

Extended Formulations for Polygons

Samuel Fiorini · Thomas Rothvoß ·
Hans Raj Tiwary

Received: 13 August 2011 / Revised: 24 February 2012 / Accepted: 28 February 2012 /
Published online: 16 March 2012
© Springer Science+Business Media, LLC 2012

Abstract The extension complexity of a polytope P is the smallest integer k such that P is the projection of a polytope Q with k facets. We study the extension complexity of n -gons in the plane. First, we give a new proof that the extension complexity of regular n -gons is $O(\log n)$, a result originating from work by Ben-Tal and Nemirovski (Math. Oper. Res. 26(2), 193–205, 2001). Our proof easily generalizes to other permutahedra and simplifies proofs of recent results by Goemans (2009), and Kaibel and Pashkovich (2011). Second, we prove a lower bound of $\sqrt{2n}$ on the extension complexity of generic n -gons. Finally, we prove that there exist n -gons whose vertices lie on an $O(n) \times O(n^2)$ integer grid with extension complexity $\Omega(\sqrt{n}/\sqrt{\log n})$.

Keywords Extended formulations · Polygon · Polytope · Lower bound

1 Introduction

Consider a (convex) polytope P in \mathbb{R}^d . An *extension* (or *extended formulation*) of P is a polytope Q in \mathbb{R}^e such that P is the image of Q under a linear projection from \mathbb{R}^e to \mathbb{R}^d . The main motivation for seeking extensions Q of the polytope P is perhaps that the number of facets of Q can sometimes be significantly smaller than


S. Fiorini · H.R. Tiwary (✉)

Department of Mathematics, Université Libre de Bruxelles, Brussels, Belgium
e-mail: htiwary@ulb.ac.be

S. Fiorini
e-mail: sfiorini@ulb.ac.be

T. Rothvoß
Department of Mathematics, MIT, Cambridge, USA
e-mail: rothvoss@math.mit.edu

Fig. 1 Proof by picture that the extension complexity of a regular 8-gon is at most 6. Here $P \subseteq \mathbb{R}^2$ is a regular 8-gon, $Q \subseteq \mathbb{R}^3$ is a polytope combinatorially equivalent to a 3-cube, and $\pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is a linear projection map such that $\pi(Q) = P$



that of P . This phenomenon has already found numerous applications in optimization, and in particular linear and integer programming. To our knowledge, systematic investigations began at the end of the 1980s with the work of Martin [13] and Yannakakis [17], among others. Recently, the subject is receiving an increasing amount of attention. See, e.g., the surveys by Conforti, Cornuéjols and Zambelli [4], Vanderbeck and Wolsey [16], and Kaibel [10].

A striking example, which is relevant to this paper, arises when P is a regular n -gon in \mathbb{R}^2 . As follows from results of Ben-Tal and Nemirovski [2], for such a polytope P , one can construct an extension Q with as few as $O(\log n)$ facets. It remained an open question to determine to which extent such a dramatic decrease in the number of facets is possible when P is a *non-regular* n -gon.¹ This is the main question we address in this paper.

Before giving an outline of the paper, we state a few more definitions. The *size* of an extension Q is simply the number of facets of Q . The *extension complexity* of P is the minimum size of an extension of P , denoted as $xc(P)$. See Fig. 1 for an illustration.

Notice that the extension complexity of every n -gon is $\Omega(\log n)$. This follows from the fact that any extension Q with k facets has at most 2^k faces. Since each face of P is the projection of a face of the extension Q , it follows that Q must have at least $\log_2 f$ facets if P has f faces [7]. Thus if P is an n -gon, we have $xc(P) \geq \log_2(2n + 2) = \Omega(\log n)$. When P is a regular n -gon, we have $xc(P) = \Theta(\log n)$.

One of the fundamental results that can be found in Yannakakis’ groundbreaking paper [17] is a characterization of the extension complexity of a polytope in terms of the non-negative rank of its slack matrix. Although this is discussed in detail in Sect. 2, we include a brief description here. To each polytope P one can associate a matrix $S(P)$ that records, in the entry that is in the i th row and j th column, the slack of the j th vertex with respect to the i th facet. This matrix is the ‘slack matrix’ of P . It turns out that computing $xc(P)$ amounts to determining the minimum number r such that there exists a factorization of the slack matrix of P as $S(P) = TU$, where T is a non-negative matrix with r columns and U is a non-negative matrix with r rows. Such a factorization is called a ‘rank r non-negative factorization’ of the slack matrix $S(P)$.

