such that, for all $g \in G$, the ray $\rho(v, g-2, g-1, g, g+1, g+2)$ leaves* $\operatorname{conv} v$ *through the relative interior of the edge* $\operatorname{conv}\{v_1, v_n\}$*. Then v is G-good.*

Proof. To prove that v is G-good, we build an acute diagram Δ whose base is the G-envelope v_G of v and show that the main edge passing from

$$o_g := (v_{g-2} \wedge v_{g-1}) \wedge (v_{g+1} \wedge v_{g+2})$$

contains v_g for each $g \in G$.

[11]For simplicity, we skip the case, where $(n-v)/s - 1$ is integer. The difference would only be -1 in the final result.

[12]See Remark 25, why we can't simply use Corollary 21.

[13]In the source v is only required to be thin, but building orthogonal main edges in the proof is only guaranteed to work, if both angles at v_1 and v_n are smaller than $\pi/2$. This does not change any further proof, because this assumption is given, when this lemma gets applied.

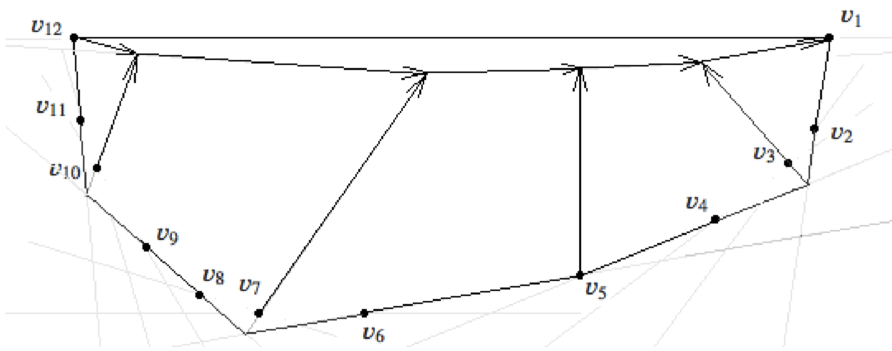

Figure 12. An instance of Lemma 31 with $n = 12$ and $G = \{3, 7, 10\}$. [15, Figure 7]

First we build an acute diagram Δ using Corollary 21 with base v_G and
(1) for each $g \in G$ the main edge from o_g goes towards v_g,
(2) for each $k \notin \{1, n, g-1, g, g+1\}$ the main edge from v_k is orthogonal to $v_1 \wedge v_n$,
(3) the main edge from v_n passes towards a point u_0 that lies sufficiently close to the middle of $\operatorname{conv}\{v_1, v_n\}$ in the interior of $\operatorname{conv} v$.

Now it remains to show that the edge passing from o_g actually contains v_g. We will prove this by contradiction: Assume there is some $j \in G$, for which v_j is not contained in the edge passing from o_j. That is, there is an edge in Δ, which has its endpoint e in the interior of $\operatorname{conv}\{o_j, v_j\}$, thus outside of $\operatorname{conv} v$.

Case 1. There is a path of inner arcs in Δ from v_n to v_1, including e.
This is impossible, because the main edge passing from v_n leaves towards u_0, thus not leaving $\operatorname{conv} v$. This follows from property (v) in Lemma 22, because an inner arc of Δ on the path from v_n to v_1 separating the face containing $v_1 \wedge v_n$ and the face containing $v_i \wedge v_{i+1}$ continues to arc "towards" v_1 for decreasing i, because v itself is convex.[14]

Case 2. There is *no* path of inner arcs in Δ from v_n to v_1, including e.
All main edges of Δ go towards the interior of $\operatorname{conv}\{v_1, v_n\}$ (except the two at v_1 and v_n). And any path from a base vertex of Δ to v_1 that includes e has to leave $\operatorname{conv} v$. This path can only leave through the interior of $\operatorname{conv}\{v_1, v_n\}$ because of (v) in Lemma 22 and the convexity of v.[15] So it has to intersect the path from v_n to v_1 before leaving $\operatorname{conv} v$. This contradicts the assumption of this case.

This shows that the edge passing from o_g contains v_g for all $g \in G$. □

We now need to define one more property of a sequence, by which we split upcoming considerations.

Definition 32 (Slanted sequence)**.** Let $v = (v_1, \dots, v_t)$ be a thin clockwise sequence, $\beta \in (0, \pi/2)$, $\delta \geq 0$. We say that v is *clockwise-slanted* to an angle β with tolerance δ if,

[14] The line through this arc has to be concurrent with $(v_1 \wedge v_n) \wedge (v_i \wedge v_{i+1})$, which is a point on $v_1 \wedge v_n$ outside $\operatorname{conv}\{v_1, v_n\}$, because of the convexity of v. This arc can't be moving away from v_n steeper than $v_n \wedge u_0$.

[15] For any intersection point, there is a leftmost incoming arc f and a rightmost incoming arc g, which are concurrent with the base edges f_l, f_r or g_l, g_r respectively. So the outgoing arc h is concurrent with f_l and g_r, which sets it between the continuations of f and g.

for all $i, \hat{\imath}, \hat{\jmath}, j$ satisfying $1 \leq i < \hat{\imath} < \hat{\jmath} < j \leq t$, the ray $\rho(v, i, \hat{\imath}, k, \hat{\jmath}, j)$ satisfies

$$\angle\left(\overrightarrow{v_{\hat{\jmath}} \wedge v_j},\ \rho(v, i, \hat{\imath}, k, \hat{\jmath}, j)\right) < \beta$$

for all k satisfying $\hat{\imath} < k < \hat{\jmath}$ except for at most δ such values of k.

If a counterclockwise sequence satisfies above assumptions, we call the reversed (clockwise) sequence *counterclockwise-slanted.*

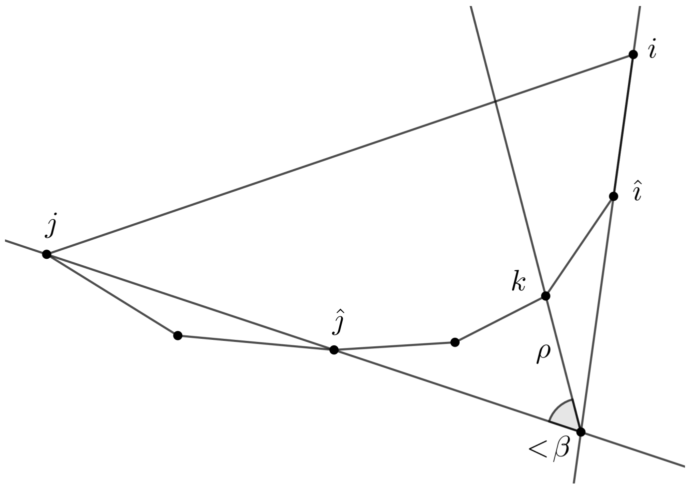

Figure 13. Points and angles used in the definition of slanted sequences.

