

Monotone paths in planar convex subdivisions and polytopes*

Adrian Dumitrescu[†] Günter Rote[‡] Csaba D. Tóth[§]

November 18, 2012

Abstract

Consider a connected subdivision of the plane into n convex faces where every vertex is incident to at most Δ edges. Then, starting from every vertex there is a path with at least $\Omega(\log_{\Delta} n)$ edges that is monotone in some direction. This bound is the best possible. Consider now a connected subdivision of the plane into n convex faces where exactly k faces are unbounded. Then, there is a path with at least $\Omega(\log(n/k)/\log \log(n/k))$ edges that is monotone in some direction. This bound is also the best possible. Our methods are constructive and lead to efficient algorithms for computing monotone paths of lengths specified above.

In 3-space, we show that for every $n \geq 4$, there exists a polytope P with n vertices, bounded vertex degrees, and triangular faces such that every monotone path on the 1-skeleton of P has at most $O(\log^2 n)$ edges. We also construct a polytope Q with n vertices, and triangular faces, (with unbounded degree however), such that every monotone path on the 1-skeleton of Q has at most $O(\log n)$ edges.

Keywords: Monotone path, convex subdivision, geometric graph, convex polytope.

Mathematics Subject Classification: 52B10, 52C45.

1 Introduction

A *geometric graph* $G = (V, E)$ in Euclidean d -space is a set V of distinct points (vertices) in Euclidean d -space \mathbb{R}^d , and a set E of line segments (edges) between vertices such that no vertex lies in the relative interior of any edge. For our investigation, it is convenient to define an *extended geometric graph* $G = (V, E)$, where E may also contain rays, each emitted by a vertex, and lines (disjoint from vertices). A directed path p in an extended geometric graph G is *monotone* (resp., *weakly monotone*) if there exists a unit vector \vec{u} such that the inner product $\vec{e} \cdot \vec{u}$ is positive (non-negative) for every directed edge \vec{e} of p . In \mathbb{R}^2 , in particular, the direction of a unit vector $\vec{u} = (\cos \theta, \sin \theta)$ is determined by the angle $\theta \in (-\pi, \pi]$. A directed path p is *x -monotone* (resp., *y -monotone*) if it is monotone in direction 0 (resp., $\frac{\pi}{2}$). The *size* (or *length*) of a path is the number

*Preliminary results were reported by the authors in [6, 14, 7].

[†]Department of Computer Science, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA. Email: dumitres@uwm.edu. Supported in part by NSF grant DMS-1001667.

[‡]Institut für Informatik, Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany. Email: rote@inf.fu-berlin.de. Supported in part by CIB (Centre Interfacultaire Bernoulli) in Lausanne, by the NSF (National Science Foundation), and by the ESF EUROCORES programme EuroGIGA-VORONOI, Deutsche Forschungsgemeinschaft (DFG): RO 2338/5-1. This work was initiated while the author was a guest of CIB in Fall 2010.

[§]Department of Mathematics and Statistics, University of Calgary, Calgary, AB, Canada, T2N 1N4. E-mail: cdtoth@ucalgary.ca. Supported in part by NSERC grant RGPIN 35586. Research by this author was conducted at the Fields Institute for Research in Mathematical Sciences, Toronto, ON, Canada M5T 3J1.

of edges in the path, or equivalently, one plus the number of vertices on the path. Notice that any path (monotone or not) in an extended geometric graph contains at most two rays.

We study monotone paths in the 1-skeletons of polytopes and convex subdivisions. The *1-skeleton* $G(P)$ of a bounded polytope P in \mathbb{R}^d is the geometric graph formed by the vertices and edges of P . Similarly, the *1-skeleton* $G(P)$ of an unbounded polytope P is the extended geometric graph formed by the vertices and edges of P .