In Sect. 3, we give an explicit $O(\log n)$ rank non-negative factorization of the slack matrix of a regular n -gon. This provides a new proof that the extension complexity

¹This was posed as an open problem during the First Cargese Workshop on Combinatorial Optimization.

of every regular n -gon is $O(\log n)$. Our proof technique directly generalizes to other polytopes, such as the permutahedron. In particular, we obtain a new proof of the fact that the extension complexity of the n -permutahedron is $O(n \log n)$, a result due to Goemans [7]. Our approach builds on a new proof of this result by Kaibel and Pashkovich [11] but is different because it works by directly constructing a non-negative factorization of the slack matrix.

In Sect. 4, we prove that there exist n -gons whose extension complexity is at least $\sqrt{2n}$. However, the proof uses polygons whose coordinates are transcendental numbers, which is perhaps not entirely satisfactory. For instance, one might ask whether a similar result holds when the encoding length of each vertex of the polygon is $O(\log n)$.

In Sect. 5, we settle this last question by proving the existence of n -gons whose vertices belong to an $O(n) \times O(n^2)$ integer grid and with extension complexity $\Omega(\sqrt{n}/\sqrt{\log n})$. This is inspired by recent work of one of the authors on the extension complexity of 0/1-polytopes [14].

2 Slack Matrices and Non-negative Factorizations

Consider a polytope P in \mathbb{R}^d with m facets and n vertices. Let $A_1x \leq b_1, \dots, A_mx \leq b_m$ denote the facet-defining inequalities of P , where A_1, \dots, A_m are row vectors. Let also v_1, \dots, v_n denote the vertices of P . The *slack matrix* of P is the non-negative $m \times n$ matrix $S = S(P)$ with $S_{ij} = b_i - A_iv_j$.

A *rank r non-negative factorization* of a non-negative matrix S is an expression of S as product $S = TU$ where T and U are non-negative matrices with r columns and r rows, respectively. The *non-negative rank* of S , denoted by $\text{rank}_+(S)$, is the minimum number r such that S admits a rank r non-negative factorization [3].

The following theorem is (essentially) due to Yannakakis, see also [6].

Theorem 1 (Yannakakis [17]) *For all polytopes P ,*

$$\text{xc}(P) = \text{rank}_+(S(P)).$$

To conclude this section, we briefly indicate how to obtain extensions from non-negative factorizations, and prove half of Theorem 1. Assuming $P = \{x \in \mathbb{R}^d : Ax \leq b\}$, consider a rank r non-negative factorization $S(P) = TU$ of the slack matrix of P . Then it can be shown that the image of the polyhedron $Q := \{(x, y) \in \mathbb{R}^{d+r} \mid Ax + Ty = b, y \geq 0\}$ under the projection $\mathbb{R}^{d+r} \rightarrow \mathbb{R}^d : (x, y) \mapsto x$ is exactly P . Notice that Q has at most r facets. Now if we take $r = \text{rank}_+(S(P))$, then Q is actually a polytope [5]. Thus Q is an extension of P with at most $\text{rank}_+(S(P))$ facets, and hence $\text{xc}(P) \leq \text{rank}_+(S(P))$.

3 Regular Polygons

First, we give a new proof of the tight logarithmic upper bound on the extension complexity of a regular n -gon. This result is implicit in work by Ben-Tal and Nemirovski [2] (although for n being a power of two). Another proof can be found in

Kaibel and Pashkovich [11]. Then, we discuss a generalization of the proof to related higher-dimensional polytopes.

Theorem 2 *Let P be a regular n -gon in \mathbb{R}^2 . Then $xc(P) = O(\log n)$.*

Proof Without loss of generality, we may assume that the origin is the barycenter of P . After numbering the vertices of P counterclockwise as v_1, \dots, v_n , we define a sequence $\ell_0, \dots, \ell_{q-1}$ of axes of symmetry of P , as follows.

Initialize i to 0, and k to n . While $k > 1$, repeat the following steps:


- define ℓ_i as the line through the origin and the midpoint of vertices $v_{\lfloor \frac{k}{2} \rfloor}$ and $v_{\lfloor \frac{k+1}{2} \rfloor}$;
- replace k by $\lfloor \frac{k+1}{2} \rfloor$;
- increase i by one.