Note that we couldn't come up with an intuitive understanding of slanted sequences, but there are some remarks: If k is moved in clockwise direction, the measured angle decreases. Therefore, vertices failing the condition are consecutive starting after $\hat{\imath}$. And, if one fixes k, $\hat{\jmath}$ and j, moving i or $\hat{\imath}$ in clockwise direction will increase the measured angle. However, the same can't be done for j and $\hat{\jmath}$.

We split the search for a good subsequence into two cases:

1. The original sequence has a large enough slanted subsequence.
2. The original sequence does not have such a subsequence.

We can show that, for each case, there is a large enough subsequence, for which Theorem 24 can be applied well. This is done in the following lemmata, for which we will skip the technical proofs, since they provide little insight.

Lemma 33 ([15, Lemma 49])**.** *Let δ, τ, m be positive integers and $n = 8\tau m$. Let α, β be positive reals with $\pi/2 > \beta \geq 2\alpha$. Assume $v = (v_1, \ldots, v_n)$ is an α-thin sequence such that any subsequence with m points is neither clockwise-slanted nor counterclockwise-slanted relative to the angle β and tolerance 2δ. Then v has a subsequence with $(6 + \delta)\tau$ points with extension complexity not exceeding $6\tau + \delta + 1$.*

Lemma 34 ([15, Theorem 57])**.** *Let $\delta > 1$ be an integer and $n = 8\delta^2$. Let $v = (v_0, \ldots, v_n)$ be a $\pi/6$-thin sequence that is clockwise-slanted to the angle $\pi/3$ with tolerance 2δ. Then v has a subsequence of size at least $0.25\delta^2$ and extension complexity at most 3δ.*

Now we have all premises for proving Theorem 12, which states that a $\pi/6$-thin sequence v with $n = 1024\tau^3 + 8\tau$ vertices contains a subsequence u with $|u| \geq 4\tau^2$ and $\mathrm{xc}(u) \leq 12\tau$. In light of 26 this subsequence makes optimal use of Theorem 24 in an asymptotical way.

Proof outline of Theorem 12. We try to apply Lemma 33 with $\beta = \pi/3$, $\delta = 4\tau$ and $m = 8\delta^2 + 1$.

If it is applicable, v contains a subsequence with the desired properties.

If it isn't applicable, v has a subsequence u, which is slanted to the angle $\pi/3$ with tolerance 2δ. By Lemma 34, u contains a subsequence with the desired properties. □

2.1.4 Conclusion

We proved that $\mathrm{xc}(P) \in O(n^{2/3})$ for any n-gon P, which leaves us wondering if this bound can be improved further, since the lower bound for some n-gon P is $\mathrm{xc}(P) \in \Omega(n^{1/2})$.

Theorem 12 shows that we can find a subsequence with $m \in \Omega(n^{2/3})$ vertices with extension complexity in $O(m^{1/2})$. Referring to Observation 26, Theorem 24 is applied asymptotically optimal. If we wanted to improve the upper bound for $\mathrm{xc}(P)$ with this approach, we would have to find a subsequence with more vertices, for which we could apply Theorem 24 optimally. If this subsequence had $\Theta(n)$ vertices, the ensuing result of $\mathrm{xc}(P) \in O(n^{1/2})$ for any n-gon P would close the gap between lower and upper bound.

2.2 Cyclic n-Gons

In this section we give an overview of *Extension complexity of low-dimensional polytopes*, Kwan, Sauermann, and Zhao [8], which proves $\mathrm{xc}(P) \in O(n^{1/2})$ for any n-gon P with vertices on a circle.

The type of approach in this paper differs form Shitov [15], since it applies the linear-algebraic method using slack matrices.

One important insight is the "lampshade argument", which the authors called this way because of the geometric interpretation in three dimensions (see Figure 14). In the two-dimensional case, the simplified form can be interpreted like this:

Given a polygon P, consider a set of consecutive facets F' and the set of vertices V', which are not endpoints of F'. Then if we consider the $V' \times F'$ submatrix $M[V', F']$ of a slack matrix M, we have $\mathrm{rank}_+ M[V', F'] = O(1)$.
The reason for this is that we can enclose V' inside a polygon Q, such that all vertices of Q are on the "positive slack" side of F' (meaning Q and P lie on the same side of every $f \in F'$). This polygon Q can be thought of as a "polyhedral lampshade" and one can always build such a polygon Q with constantly many vertices.
Since Q encloses V', every vertex $v \in V'$ is a convex combination of the vertices in Q.
For every vertex q of Q consider the vector $u_q \in \mathbb{R}^{|F'|}$ of slacks for all $f \in F'$, which is positive by construction of Q.
And since the slack function is affine-linear, one can convex combine the slack vectors for all $v \in V'$ (regarding F') from those u_q.
In other words, every row of $M[V', F']$ is a convex combination of these constantly many vectors u_q, which shows $\mathrm{rank}_+ M[V', F'] = O(1)$.

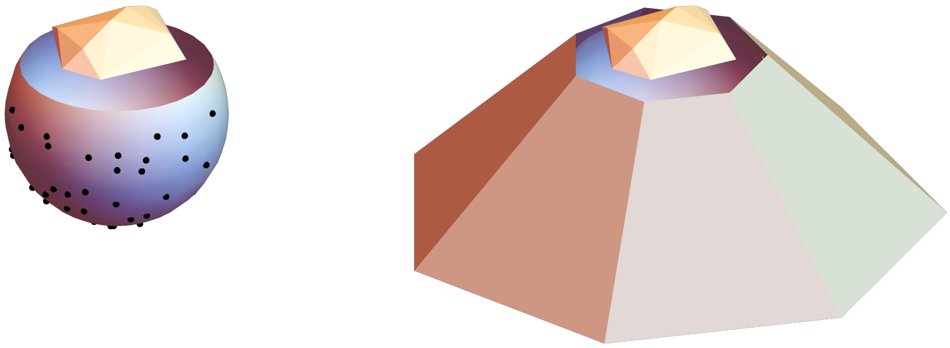

Figure 14. On the left, a small patch of facets near the north pole is far away from a collection of vertices. On the right, a "polyhedral lampshade" encloses all the vertices in our collection, and lies entirely on the "positive slack" side of each of the facets in our patch. [8, Figure 2]

Kwan, Sauermann, and Zhao developed this argument even for higher dimensions. Since we focus on polygons in this paper, we will only cover the two-dimensional theorem.