A *convex subdivision* (for short, *subdivision*) of \mathbb{R}^d is a set Π of (bounded or unbounded) convex polytopes (called *faces*) that tile \mathbb{R}^d . The 1-skeleton of a subdivision Π of the plane \mathbb{R}^2 is the extended geometric graph $G(\Pi)$ whose vertices are the points incident to 3 or more edges, and whose edges are the line segments, rays, and lines lying on the common boundary of two faces. To exclude some trivial cases, we always consider convex subdivisions whose 1-skeleton is connected, referred to as *connected subdivisions* for short. We are looking for long monotone paths in the 1-skeletons of polytopes and convex subdivisions of the plane. Our results are the following.

Theorem 1. *Let Π be a connected subdivision of the plane into n convex cells in which every vertex is incident to at most Δ edges. Then, for every vertex v , there is a weakly monotone path with at least $c \log_{\Delta} n$ edges starting from v , where $c > 0$ is an absolute constant. Apart from the constant c , this bound is the best possible.*

Theorem 2. *Let Π be a connected subdivision of the plane into n convex cells, k of which are unbounded with $n > k \geq 3$. Then $G(\Pi)$ contains a monotone path with at least $c \log \frac{n}{k} / \log \log \frac{n}{k}$ edges, where $c > 0$ is an absolute constant. Apart from the constant c , this bound is the best possible.*

We also consider long monotone paths in the 1-skeleton of a convex polytope in 3-space. We present two constructions, one with bounded vertex degrees and one with arbitrary degrees.

Theorem 3. *For every $n \geq 4$, there is a polytope P in \mathbb{R}^3 with n vertices, bounded vertex degrees, and triangular faces such that every monotone path in $G(P)$ has at most $O(\log^2 n)$ edges.*

Theorem 4. *For every $n \geq 4$, there is a polytope Q in \mathbb{R}^3 with n vertices and triangular faces such that every monotone path in $G(Q)$ has at most $O(\log n)$ edges.*

We do not know whether the bounds in Theorems 3 and 4 are asymptotically tight. The diameter of a bounded degree graph on n vertices is $\Omega(\log n)$. It follows that every monotone path connecting a diametral pair of vertices of a polytope with n vertices of bounded degree has $\Omega(\log n)$ edges. If the maximum vertex degree of the polytope is not bounded, then a lower bound of $\Omega(\log n / \log \log n)$ follows from the result of Chazelle et al. [4] (see below), applied to the dual graph of a plane projection, using reciprocal diagrams and the Maxwell-Cremona correspondence.

Related work. It is well known that the classical simplex algorithm in linear programming produces a monotone path on the 1-skeleton of a d -dimensional polytope of feasible solutions; it is called a *parametric simplex path*. According to the old *monotone Hirsch conjecture* [18], for any vector \vec{u} , the 1-skeleton of every d -dimensional polytope with n facets contains a \vec{u} -monotone path of at most $n - d$ edges from any vertex to a \vec{u} -maximal vertex. For the weakly monotone version, counterexamples have already been found by Todd [16] in the 1980s. Recent counterexamples for this conjecture found by Santos [15] show that the monotone variant is also false. It is not known whether the Hirsch conjecture can be relaxed so that it holds when the length $n - d$ is replaced by some appropriate polynomial in d and n .

The *monotone upper bound problem* posed by Klee [8] asks for the maximum number of edges $M(d, n)$ of a monotone path on the 1-skeleton of a polytope with $n \geq d + 1$ facets in \mathbb{R}^d . An obvious

upper bound for $M(d, n)$ (within 1) is the maximum number of vertices such a polytope can have, as given by the Upper Bound Theorem [9]. Pfeifle [10] and Pfeifle and Ziegler [11] have shown that $M(d, n)$ always attains this bound for $1 \leq d \leq 4$, but it does not always reach it for $d = 6$.