Define q as the final value of i . Thus, q is the number of axes of symmetry ℓ_i defined. Note that when $k = k(i)$ is odd, then ℓ_i passes through one of the vertices of P . Note also that $q = O(\log n)$. For each $i = 0, \dots, q - 1$, one of the two closed half-planes bounded by ℓ_i contains v_1 . We denote it ℓ_i^+ . We denote the other by ℓ_i^- .

Now, consider a vertex v of P . We define the *folding sequence* $v^{(0)}, v^{(1)}, \dots, v^{(q)}$ of v as follows. We let $v^{(0)} := v$, and for $i = 0, \dots, q - 1$, we let $v^{(i+1)}$ denote the image of $v^{(i)}$ by the reflection with respect to ℓ_i if $v^{(i)}$ is not in the half-space ℓ_i^+ , and we let $v^{(i+1)} := v^{(i)}$ otherwise. In other words, $v^{(i+1)}$ is the image of $v^{(i)}$ under the *conditional reflection* with respect to half-plane ℓ_i^+ . By construction, we always have $v^{(q)} = v_1$.

Next, consider a facet F of P . The folding sequence $F^{(0)}, F^{(1)}, \dots, F^{(q)}$ of facet F is defined similarly as the folding sequence of vertex v . Pick any inequality $a^T x \leq \beta$ defining F . We let $a^{(0)} := a$, and for $i = 0, \dots, q - 1$, we let $a^{(i+1)}$ denote the image of $a^{(i)}$ under the conditional reflection with respect to ℓ_i^+ . Then $F^{(i)}$ is the facet of P defined by $(a^{(i)})^T x \leq \beta$. The last facet $F^{(q)}$ in the folding sequence is always either the segment $[v_1, v_2]$ or the segment $[v_1, v_n]$. See Fig. 2 for an illustration with $n = 15$, and thus $q = 4$.

Fig. 2 A 15-gon with four axes of symmetry, a vertex- and a facet-folding sequence



Finally, we define a non-negative factorization $S(P) = TU$ of the slack matrix of P , of rank $2q = O(\log n)$. Below, let $d(x, \ell_i)$ denote the distance of $x \in \mathbb{R}^2$ to line ℓ_i .

In the left factor of the factorization, the row corresponding to facet F is of the form (t_0, \dots, t_{q-1}) , where $t_i := (\sqrt{2}d(a^{(i)}, \ell_i), 0)$ if $a^{(i)}$ is not in ℓ_i^+ and $t_i := (0, \sqrt{2}d(a^{(i)}, \ell_i))$ otherwise. Similarly, in the right factor, the column corresponding to vertex v is of the form $(u_0, \dots, u_{q-1})^T$, where $u_i := (0, \sqrt{2}d(v^{(i)}, \ell_i))^T$ if $v^{(i)}$ is not in ℓ_i^+ and $u_i := (\sqrt{2}d(v^{(i)}, \ell_i), 0)^T$ otherwise.

The correctness of the factorization rests on the following simple observation: for $i = 0, \dots, q - 1$ the slack of $v^{(i+1)}$ with respect to $F^{(i+1)}$ equals the slack of $v^{(i)}$ with respect to $F^{(i)}$ plus some correction term. If $a^{(i)}$ and $v^{(i)}$ are on opposite sides of ℓ_i , then the correction term is $2d(a^{(i)}, \ell_i)d(v^{(i)}, \ell_i)$. Otherwise, it is zero (no correction is necessary). Indeed, letting n_i denote a unit vector normal to ℓ_i , and assuming that $v^{(i)}$ and $a^{(i)}$ are on opposite sides of ℓ_i , we have

$$\begin{aligned} \beta - (a^{(i)})^T v^{(i)} &= \beta - (a^{(i)})^T (v^{(i)} - 2(n_i^T v^{(i)})n_i + 2(n_i^T v^{(i)})n_i) \\ &= \beta - (a^{(i+1)})^T v^{(i+1)} - 2((a^{(i)})^T n_i)(n_i^T v^{(i)}) \\ &= \beta - (a^{(i+1)})^T v^{(i+1)} + 2d(a^{(i)}, \ell_i)d(v^{(i)}, \ell_i). \end{aligned}$$

When $v^{(i)}$ and $a^{(i)}$ are on the same side of ℓ_i , we obviously have

$$\beta - (a^{(i)})^T v^{(i)} = \beta - (a^{(i+1)})^T v^{(i+1)}.$$

Observe that the slack of $v^{(q)}$ with respect to $F^{(q)}$ is always 0. The theorem follows. □

The n -permutahedron is the polytope of dimension $n - 1$ in \mathbb{R}^n whose $n!$ vertices are the points obtained by permuting the coordinates of $(1, 2, \dots, n)^T$. It has $2^n - 2$ facets, defined by the inequalities $\sum_{j \in S} x_j \leq g(|S|)$ for all proper non-empty subsets S of $[n] := \{1, 2, \dots, n\}$, where $g(S) := \binom{n+1}{2} - \binom{n-|S|+1}{2}$.