Definition 35. A polygon is called *cyclic* if all its vertices lie on a common circle.

Theorem 36 ([8, Theorem 1.3])**.** *Let P be a cyclic polygon with $n \geq 18^2$ vertices. Then* $\mathrm{xc}(P) \leq 24\, n^{1/2}$.

The approach follows these steps:

1. Split the circle into arcs, containing approximately $n^{1/2}$ facets each.
2. Color these arcs with constantly many colors, such that two arcs of the same color are "well-separated".
3. Build a nonnegative matrix deduced from the slack matrix by rescaling rows and adding approximately $n^{1/2}$ vectors, such that the entries for vertices and facets of the same arc are zero.
4. Apply the "lampshade argument" (in a more general form) for the rest of this matrix and obtain the desired bound.

We will now go through these steps in more detail.

Let P be a cyclic polygon, V its set of vertices, F its set of facets and M a fixed slack matrix of P.

We divide the circle into a set $\mathcal{X}$ of $\lceil n^{1/2} \rceil$ arcs, such that each arc spans from one vertex to another vertex (including both) and contains at most $\lceil n^{1/2} \rceil$ facets, i.e. facets with both endpoints in the arc, and therefore at most $N := \lceil n^{1/2} \rceil + 1$ vertices.

Definition 37 (Well-separated)**.** We say that two arcs $X, X' \in \mathcal{X}$ with arc lengths ε, ε' are *well-separated*, if the arc-distance between any two points $x \in X$ and $x' \in X'$ is at least $5 \min\{\varepsilon, \varepsilon'\}$.

Lemma 38. *For every arc $X \in \mathcal{X}$, there are at most 13 arcs $X' \in \mathcal{X}$, which are not well-separated from X and at least as long as X.*

Proof. We count how many such arcs fit into the "not well-separated space".
Let ε be the length of X. If we put 5 arcs of length at least ε to each side of X, each other arc would have arc-distance at least 5ε to X. So the maximal number of not well-separated arcs with length at least ε distinct to X is 10.[16] □

Lemma 39. *We can color the arcs in $\mathcal{X}$ with at most 14 colors, such that arcs of the same color are well-separated.*

Proof. Order the arcs by decreasing length and color them in that order. For an arc X, there are, by Lemma 38, at most 13 arcs of at least the same length not well-separated. So just pick a color (from a set of 14 colors), which wasn't assigned to those arcs and proceed until all arcs are colored. □

We color the arcs in $\mathcal{X}$ with colors labeled by $c = 1, \dots, 14$ and define $\mathcal{X}_c$ as the set of arcs with color c.

We also split the columns of the slack matrix M by color, i.e. into slices $M[V, F_c]$, where F_c are all facets in arcs with color c. This allows us to handle each color separately, since we can estimate $\operatorname{rank}_+ M \le \sum_{c=1}^{14} \operatorname{rank}_+ M[V, F_c]$.

For an arc $X \in \mathcal{X}$ define F^X or V^X as the facets or vertices inside that arc. We also call a vertex v *local* to a facet f, if $x \in V^X$ and $f \in F^X$, i.e. if they are in the same arc.

Further, we enumerate all vertices for each arc of color c, that is for each $X \in \mathcal{X}_c$ we define a bijection $\phi_X : V^X \to \{1, \dots, |V^X|\}$. We define $\phi : V_c \to \{1, \dots, N\}$ as the union of all V_X, which is well-defined, since the sets $V_X \in V_c$ are disjoint as they are well-separated.

Lemma 40. *For all $v \in V_c$ there is $\alpha_v > 0$, such that for all $X \in \mathcal{X}_c$, $f \in F^X$, $w \in V^X$ with $\phi(v) = \phi(w)$ we have $\alpha_v M_{v,f} \ge \alpha_w M_{w,f}$.*

In other words: For each vertex of color c there is a scalar, such that for each facet the scaled slack of the local vertex is the smallest for all vertices of the same index and color.

Proof outline. We can split our considerations by ϕ, that is we can find $\alpha_v > 0$ separately for each set of vertices $v \in V_c$ with the same index $\phi(v)$.

Let's fix some index $i \in \{1, \dots, N\}$. Now order all arcs X of color c with at least i vertices decreasing by arc-length as $X_1, \dots, X_m$. Also define $v_j \in V^{X_j}$ with index $i = \phi(v_j)$.

Now look at the submatrix $M\left[\{v_1, \dots, v_m\}, F^{X_1} \cup \dots \cup F^{X_m}\right]$ of the slack matrix, which we will treat kind of like a square matrix with one entry being $M\left[\{v_j\}, F^{X_j}\right]$.

For each $j \in \{1, \dots, m\}$ and $f \in (F^{X_1} \cup \dots \cup F^{X_m}) \setminus F^{X_j}$ we have $M_{v_j,f} > 0$, since v_j cannot be a vertex of f, because all X_k disjoint as they are well-separated.

Further, for each $j \in \{1, \dots, m\}$, $k \in \{1, \dots, j-1\}$, $f \in F^{X_j}$ and $g \in F^{X_1} \cup \dots \cup F^{X_{j-1}}$ we have $M_{v_j,f} M_{v_k,g} \le M_{v_j,g} M_{v_k,f}$.
To show this, first observe that X_j is the shortest of the regarded arcs. If we denote its length with ε, it has arc-distance at least 5ε of every arc X_k, because they are well-separated and ordered by decreasing length.

[16] Our results in this proof differ from the source. But since this number only changes a constant in the bound, we will go on with the original numbers.