Balogh et al. [3] showed that there is a convex subdivision Π_n generated by n lines in the plane with $O(n^2)$ faces such that $G(\Pi_n)$ contains a monotone path of length $\Omega(n^2/c^{\sqrt{n}})$ for some constant $c > 1$. In such a monotone path, many consecutive edges are collinear. Dumitrescu [5] proved that every simple arrangement of n lines admits a monotone path of length at least n in the convex subdivision generated by n lines. On the other hand, for every $n \geq 2$, there exists a line arrangement in which no monotone path is longer than $3n/2 + O(\log n)$. Again, consecutive collinear edges are allowed in such a path. If these are disallowed, the lower bound is unchanged but the upper bound drops to $4n/3 + O(\log n)$. As regarding algorithms for monotone paths, Arkin et al. [1] devised a polynomial-time algorithm that computes an x -monotone path between two given points in the plane in the presence of polygonal obstacles, or reports that none exists.

A *monotone face sequence* in a convex subdivision Π is a sequence of faces such that there is a direction \vec{u} such that any two consecutive faces, f_1 and f_2 , are adjacent and a vector of direction \vec{u} crosses their common boundary from f_1 to f_2 . Chazelle et al. [4] showed that in a subdivision of the plane into n convex faces in which every face is adjacent to at most d other faces, there is a monotone face sequence of length

$$\Omega(\log_d n + \log n / \log \log n), \tag{1}$$

and this bound is tight. Moreover, a monotone face sequence of this length can even be achieved by faces stabbed by a line. The latter result was generalized to d -dimensions by Tóth [17]: for every subdivision of \mathbb{R}^d into n convex faces, there is a line that stabs $\Omega((\log n / \log \log n)^{\frac{1}{d-1}})$ faces, and this bound is the best possible.

Motivation. We were led to the above result of Chazelle et al. [4] by the complexity analysis of an algorithm for partial matching between two finite planar point sets under translations [13]. There, one could show that a certain subdivision Q contained no monotone face sequence longer than some polynomial in the given parameters. If the lower bound $\Omega(\log_d n + \log n / \log \log n)$ could be strengthened to a polynomial bound of the form $\Omega(n^\alpha)$ for some $\alpha > 0$, this would have implied a polynomial bound on the number of faces of Q .

If P is the vertical projection of a piecewise linear convex terrain \widehat{P} in \mathbb{R}^3 , one can apply a polarity transform to this terrain with respect to the paraboloid $z = x^2 + y^2$, yielding another piecewise linear convex terrain \widehat{P}^* . The projection P^* of this terrain is a *reciprocal diagram* of P (cf. [2]): its graph is dual to the graph of P , in the sense that vertices of P^* correspond to faces of P and vice-versa. Each edge in P^* has a corresponding edge in P , and moreover, these two edges are perpendicular. (This last property distinguishes a reciprocal diagram from a general drawing of the dual graph of P .) A monotone path in P becomes a monotone face sequence in P^* .

Thus, for subdivisions P that are projections of convex terrains, the question about long monotone paths is completely answered by applying the tight bound of (1) for $Q = P^*$. However, for general subdivisions, the problems are not directly related, and in fact, the answers are different: We will see in Section 2 that one cannot add any other term that grows to infinity (such as the term $\Omega(\log n / \log \log n)$ in (1)) to the lower bound in Theorem 1.

Outline. We start with the proof of Theorem 1 in Section 2. We study convex subdivisions of simple polygons in Section 3.1. The tools developed there are instrumental in the proof of

Theorem 2 in Section 3.2. The proofs of Theorems 3 and 4 regarding monotone paths on polytopes appear in Sections 4 and 5, respectively. We conclude with some open problems in Section 6.

2 Proof of Theorem 1

Lower bound. The lower bound in Theorem 1 follows from the following lemma in a straightforward way (by counting, or inductively).

Lemma 1. *Let v be a vertex in a connected convex planar subdivision Π . Then $G(\Pi)$ contains a spanning tree rooted at v such that all paths starting at v are weakly monotone.*

Proof. For a generic direction \vec{u} , we define the *rightmost path* $R(\vec{u})$ starting at v as follows; see Fig. 1(left): start at v and always follow the rightmost outgoing edge that is weakly monotone in direction \vec{u} until we arrive at an unbounded ray.