Let j and k denote two elements of $[n]$ such that $j < k$. We denote $H_{j,k}$ the hyperplane defined by $x_j = x_k$, and $H_{j,k}^+$ the closed half-space defined by $x_j \leq x_k$. Applying the conditional reflection with respect to $H_{j,k}^+$ to a vector $x \in \mathbb{R}^n$ amounts to swapping the coordinates x_j and x_k if and only if $x_j > x_k$. Intuitively, the conditional reflection with respect to $H_{j,k}^+$ sorts the coordinates x_j and x_k .

The proof of Theorem 2 can be modified to give a new proof of the existence of $O(n \log n)$ size extension of the n -permutahedron [7], as follows. Since there exists a sorting network of size $O(n \log n)$ for sorting n inputs, a celebrated result of Ajtai, Komlós and Szemerédi [1], there exist $q = O(n \log n)$ half-spaces $H_{j_0, k_0}^+, H_{j_1, k_1}^+, \dots, H_{j_{q-1}, k_{q-1}}^+$ such that sequentially applying the conditional reflection with respect to H_{j_i, k_i}^+ for $i = 0, \dots, q - 1$ to any point $x \in \mathbb{R}^n$, sorts this point x .

Therefore, the folding sequence of any vertex v of the n -permutahedron always ends with the vertex $(1, 2, \dots, n)^T$. Moreover, the folding sequence of the facet defined by $\sum_{j \in S} x_j \leq g(|S|)$ always ends with the facet defined by $\sum_{j=n-|S|+1}^n x_j \leq$

$g(|S|)$. Note that this last facet contains the vertex $(1, 2, \dots, n)^T$. Hence the proof technique used above for a regular n -gon extends to the n -permutahedron.

In fact, it turns out that the proof technique further extends to the permutahedron of any finite reflection group. One simply has to choose the right sequence of conditional reflections. Such sequences were constructed by Kaibel and Pashkovich [11], with the help of Ajtai–Komlós–Szemerédi sorting networks. Thus we can re-prove their main results about permutahedra of finite reflection groups. Our proof is different in the sense that we explicitly construct a non-negative factorization of the slack matrix.

4 Generic Polygons

We begin by recalling some basic facts about field extensions (see, e.g., Hungerford [9], Lang [12], or Stewart [15]). Let L be a field and K be a subfield of L . Then L is an *extension field* of K , and L/K is a *field extension*. We say that the field extension L/K is *algebraic* if every element of L is algebraic over K , that is, for each element of L there exists a non-zero polynomial with coefficients in K that has the element as one of its roots.

For $\alpha_1, \dots, \alpha_q \in L$, the inclusion-wise minimal subfield of L that contains both K and $\{\alpha_1, \dots, \alpha_q\}$ is denoted by $K(\{\alpha_1, \dots, \alpha_q\})$, or simply $K(\alpha_1, \dots, \alpha_q)$. It is also the subfield formed by all fractions $\frac{f(\alpha_1, \dots, \alpha_q)}{g(\alpha_1, \dots, \alpha_q)}$ where f and g are polynomials with coefficients in K and $g(\alpha_1, \dots, \alpha_q) \neq 0$.

A subset X of L is said to be *algebraically independent* over K if no non-trivial polynomial relation with coefficients in K holds among the elements of X . The *transcendence degree* of the field extension L/K is defined as the largest cardinality of an algebraically independent subset of L over K . It is also the minimum cardinality of a subset Y of L such that $L/K(Y)$ is algebraic.

We say that a polygon in \mathbb{R}^2 is *generic* if the coordinates of its vertices are distinct and form a set that is algebraically independent over the rationals.

Theorem 3 *If P is a generic convex n -gon in \mathbb{R}^2 then $xc(P) \geq \sqrt{2n}$.*

Proof Let $\alpha_1, \dots, \alpha_{2n}$ denote the coordinates of the n vertices of P , listed in any order. Thus $X := \{\alpha_1, \dots, \alpha_{2n}\}$ is algebraically independent over \mathbb{Q} .