The gist of the proof is following. Since M is representation dependent, reformulate the statement as a ratio of distances:

$$\frac{d(v_j, f)}{d(v_k, f)} \leq \frac{d(v_j, g)}{d(v_k, g)}$$

This can be done, since $M_{v_k,g} = 0$ verifies the statement and $M_{v_k,f} > 0$ holds, as the arcs are well-separated.
The idea is as follows: Because v_j is local to f, but v_k and g are relatively far away (at least 5ε in arc-distance), $d(v_j, f)$ is much smaller than $d(v_k, f)$ and $d(v_j, g)$. $d(v_k, g)$ may be large, but then we can argue using the triangle inequality that then also one of $d(v_k, f)$ or $d(v_j, g)$ has to be large.[17]

After those two properties of M are shown to be true, we can apply a technical lemma [8, Lemma 10.3], which gives us factors $\alpha_{v_1}, \ldots, \alpha_{v_m}$, such that $\alpha_{v_i} M_{v_i,f} \leq \alpha_{v_j} M_{v_j,f}$ for all $f \in F^{X_i}$. □

For each $i \in \{1, \ldots, N\}$ we define vectors $t^{(i)} \in \mathbb{R}^{|F_c|}$. For each $X \in \mathcal{X}_c$, $f \in F_X$:

$$t_f^{(i)} = \begin{cases} 0 & \text{if } i > |V^X| \\ \alpha_v M_{v,f} & \text{else, with } v = \phi_X^{-1}(i) \end{cases}$$

The vectors $t_f^{(i)}$ are the scaled slacks for the local vertices (regarding f) with index i.

We define the matrix $K \in \mathbb{R}^{|V| \times |F_c|}$, for which we will later apply the (general) "lampshade argument". For all $f \in F_c$ we set

$$K_{v,f} = \begin{cases} \alpha_v M_{v,f} - t_f^{(\phi(v))} & \text{if } v \in V_c, \\ M_{v,f} & \text{else.} \end{cases}$$

In other words, we manipulate the rows of K with vertices of color c, by scaling them, such that for vertices with the same index the slack of the local vertex is smallest, and then subtracting that smallest scaled slack.

The purpose of this manipulation is that for each $X \in \mathcal{X}_c$, any vertex $v \in V^X$ and any facet $f \in F^X$, we have $K_{v,f} = 0$, i.e. all entries of $K[V^X, F^X]$ are zero, while K is still nonnegative (that's why we had to scale all rows first).

Lemma 41. *The matrix K has nonnegative entries and satisfies* $\operatorname{rank}_+ K \leq 8|\mathcal{X}_c|$.

For the proof we need the generalized "lampshade argument", how it was originally proved by Shitov [14, Lemma 3.1]. It was simplified for this use case to allow simpler notation.

Lemma 42. *Let P be a polygon, and let V, F and M be defined like above.*

Let $X \subseteq V$ be a set of consecutive vertices of P, and let $F' \subseteq F$ be the set of facets of P with both endpoints in X.

Let M' be a matrix with rows indexed by $V \setminus X$ and columns indexed by F', such that the following condition holds for each vertex $v \in V \setminus X$:

There are real numbers $\alpha_v > 0$, $\beta_v \geq 0$ and a vertex $x_v \in X$ such that
$M'_{v,f} = \alpha_v M_{v,f} - \beta_v M_{x_v,f}$ for all $f \in F'$.

Then, if all the entries of M' are nonnegative, we have $\operatorname{rank}_+ M' \leq 8$.

[17] Remark that this is the only place we use that P is cyclic beside Definition 37.

We will not prove this lemma here. But remark, that it does not require P to be cyclic.

Proof outline of Lemma 41. We partition K into submatrices $K[V, F^X]$ for $X \in \mathcal{X}_c$ and show that each of those $|\mathcal{X}_c|$ slices has nonnegative rank at most 8.

We know $K[V^X, F^X] = 0$ by definition of K. For estimating the nonnegative rank, we only have to consider $K[V \setminus V^X, F^X]$. There are three possibilities for rows of this submatrix:

- If $v \notin V_c$, then $K_{v,f} = M_{v_f}$ for all $f \in F^X$.
- If $v \in V_c$ and $\phi(v) > |V^X|$, then $K_{v,f} = \alpha_v M_{v,f}$ for all $f \in F^X$, because $t_f^{(\phi(v))} = 0$.
- If $v \in V_c$ and $\phi(v) \leq |V^X|$, then $K_{v,f} = \alpha_v M_{v,f} - \alpha_{x_v} M_{x_v,f}$ for all $f \in F^X$, where $x_v \in V^X$ with $\phi(x_v) = \phi(v)$ is the vertex local to f with the same index as v.

From this we can see that $K[V \setminus V^X, F^X]$ is nonnegative by the choice of α_v in Lemma 40.

To show $\operatorname{rank}_+ K[V \setminus V^X, F^X] \leq 8$ we apply Lemma 42 for each of the above cases with V^X as our set of consecutive vertices.
We have to choose α_v, β_v and x_v for each vertex $v \in V \setminus V^X$.

- $v \notin V_c$: $\alpha_v = 1$, $\beta_v = 0$ and $x_v \in V^X$ arbitrary.
- $v \in V_c$ and $\phi(v) > |V^X|$: α_v already defined, $\beta_v = 0$ and $x_v \in V^X$ arbitrary.
- $v \in V_c$ and $\phi(v) \leq |V^X|$: α_v already defined, choose $x_v \in V^X$ with $\phi(x_v) = \phi(v)$ like above and set $\beta_v = \alpha_{x_v}$.

By Lemma 42 and $K[V^X, F^X] = 0$ we can conclude:

$$\operatorname{rank}_+ K[V, F^X] = \operatorname{rank}_+ K[V \setminus V^X, F^X] \leq 8$$

□

As a result, there are $8\,|\mathcal{X}_c|$ nonnegative vectors, such that each row of K can be written as nonnegative linear combination of them. Together with the N vectors $t^{(i)}$ we can write every row of $M[V, F_c]$ as a nonnegative linear combination of $8\,|\mathcal{X}_c| + N$ vectors. Therefore, $\operatorname{rank}_+ M[V, F_c] \leq 8\,|\mathcal{X}_c| + N$.

In summary:

$$\operatorname{xc}(P) = \operatorname{rank}_+ M \leq \sum_{c=1}^{14} \operatorname{rank}_+ M[V, F_c] \leq 8|\mathcal{X}| + 14N \leq 22\, n^{1/2} + 36 \overset{(n \geq 18^2)}{\leq} 24\, n^{1/2}$$

2.3 Comparison

In this section we compare the results and methods from Shitov [15] and Kwan, Sauermann, and Zhao [8] by working out general similarities and differences. Further, we look at how both bounds are achieved numerically. At the end, we examine how the proofs could be altered to be used in the other's setting.