Now we start rotating \vec{u} clockwise. At some direction \vec{u}' , $R(\vec{u}')$ will be different from $R(\vec{u})$. At that point, $R(\vec{u})$ is still weakly monotone in direction \vec{u}' . Now, any vertex w (and any edge) in the region between $R(\vec{u})$ and $R(\vec{u}')$ can be reached by a weakly monotone path in direction \vec{u}' . Indeed, simply start at w and go monotonically in the direction $-\vec{u}'$ until reaching $R(\vec{u})$ or $R(\vec{u}')$. From there, follow $R(\vec{u})$ or $R(\vec{u}')$ to v . In this way, we can form a spanning tree of all vertices between $R(\vec{u})$ and $R(\vec{u}')$ with the desired properties.

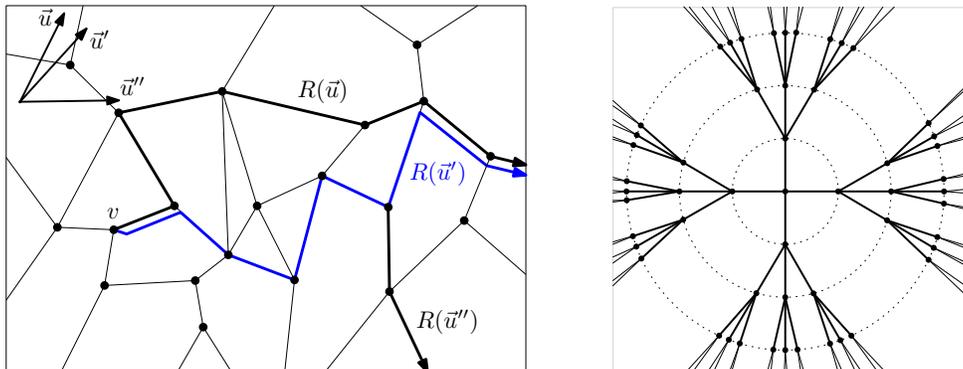


Figure 1: Left: The rightmost path $R(\vec{u})$ starting from vertex v in direction \vec{u} . Right: An balanced tree of maximum degree 4, which is the extended graph $G(\Pi)$ of a convex subdivision Π .

Continuing the rotation in this way, we eventually reach all vertices and all infinite rays. \square

The subdivision in Fig. 2 shows that the lemma does not hold with (strictly) monotone paths. However, if there are no angles of 180° , the statement extends to strictly monotone paths.

Let v be a vertex in a connected convex planar subdivision Π . By Lemma 1, $G(\Pi)$ contains a spanning tree rooted at v such that all paths starting at v are weakly monotone. The maximum degree in $G(\Pi)$ is at most Δ . Hence the spanning tree contains a path of size $\Omega(\log_\Delta n)$ from v to some vertex of $G(\Pi)$.

Upper bound. If the maximum degree Δ is $\Theta(n)$, Theorem 1 gives only a trivial statement. Dividing the plane into n convex sectors, by n rays starting from the origin shows that, indeed, there is no non-constant lower bound on the length of monotone paths in this case.

This construction generalizes for any $\Delta \in \mathbb{N}$, and $n \geq 3$. We construct a convex subdivision Π such that $G = G(\Pi)$ is a tree with maximum degree Δ and diameter $O(\log_\Delta n)$. See Fig. 1(right).

Recall that a y -monotone polygon is bounded from the left and from the right by two y -monotone directed paths. The common start (resp., end) point of the two boundary paths is called the *bottom* (resp., *top*) vertex of P .