Now suppose that P is the projection of a d -dimensional polytope Q with k facets. Without loss of generality, we may assume that Q lives in \mathbb{R}^d and that the projection is onto the two first coordinates.

Consider any linear description of Q . This description is defined by $k(d + 1)$ real numbers: the kd entries of the constraint matrix and the k right-hand sides. We denote these reals as $\beta_1, \dots, \beta_{k(d+1)}$. By Cramer’s rule, each α_i can be written as $\alpha_i = \frac{f_i(\beta_1, \dots, \beta_{k(d+1)})}{g_i(\beta_1, \dots, \beta_{k(d+1)})}$ where f_i and g_i are polynomials with rational coefficients and $g_i(\beta_1, \dots, \beta_{k(d+1)}) \neq 0$. In particular, this means that each α_i is in the extension field $L := \mathbb{Q}(\beta_1, \dots, \beta_{k(d+1)})$.

Since X is algebraically independent over \mathbb{Q} and $X \subseteq L$, the transcendence degree of L/\mathbb{Q} is at least $2n$. But on the other hand, the transcendence degree of L/\mathbb{Q} is at most $k(d + 1)$. Indeed, letting $Y := \{\beta_1, \dots, \beta_{k(d+1)}\}$, we have $\mathbb{Q}(Y) = L$ and thus

$L/\mathbb{Q}(Y)$ is algebraic. It follows that $k(d + 1) \geq 2n$. Since $k \geq d + 1$, we see that $k^2 \geq 2n$, hence $k \geq \sqrt{2n}$. \square

5 Polygons with Integer Vertices

Since encoding transcendental numbers would require an infinite number of bits, an objection might be raised that Theorem 3 is not very satisfying. In this section we provide a slightly weaker lower bound with polygons whose vertices can be encoded efficiently. In particular we will now show that for every n there exist polygons with vertices on an $O(n) \times O(n^2)$ grid and whose extension complexity is large. To do this we will need a slightly modified version of a rounding lemma proved by Rothvoß [14], see Lemma 5 below.

For a matrix A let A_ℓ (resp. A^ℓ) denote the ℓ th row (resp. ℓ th column) of A . Similarly, for a subset I of row indices of A , let A_I denote the submatrix of A obtained by picking the rows indexed by the elements of I .

Let T and U be $m \times r$ and $r \times n$ non-negative matrices. Since below T and U will be respectively the left and right factor of a factorization of some slack matrix, we can assume that no column of T is identically zero and, similarly, no row of U is identically zero. The pair T, U is said to be *normalized* if $\|T^\ell\|_\infty = \|U_\ell\|_\infty$ for every $\ell \in [r]$. Since multiplying a column ℓ of T by $\lambda > 0$ and simultaneously dividing row ℓ of U by λ leaves the product TU unchanged, we can always scale the rows and columns of two matrices so that they are normalized without changing TU .

Lemma 4 (Rothvoß [14]) *If the pair T, U is normalized, then $\max\{\|T\|_\infty, \|U\|_\infty\} \leq \sqrt{\|TU\|_\infty}$.*

Proof Let $S := TU$. Suppose, for the sake of contradiction, that the assertion does not hold. Without loss of generality, we may assume that $\|T\|_\infty > \sqrt{\|TU\|_\infty}$. Thus $T_{i\ell} > \sqrt{\|TU\|_\infty}$ for some indices i and ℓ . Since T, U is normalized, $\|U_\ell\|_\infty = \|T^\ell\|_\infty > \sqrt{\|TU\|_\infty}$ and there must be an index j such that $U_{\ell j} > \sqrt{\|TU\|_\infty}$. Then $S_{ij} \geq T_{i\ell}U_{\ell j} > \|TU\|_\infty$, which is a contradiction. \square

Consider a set of n convex independent points V in the plane lying on an integer grid of size polynomial in n , its convex hull $P := \text{conv}(V)$, and $X := \mathbb{Z}^2 \cap P$. The next crucial lemma (adapted from a similar result in [14]) implies that the description of an extension $Q := \{(x, y) \mid Ax + Ty = b, y \geq 0\}$ for P —potentially containing irrational numbers—can be rounded such that an integer point x is in X if and only if there is a $y \geq 0$ such that $\bar{A}x + \bar{T}y \approx \bar{b}$ holds for the rounded system. Moreover, all coefficients in the rounded system come from a domain which is bounded by a polynomial in n .