We start with a quick overview of both results:

- Shitov [15] proves Theorem 6, stating that every n-gon P has $\mathrm{xc}(P) \in O(n^{2/3})$.
- Kwan, Sauermann, and Zhao [8] prove Theorem 36, which states that every cyclic n-gon P has $\mathrm{xc}(P) \in O(n^{1/2})$.

First we list the general similarities of both approaches:

- Both follow the strategy of splitting the polygon into smaller slices.
  Shitov splits the polygon into 12 slices which are treated separately.
  Kwan, Sauermann, and Zhao split the polygon into $O(n^{1/2})$ arcs containing $O(n^{1/2})$ vertices, but they are handled interdependently.
- There is some notion of enveloping vertices in both.
  Shitov requires an envelope around vertices in the central Theorem 24. It is used to build an extended formulation for this specific polygon.
  Kwan, Sauermann, and Zhao use a "lampshade argument", which builds a polygon around vertices away from a set of facets, proving that a part of the slack matrix has constant nonnegative rank.

Even though both methods have some ideas in common, they are very different.

- Shitov uses a purely geometric approach.
  Kwan, Sauermann, and Zhao use the linear-algebraic strategy using slack matrices for their main reasoning. But lemmata are often proved geometrically. The approach has the advantage that transformations on the slack matrix, which don't alter the nonnegative rank, can not always be represented geometrically.
- There is another difference in how they treat subsets of vertices:
  Shitov inductively extracts a large subsequence with small extension complexity (ignoring all other vertices).
  Kwan, Sauermann, and Zhao split their consideration by constantly many colors, but always handle the dependencies to all other vertices.

We continue by analyzing how each approach gets to its asymptotical bound numerically.

Shitov's result of $\mathrm{xc}(P) \in O(n^{2/3})$ for arbitrary n-gons P:

- There is a subsequence u with $m \in O(n^{2/3})$ vertices for every n-gon, for which the main theorem can be applied asymptotically optimal.

- It provides the bound $\operatorname{xc}(u) \in O(m^{1/2}) = O(n^{1/3})$ for the extension complexity of that subsequence.
- Inductively extracting such a subsequence proves the bound $\operatorname{xc}(P) \in O(n^{2/3})$.

Kwan, Sauermann, and Zhao's result of $\operatorname{xc}(P) \in O(n^{1/2})$ for cyclic n-gons P:

- The circle is split into $|\mathcal{X}| \in O(n^{1/2})$ arcs.
- The slack matrix is split into $O(1)$ "color columns", for which a matrix K with $\operatorname{rank}_+ K \in O(|\mathcal{X}_c|)$ is constructed ($\mathcal{X}_c$ are the arcs of color c).
- Each "color column" in the slack matrix has $\operatorname{rank}_+ M[V, F_c] \in O(|\mathcal{X}_c| + n^{1/2})$.
- Joining these columns results in $\operatorname{rank}_+ M \in O(|\mathcal{X}| + n^{1/2}) = O(n^{1/2})$.

Finally, we have a look at how each procedure could be used in the other's setting:

- Even though Kwan, Sauermann, and Zhao state that Lemma 10.2 is "the only place where we use this assumption that P is cyclic" [8, p. 22], the definition of well-separateness uses arc-distance. Generalizing the proof requires something like a boundary-distance, where the distance between two vertices is the sum of the lengths of the edges separating them. In this setting, a counterexample for Lemma 10.2 can be found easily, because well-separated facets and vertices can still have small slacks (see Figure 15).
  We conclude that this approach can't be adopted easily for the general case, since it relies on P being cyclic in central parts. Still, the methods used may be of value for future approaches. This also aligns well with the authors' opinion [8, p. 28].
- If we want to improve the bound achieved by Shitov for cyclic polygons, we have to find a larger subsequence for which we can apply Theorem 24 optimally. In the best case, we can find a subsequence with $\Theta(n)$ vertices, where $|V|, |S|, \delta \in \Theta(n^{1/2})$ in Theorem 24. Then we could obtain the optimal upper bound of $O(n^{1/2})$ for the extension complexity.
  Here, we were not able to provide such an improvement.

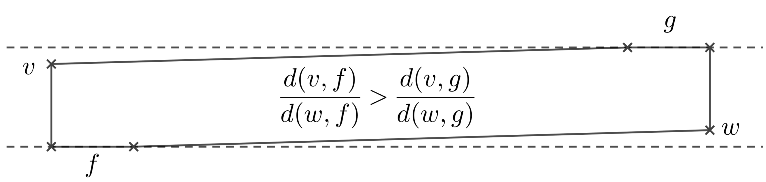

Figure 15. Counterexample for Lemma 10.2 [8], if P weren't cyclic.

3 Lower Bounds for the Extension Complexity

In this section we prove that there is a polygon P with $\mathrm{xc}(P) \in \Omega(n^{1/2})$.

After the proof, we will explain the key concepts of the applied theorem and show how to apply it in general.

The central piece of the proof is Theorem 1 from Averkov, Kaibel, and Weltge [1], which we adopt for the case of linear extended formulations.

We define $\|\cdot\|$ to be the Euclidean norm and $\mathbb{B}^d := \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}$ to be the d-dimensional unit ball.

3.1 Introducing the Applied Theorem

First we have to define the *Hausdorff distance* of two compact sets, which can be thought of "the longest distance one can be forced to travel from a point in one of the two sets to the other set". See Figure 16 for an example.

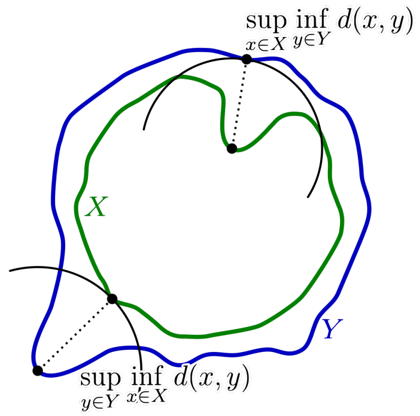

Figure 16. Hausdorff distance example. [11]

Definition 43 (Hausdorff distance)**.** The *Hausdorff distance* of two non-empty compact sets $X, Y \subseteq \mathbb{R}^d$ is defined by

$$\mathrm{dist}_H(X,Y) := \max\left\{\sup_{x\in X}\inf_{y\in Y}\|x-y\|, \sup_{y\in Y}\inf_{x\in X}\|x-y\|\right\}.$$