Observation 2. *Let $\Pi = \Pi(P)$ be a convex subdivision of a y -monotone polygon P with no horizontal edges. For every vertex v of $G(\Pi)$, there is a y -monotone path from the bottom vertex of P to the top vertex of P which is incident to v .*

Proof. The claim is immediate if v is on the boundary of P . Assume that v is in the interior of P . By Observation 1, $G(\Pi)$ contains monotone paths from v to the boundary of P in the two opposite directions $\pi/2$ and $-\pi/2$. The combination of these paths (together with parts of the boundary of P , if needed) gives the required y -monotone path through v from the bottom to the top vertex of P . \square

Criterion for y -monotone polygons. To prove the lower bound in Theorem 2, we constructively build a monotone path of the required length for any given convex partition. In our recursive construction, we successively subdivide a y -monotone polygon P into smaller polygons along certain y -monotone paths in $G(\Pi(P))$. Lemma 2 below provides a criterion for producing y -monotone pieces. A vertex v in a geometric graph G is called *y -maximal* (resp., *y -minimal*) if all edges of G incident to v lie in a closed halfplane below (resp., above) the horizontal line passing through v . A vertex v is *y -extremal* if it is y -maximal or y -minimal. It is clear that the boundary of a y -monotone polygon has exactly two y -extremal vertices, namely its top vertex and its bottom vertex.

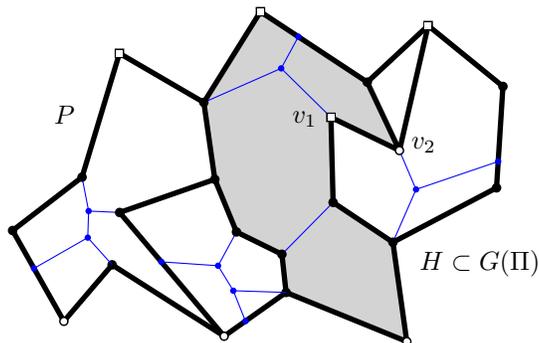


Figure 3: The graph $G(\Pi)$ of a convex partition Π of a simple polygon P . A subgraph H in bold contains all edges of P . The y -maximal (resp., y -minimal) vertices of H are marked with empty squares (resp., empty circles). Vertices v_1 and v_2 are y -extremal in H , but not convex vertices of P (here they are not even vertices of P).

Lemma 2. *Let P be a simple polygon with a convex subdivision Π such that no edge in $G = G(\Pi)$ is horizontal. Let H be a subgraph of G that contains all edges and vertices of P . Then all bounded faces of H are y -monotone polygons if and only if all y -extremal vertices of H are convex vertices of P .*

Proof. Assume that all bounded faces of H are y -monotone polygons. Suppose that H has a y -extremal vertex v . We may assume without loss of generality that all edges of H incident to v are in the halfplane above v . Let f_v be the face of H incident to v that lies directly below v . Face f_v has a reflex interior angle at v , and v is neither the top nor the bottom vertex of f_v . Hence f_v is not y -monotone, and so it has to be an unbounded face of H . Since H contains all boundary edges of P , the face f_v is the unbounded face of G , as well. It follows that v is a convex vertex of P .

Assume that all y -extremal vertices of H are convex vertices of P . Consider a bounded face f of H . Let p_1 and p_2 be edge-disjoint directed paths on the boundary of f from a bottom (lowest) vertex of f to a top (highest) vertex of f (ties are broken arbitrarily), such that f lies on the right side of p_1 and on the left side of p_2 . Suppose for contradiction that f is not y -monotone. We may assume without loss of generality that p_2 is not y -monotone, as in Fig. 3 (where f is the shaded face). The first (resp., last) edge of p_2 has a positive inner product with $(0, 1)$ by construction. There are two consecutive edges \vec{e}_1 and \vec{e}_2 in p_2 such that $\vec{e}_1 \cdot (0, 1) > 0 > \vec{e}_2 \cdot (0, 1)$, since otherwise p_2 would be y -monotone. Let v be the common vertex of these two edges. Since f is on the left side of both \vec{e}_1 and \vec{e}_2 , these two edges are consecutive in the counterclockwise rotation order of the edges of H incident to v . Hence, H has no edge incident to v in the halfplane above v , and so v is y -extremal (y -maximal) in H . However, v is a reflex vertex of face f , hence it cannot be a convex vertex of P (specifically, v is either interior to P , or a reflex vertex of P). This contradicts our initial assumption and completes the proof. \square

Subdividing a polygon into y -monotone pieces. Our upper bound relies on the following two lemmas. In Lemma 3, we partition the bounded faces of a convex subdivision of the plane into monotone polygons. In Lemma 4, we subdivide a y -monotone polygon P into smaller y -monotone polygons which are not incident to both the top and the bottom vertex of P .