Lemma 5 *For $d, N \geq 2$ let $V = \{v_1, \dots, v_n\} \subseteq \mathbb{Z}^d$ be a convex independent and non-empty set of points with $\|v_i\|_\infty \leq N$ for $i \in [n]$. Let $P := \text{conv}(V)$ and let $X := P \cap \mathbb{Z}^d$. Denote $r := \text{xc}(P)$ and $\Delta := ((d + 1)N)^d$. Then there are*

matrices $\bar{A} \in \mathbb{Z}^{(d+r) \times d}$, $\bar{T} \in (\frac{1}{4r(d+r)\Delta} \mathbb{Z}_+)^{(d+r) \times r}$ and a vector $\bar{b} \in \mathbb{Z}^{d+r}$ with $\|\bar{A}\|_\infty, \|\bar{b}\|_\infty, \|\bar{T}\|_\infty \leq \Delta$ such that

$$X = \left\{ x \in \mathbb{Z}^d \mid \exists y \in [0, \Delta]^r : \|\bar{A}x + \bar{T}y - \bar{b}\|_\infty \leq \frac{1}{4(d+r)} \right\}.$$

Proof Let $Ax \leq b$ be a non-redundant description of P with integral coefficients. We may assume (see, e.g., [8, Lemma D.4.1]) that $\|A\|_\infty, \|b\|_\infty \leq \Delta = ((d+1)N)^d$. Since $\text{xc}(P) = r$, by Yannakakis’ Theorem 1 there exist matrices $T \in \mathbb{R}_+^{m \times r}$ and $U \in \mathbb{R}_+^{r \times n}$ such that $S := TU$ is the slack matrix of P , and $P = \{x \in \mathbb{R}^d \mid \exists y \in \mathbb{R}^r : Ax + Ty = b, y \geq 0\}$. Without loss of generality assume that the pair T, U is normalized. Note that

$$\|S\|_\infty = \max_{\substack{i \in [m] \\ j \in [r]}} (b_i - A_i v_j) \leq \Delta + dN\Delta \leq \Delta^2.$$

Since T, U are normalized, using Lemma 4, we have that $\|T\|_\infty \leq \Delta$ and $\|U\|_\infty \leq \Delta$.

Let $W := \text{span}(\{(A_i, T_i) \mid i \in [m]\})$ be the row span of the constraint matrix of the system $Ax + Ty = b$ and let $k := \dim(W)$ be the dimension of W . Choose $I \subseteq \{1, \dots, m\}$ of size $|I| = k$ such that the volume of the parallelepiped spanned by the vectors $\{(A_i, T_i) \mid i \in I\}$, denoted by $\text{vol}(\{(A_i, T_i) \mid i \in I\})$, is maximized. Let T'_I be the matrix obtained from rounding the coefficients of T_I to the nearest multiple of $\frac{1}{4r(d+r)\Delta}$. Our choice will be $\bar{A} := A_I, \bar{T} := T'_I$ and $\bar{b} := b_I$. Let

$$Y := \left\{ x \in \mathbb{Z}^d \mid \exists y \in [0, \Delta]^r : \|A_I x + T'_I y - b_I\|_\infty \leq \frac{1}{4(d+r)} \right\}.$$

Then it is sufficient to show that $X = Y$.

Claim 6 $X \subseteq Y$.

Proof of claim Consider an arbitrary vertex $v_j \in V$. Since, $S = TU$, we can choose $y := U^j \geq 0$ such that $Av_j + Ty = b$. Since T, U are normalized, we have that $\|y\|_\infty \leq \|U\|_\infty \leq \Delta$. Note that $\|T - T'\|_\infty \leq \frac{1}{4r(d+r)\Delta}$. By the triangle inequality,

$$\begin{aligned} \|A_I v_j + T'_I y - b_I\|_\infty &\leq \underbrace{\|A_I v_j + T_I y - b_I + (T'_I - T_I)y\|_\infty}_{=0} \\ &\leq r \cdot \underbrace{\|T'_I - T_I\|_\infty}_{\leq \frac{1}{4r(d+r)\Delta}} \cdot \underbrace{\|y\|_\infty}_{\leq \Delta} \leq \frac{1}{4(d+r)}. \end{aligned}$$

Thus $v_j \in Y$ and hence $V \subseteq Y$. It follows that $X \subseteq Y$. □

Claim 7 $X \supseteq Y$.