Theorem 44 ([1, Theorem 1], adopted)**.** *Let $\mathcal{P}$ be a family of polytopes in $\mathbb{R}^d$ of dimensions at least one with $2 \leq |\mathcal{P}| < \infty$ such that each $P \in \mathcal{P}$ has an extended formulation of size m. Let $\rho > 0$ and $\Delta > 0$ be such that each $P \in \mathcal{P}$ is contained in the ball $\rho\mathbb{B}^d$ and, for every two distinct polytopes $P \in \mathcal{P}$ and $P' \in \mathcal{P}$, one has $\mathrm{dist}_H(P, P') \geq \Delta$. Then*

$$m^2 \geq \frac{\log_2|\mathcal{P}|}{8d\,(1 + \log_2(2\rho/\Delta) + \log_2\log_2|\mathcal{P}|)} =: B.$$

In particular, we have

$$\max\{\mathrm{xc}(P) \mid P \in \mathcal{P}\} \geq \sqrt{B}.$$

3.2 Proving the Lower Bound

Corollary 45. *There exists a (cyclic) polygon P with* $\mathrm{xc}(P) \in \Omega(n^{1/2})$.

Proof. To apply Theorem 44 we have to pick a family of polytopes $\mathcal{P}$. Then we have to determine ρ and Δ and bound $\log_2|\mathcal{P}|$ from below and $\log_2 \log_2|\mathcal{P}|$ from above.

So we choose n^2 fixed points evenly on the unit circle such that they would form a regular polygon. Let $\mathcal{P}_n$ be the family of polygons, where each polygon is the convex hull of n points chosen from the n^2 points on the unit circle.

Then $\rho = 1$ holds, since all polygons are contained in the unit circle.

For two distinct polygons $P, P' \in \mathcal{P}_n$ one of them has a vertex v, which the other one does not have. W.l.o.g. $v \notin P, v \in P'$. As a result, $\mathrm{dist}_H(P, P') \geq \inf_{p \in P} \|v - p\| \geq d_{min}$. See Figure 17 for the definition of d_{min}.

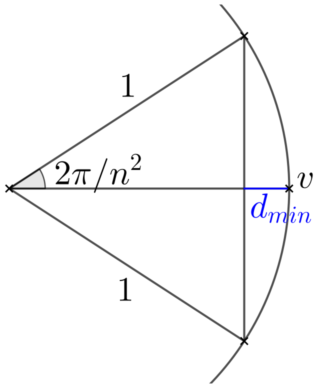

Figure 17. Definition of d_{min}.

$$1 - d_{min} = \cos\left(\frac{2\pi}{n^2}\right)$$

$$d_{min} \geq 1 - \left(1 - \frac{\left(\frac{2\pi}{n^2}\right)^2}{2}\right) = 2\frac{\pi^2}{n^4} =: \Delta$$

For the estimate we used $\cos(x) \geq 1 - \frac{x^2}{2!}$.

Since each polygon in $\mathcal{P}_n$ is defined by n points, chosen from a set of n^2 points, $|\mathcal{P}_n| = \binom{n^2}{n}$. With $(\frac{n}{k})^k \leq \binom{n}{k} \leq n^k$ we can estimate $n^n = \left(\frac{n^2}{n}\right)^n \leq |\mathcal{P}_n| \leq (n^2)^n = n^{2n}$, $\log_2|\mathcal{P}_n| \geq n \log_2 n$ and $\log_2 \log_2|\mathcal{P}_n| \leq \log_2(2n \log_2 n) = 1 + \log_2 n + \log_2 \log_2 n$.

Now we can estimate B from Theorem 44:

$$\begin{aligned} B &= \frac{\log_2|\mathcal{P}_n|}{8d\,(1 + \log_2(2\rho/\Delta) + \log_2 \log_2|\mathcal{P}_n|)} \\ &\geq \frac{n \log_2 n}{16\,(1 + \log_2(n^4/\pi^2) + 1 + \log_2 n + \log_2 \log_2 n)} \\ &= \frac{n \log_2 n}{16\,(2 - 2\log_2 \pi + 5\log_2 n + \log_2 \log_2 n)} \\ &\geq \frac{n}{16 \cdot 6} \end{aligned}$$

For the last inequality we used $2 - 2\log_2 \pi \leq 0$ and $\frac{5\log_2 n + \log_2 \log_2 n}{\log_2 n} \leq 6$ for $n \geq 1$.

Therefore, we can conclude

$$\max\{\mathrm{xc}(P) \mid P \in \mathcal{P}_n\} \overset{\text{(Th. 44)}}{\geq} \sqrt{B} \geq \frac{1}{12}\sqrt{n}.$$

□

3.3 Key Ideas Behind the Applied Theorem

The central idea of Theorem 44 is to encode the extended formulations. Then one counts how many extended formulations fit into the containing volume, since those formulations are separated themselves.

Here is a short outline of the proof:

1. Each extended formulation can be normalized, i.e. $P = \varphi(Q) + t$ with $Q = \{x \mid Ax \leq \mathbb{1}\}$, $\mathbb{B}^n \subseteq Q \subseteq n\mathbb{B}^n$ and φ being linear.
2. Each polytope $P_i \in \mathcal{P}$, where $P_i = \varphi_i(Q_i) + t_i$ and $Q_i = \{x \in \mathbb{R}^{n_i} \mid A_i x \leq \mathbb{1}\})$, is encoded trough $P_i \mapsto (A_i, \varphi_i, t_i)$.
   From now on everything takes place in the encoding vector spaces $\mathcal{V}^n$, where n is the dimension of the extended formulation.
3. The encodings are partitioned into sets W_i by the dimension of Q_i.
4. The distance between two encodings $w, w' \in W_i$ can be bounded by $\|w - w'\| \geq \Delta$.
5. The overall space required can be bounded by $\|w\| \leq 3\rho n^2, \forall w \in W_i$.
6. One creates small balls around each encoding with radius $\Delta/2$ and a ball containing all small balls with radius $3\rho n^2 + \Delta/2$ around the origin.
7. By comparing volumes, one can bound the number of possible polytope encodings $|W_i|$ from above depending on ρ, Δ, m and d. Note that the dimension of $\mathcal{V}^n$ and therefore the computed volume depend on m and d.
8. These bounds are added by $|\mathcal{P}| = \sum |W_i|$ and solved for m.

The lower bounding for the extension complexity, i.e. $\max\{\mathrm{xc}(P) \mid P \in \mathcal{P}\} \geq \sqrt{B}$, is achieved by inverting the implication:

$$\big((\forall P \in \mathcal{P} : \mathrm{xc}(P) \leq m) \Rightarrow m^2 \geq B\big) \Rightarrow \big(m^2 < B \Rightarrow (\exists P \in \mathcal{P} : \mathrm{xc}(P) > m)\big)$$

This theorem can be applied like in Corollary 45 by picking a family $\mathcal{P}$ and bounding ρ from above, Δ from below, $\log_2|\mathcal{P}|$ from below and $\log_2 \log_2|\mathcal{P}|$ from above. It then provides a lower bound for the largest extension complexity of one polytope $P \in \mathcal{P}$.

4 Concluding Remarks

In this paper we looked at the best known bounds for the extension complexity of polygons.

For cyclic polygons P we know $\mathrm{xc}(P) \in \Theta(n^{1/2})$.

For arbitrary polygons P, we only know $\mathrm{xc}(P) \in \Omega(n^{1/2}) \cap O(n^{2/3})$.

Kwan, Sauermann, and Zhao conjectured that cyclic polygons have worst-case extension complexity, i.e. $\mathrm{xc}(P) \in \Theta(n^{1/2})$, because they "seem to represent quite a diverse cross-section of the space of all polygons" [8, p. 3].

On the other hand, Shitov expected that $\mathrm{pc}(n) = n^{1/2} \cdot \alpha(n)$ with unbounded $\alpha(n)$ [15, Conjecture 61].

He reasons that the method developed in his paper doesn't seem to allow an $O(n^{1/2})$ upper bound for the worst-case n-gon complexity. This means there is no subsequence of $\Theta(n)$ vertices, for which we could apply Theorem 24 optimally.

And he also goes on explaining that Padrol [9] proved

$$\mathrm{wcc}(d, n) \geq 2\sqrt{dn - d} - d + 1, \tag{3}$$

where $\mathrm{wcc}(d, n)$ is the largest possible extension complexity of a polytope with n vertices in a d-dimensional space. If the conjecture were false, the bound in (3) would become asymptotically optimal for $d = 2$ This is not expected, since Kaibel and Weltge showed $\mathrm{wcc}(m^2, 2^m) \geq 1.5^m$ for the *correlation polytope* [7]. This is asymptotically much greater than $(\sqrt{2})^m$, which is the result of equation (3) in that case.

Thus, the question of worst-case extension complexity for polygons remains exciting.

A Proof Dependency Graph

The following graph gives an overview of the dependencies in the proofs from Shitov [15].

Each node corresponds to a theorem, lemma or corollary given by its number. Each edge indicates a dependency on the statement it is directed at. And the boxes group statements from different sections ("Sec." is short for "Section").

Definitions were omitted, as they made the graph too crowded. Also, some edges might be missing, because dependencies aren't always obvious, when one is immersed in a subject for some time.

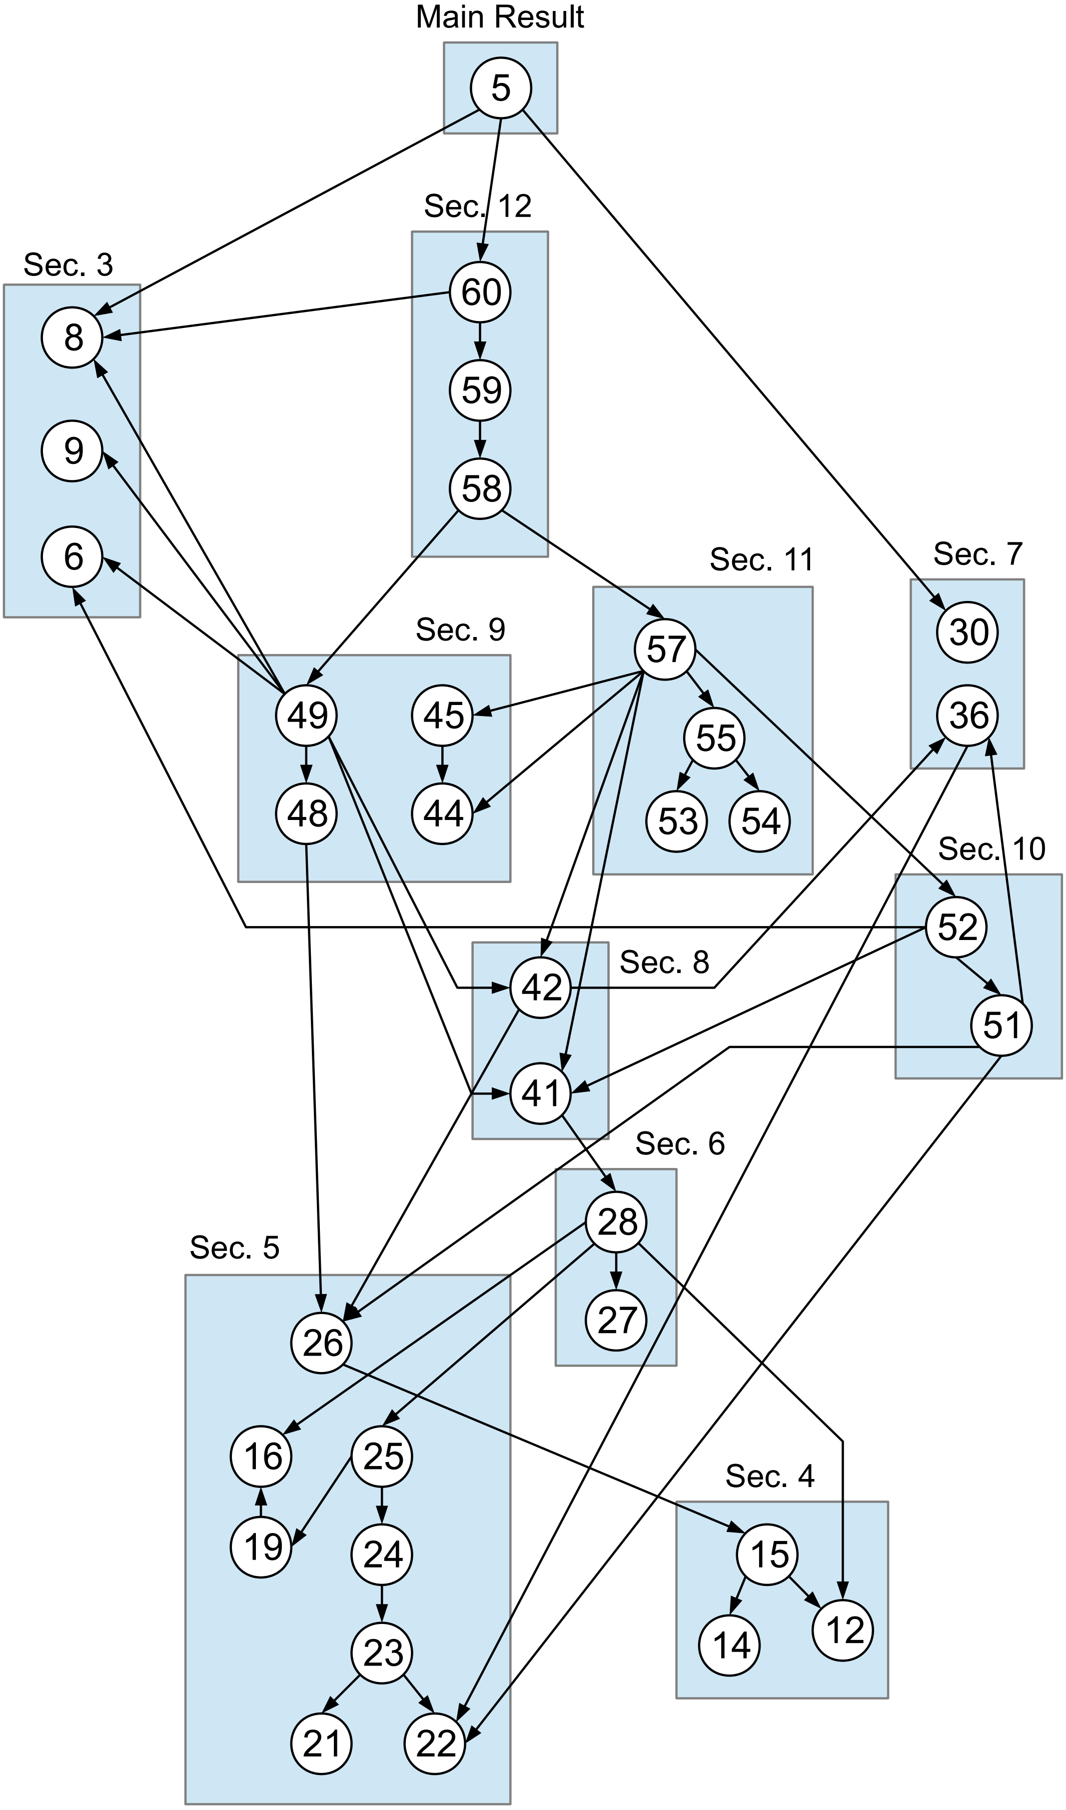
Main Result
5
Sec. 12
60
59
58
Sec. 3
8
9
6
Sec. 11
57
55
53
54
Sec. 7
30
36
Sec. 9
49
45
48
44
Sec. 10
52
51
Sec. 8
42
41
Sec. 6
28
27
Sec. 5
26
16
25
19
24
23
21
22
Sec. 4
15
14
12