Proof of claim We show that $x \in \mathbb{Z}^d \setminus X$ implies $x \notin Y$. Since $x \notin X$ and $X \subseteq P$, there must be a row ℓ with $A_\ell x > b_\ell$. Since A, b and x are integral, one even has $A_\ell x \geq b_\ell + 1$. Note that in general ℓ is not among the selected constraints with row indices in I . But there are unique coefficients $\lambda \in \mathbb{R}^k$ such that we can express constraint $A_\ell x + T_\ell y = b_\ell$ as a linear combination of those with indices in I , i.e.

$$(A_\ell, T_\ell) = \sum_{i \in I} \lambda_i (A_i, T_i).$$

It is easy to see that $\sum_{i \in I} \lambda_i b_i = b_\ell$, since otherwise the system $Ax + Ty = b$ could not have any solution (x, y) at all and $P = \emptyset$. The next step is to bound the coefficients λ_i . Here we recall that by Cramer’s rule,

$$|\lambda_i| = \frac{\text{vol}(\{(A_{i'}, T_{i'}) \mid i' \in I \setminus \{i\} \cup \{\ell\}\})}{\text{vol}(\{(A_{i'}, T_{i'}) \mid i' \in I\})} \leq 1,$$

since we picked I such that $\text{vol}(\{(A_{i'}, T_{i'}) \mid i' \in I\})$ is maximized. Fix an arbitrary $y \in [0, \Delta]^r$, then

$$\begin{aligned} 1 \leq \underbrace{|A_\ell x - b_\ell|}_{\geq 1} + \underbrace{T_\ell y}_{\geq 0} &= \left| \sum_{i \in I} \lambda_i (A_i x - b_i + T_i y) \right| \\ &\leq \sum_{i \in I} \underbrace{|\lambda_i|}_{\leq 1} \cdot |A_i x - b_i + T_i y| \\ &\leq (d + r) \cdot \|A_I x - b_I + T_I y\|_\infty \end{aligned} \tag{1}$$

using the triangle inequality and the fact that $|I| \leq d + r$. Again making use of the triangle inequality yields

$$\begin{aligned} \|A_I x - b_I + T_I y\|_\infty &= \|A_I x - b_I + T'_I y + (T_I - T'_I) y\|_\infty \\ &\leq \|A_I x - b_I + T'_I y\|_\infty + r \cdot \underbrace{\|T_I - T'_I\|_\infty}_{\leq \frac{1}{4r(d+r)\Delta}} \cdot \underbrace{\|y\|_\infty}_{\leq \Delta} \\ &\leq \|A_I x - b_I + T'_I y\|_\infty + \frac{1}{4(d+r)}. \end{aligned} \tag{2}$$

Combining (1) and (2) gives $\|A_I x - b_I + T'_I y\|_\infty \geq \frac{1}{d+r} - \frac{1}{4(d+r)} > \frac{1}{4(d+r)}$ for all $y \in [0, \Delta]^r$ and consequently $x \notin Y$. □

The theorem follows. Note that by padding zeros, we can ensure that \bar{A}, \bar{T} and \bar{b} have exactly $d + r$ rows. □

Now we are ready to prove our lower bound for the extension complexity of polygons.

Theorem 8 *For every $n \geq 3$, there exists a convex n -gon P with vertices in $[2n] \times [4n^2]$ and $\text{xc}(P) = \Omega(\sqrt{n}/\sqrt{\log n})$.*

Proof The $2n$ points of the set $Z := \{(z, z^2) \mid z \in [2n]\}$ are obviously convex independent. In other words, every subset $X \subseteq Z$ of size $|X| = n$ yields a different convex n -gon. The number of such n -gons is $\binom{2n}{n} \geq 2^n$. Let $R := \max\{\text{xc}(\text{conv}(X)) \mid X \subseteq Z, |X| = n\}$. Lemma 5 provides a map Φ which takes X as input and provides the rounded system $(\bar{A}, \bar{T}, \bar{b})$. (If the choice of A, b and I is not unique, make an arbitrary canonical choice.) By padding zeros, we may assume that this system is of size $(2 + R) \times (3 + R)$.

Also, Lemma 5 guarantees that for each system $(\bar{A}, \bar{T}, \bar{b})$, the corresponding set X can be reconstructed. In other words, the map Φ must be injective and the number of such system must be at least 2^n . Thus it suffices to determine the number of such systems: the entries in each system $(\bar{A}, \bar{T}, \bar{b})$ are integer multiples of $\frac{1}{4r(d+r)\Delta} = \frac{1}{4r(2+r)144n^4}$ for some $r \in [R]$ using $d = 2, N = 4n^2, \Delta = (12n^2)^2 = 144n^4$. Since no entry exceeds Δ , for each entry there are at most $1 + \sum_{r=1}^R (165888r(2+r)n^8) \leq cn^{11}$ many possible choices for some fixed constant c (note that $R \leq n$). Thus the number of such systems is bounded by $(cn^{11})^{(3+R) \cdot (2+R)} \leq 2^{c' \log n \cdot R^2}$ for some constant c' .

We conclude that $2^{c' \log_2 n \cdot R^2} \geq 2^n$ and thus $R = \Omega(\sqrt{n}/\sqrt{\log n})$. □

6 Concluding Remarks

Although the two lower bounds presented here on the worst case extension complexity of a n -gon are $\tilde{\Omega}(\sqrt{n})$, it is plausible that the true answer is $\tilde{\Omega}(n)$. We leave this as an open problem.

Acknowledgements We thank Stefan Langerman for suggesting the proof of Theorem 3. We also thank Volker Kaibel and Sebastian Pokutta for stimulating discussions. Finally, we thank the anonymous referee for his comments which helped improving the text. S. Fiorini supported by the *Actions de Recherche Concertées* (ARC) fund of the *Communauté française de Belgique*. T. Rothvoß supported by Feodor Lynen Fellowship of the Alexander von Humboldt Foundation, ONR Grant N00014-11-1-0053 and NSF Contract CCF-0829878. H.R. Tiwary supported by *Fonds National de la Recherche Scientifique* (F.R.S.–FNRS).

References

1. Ajtai, M., Komlós, J., Szemerédi, E.: An $O(n \log n)$ sorting network. In: Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, STOC'83, pp. 1–9. ACM, New York (1983)
2. Ben-Tal, A., Nemirovski, A.: On polyhedral approximations of the second-order cone. *Math. Oper. Res.* **26**(2), 193–205 (2001)
3. Cohen, J.E., Rothblum, U.G.: Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra Appl.* **190**, 149–168 (1993)
4. Conforti, M., Cornuéjols, G., Zambelli, G.: Extended formulations in combinatorial optimization. *4OR* **8**(1), 1–48 (2010)
5. Conforti, M., Faenza, Y., Fiorini, S., Grappe, R., Tiwary, H.R.: Extended formulations, non-negative factorizations and randomized communication protocols. <http://arxiv.org/abs/1105.4127> (2011)
6. Fiorini, S., Kaibel, V., Pashkovich, K., Theis, D.O.: Combinatorial bounds on nonnegative rank and extended formulations. Working paper (2011)
7. Goemans, M.: Smallest compact formulation for the permutahedron. <http://math.mit.edu/~goemans/PAPERS/permutahedron.pdf> (2009)

8. Hindry, M., Silverman, J.H.: *Diophantine Geometry: An Introduction*, 1st edn. Springer, Berlin (2000)
9. Hungerford, T.W.: *Algebra*. Graduate Texts in Mathematics. Springer, New York (1974)
10. Kaibel, V.: Extended formulations in combinatorial optimization. *Optima* **85**, 2–7 (2011)
11. Kaibel, V., Pashkovich, K.: Constructing extended formulations from reflection relations. In: *Proceedings of the 15th Conference on Integer Programming and Combinatorial Optimization (2011, to appear)*
12. Lang, S.: *Algebra*, Graduate Texts in Mathematics. Springer, Berlin (2002)
13. Martin, R.K.: Using separation algorithms to generate mixed integer model reformulations. *Oper. Res. Lett.* **10**(3), 119–128 (1991)
14. Rothvoß, T.: Some 0/1 polytopes need exponential size extended formulations. <http://arxiv.org/abs/1105.0036> (2011)
15. Stewart, I.: *Galois Theory*, 3rd edn. Chapman & Hall/CRC Mathematics. Chapman & Hall/CRC, Boca Raton (2004)
16. Vanderbeck, F., Wolsey, L.A.: Reformulation and decomposition of integer programs. In: Jünger, M., et al. (eds.) *50 Years of Integer Programming 1958–2008*, pp. 431–502. Springer, Berlin (2010)
17. Yannakakis, M.: Expressing combinatorial optimization problems by linear programs. *J. Comput. Syst. Sci.* **43**(3), 441–466 (1991)