References

[1] Gennadiy Averkov, Volker Kaibel, and Stefan Weltge. *Maximum Semidefinite and Linear Extension Complexity of Families of Polytopes*. 2016. arXiv: `1605.08538 [math.OC]`.

[2] Gábor Braun and Sebastian Pokutta. "The matching problem has no fully polynomial size linear programming relaxation schemes". In: *IEEE Transactions on Information Theory* 61.10 (2015), pp. 5754–5764.

[3] Samuel Fiorini, Thomas Rothvoß, and Hans Raj Tiwary. "Extended formulations for polygons". In: *Discrete & computational geometry* 48.3 (2012), pp. 658–668.

[4] Samuel Fiorini et al. "Exponential lower bounds for polytopes in combinatorial optimization". In: *Journal of the ACM (JACM)* 62.2 (2015), pp. 1–23.

[5] Lester Randolph Ford and Delbert R Fulkerson. "Maximal flow through a network". In: *Canadian journal of Mathematics* 8 (1956), pp. 399–404.

[6] Volker Kaibel and Kanstantsin Pashkovich. *Constructing Extended Formulations from Reflection Relations*. 2010. arXiv: `1011.3597 [math.CO]`.

[7] Volker Kaibel and Stefan Weltge. "A short proof that the extension complexity of the correlation polytope grows exponentially". In: *Discrete & Computational Geometry* 53.2 (2015), pp. 397–401.

[8] Matthew Kwan, Lisa Sauermann, and Yufei Zhao. *Extension complexity of low-dimensional polytopes*. 2020. arXiv: `2006.08836 [math.CO]`.

[9] Arnau Padrol. "Extension complexity of polytopes with few vertices or facets". In: *SIAM Journal on Discrete Mathematics* 30.4 (2016), pp. 2162–2176.

[10] Arnau Padrol and Julian Pfeifle. "Polygons as sections of higher-dimensional polytopes". In: *arXiv preprint arXiv:1404.2443* (2014).

[11] Rocchini. *Hausdorff distance sample*. 2007. URL: `https://commons.wikimedia.org/wiki/File:Hausdorff_distance_sample.svg` (visited on 03/28/2021). CC BY 3.0 (https://creativecommons.org/licenses/by/3.0), via Wikimedia Commons.

[12] Thomas Rothvoß. "The matching polytope has exponential extension complexity". In: *Journal of the ACM (JACM)* 64.6 (2017), pp. 1–19.

[13] Yaroslav Shitov. "An upper bound for nonnegative rank". In: *Journal of Combinatorial Theory, Series A* 122 (2014), pp. 126–132.

[14] Yaroslav Shitov. *Sublinear extensions of polygons*. 2014. arXiv: `1412.0728v1 [math.CO]`.

[15] Yaroslav Shitov. *Sublinear extensions of polygons*. 2020. arXiv: `1412.0728v2 [math.CO]`.

[16] Stefan Weltge. "Sizes of linear descriptions in combinatorial optimization". PhD thesis. Otto-von-Guericke-Universität Magdeburg, 2015.

[17] Mihalis Yannakakis. "Expressing combinatorial optimization problems by linear programs". In: *Journal of Computer and System Sciences* 43.3 (1991), pp. 441–466.