Lemma 3. *Let Π be a subdivision of the plane into n convex faces, $k \geq 3$ of which are unbounded. Then there is a subset $\Pi' \subset \Pi$ of at least $(n - k)/(k - 2)$ faces such that Π' is the convex subdivision of a monotone polygon.*

Proof. Let H denote the graph of vertices and bounded edges of $G = G(\Pi)$, without rays and lines. We proceed by induction on the number of 2-connected components, i.e., blocks of H . In the base case, H is 2-connected. Then the union of the closures of all bounded faces in Π forms a simple polygon P . Suppose that no edge of G is horizontal or vertical. Let V_0 be the set of y -extremal reflex vertices of P . If V_0 is empty, then P is a y -monotone polygon, and the $n - k$ bounded faces form a convex subdivision of P . If V_0 is nonempty, then we construct a y -monotone path $\gamma(v)$ for each $v \in V_0$, in an arbitrary order, as follows. If $v \in V_0$ is y -maximal (resp., y -minimal), then construct $\gamma(v)$ starting from v by successively appending edges in direction $\pi/2$ (resp., $-\pi/2$) until the path reaches another vertex on the boundary of P or a previously constructed path $\gamma(v')$, $v' \in V_0$. The paths $\gamma(v)$ subdivide P into $|V_0| + 1$ simple polygons, each of which is y -monotone by Lemma 2.

It remains to show that $|V_0| \leq k - 3$ after an appropriate rigid motion. Notice that every y -extremal reflex vertex of P is a y -extremal vertex of some unbounded face. An unbounded face cannot have both a top and a bottom vertex. Two of the unbounded faces, namely those containing rays in directions $(1, 0)$ and $(-1, 0)$, have neither a top nor a bottom vertex. This already implies $|V_0| \leq k - 2$. Let \vec{e}_0 be a ray edge of G emitted by an extremal vertex of the convex hull of P . Assume, by applying a reflection if necessary, that the unbounded face on the left of \vec{e}_0 is not a halfplane. Rotate the subdivision such that no edge in G is horizontal and edge \vec{e}_0 has the smallest positive slope. Now the unbounded face above \vec{e}_0 has a bottom vertex on the convex hull of P , which is not a reflex vertex of P . Therefore, we have $|V_0| \leq k - 3$.

Assume now that H has several 2-connected components. We distinguish two cases. Case 1: H is disconnected. Then G is also disconnected, and it is the disjoint union of G_1 and G_2 , which are incident to k_1 and k_2 unbounded faces, respectively. Then there is at most one face (a parallel strip) incident to both G_1 and G_2 . Hence $k_1 + k_2 \leq k + 1$, and so $(k_1 - 2) + (k_2 - 2) < k - 2$. Induction completes the proof as follows: Let G_1 and G_2 have \bar{n}_1 and \bar{n}_2 bounded faces, with

$\bar{n}_1 + \bar{n}_2 = \bar{n} = n - k$. Then we have $\max(\frac{\bar{n}_1}{k_1-2}, \frac{\bar{n}_2}{k_2-2}) \geq \frac{\bar{n}}{k-2}$, and hence it is sufficient to find the required subdivision in G_1 or G_2 . Case 2: H is connected but has a cut vertex v . Then G decomposes into subgraphs G_1, \dots, G_ℓ , for some $\ell \geq 2$, whose only common vertex is v . G_i can have a reflex angle incident to v ; thus we augment it to a convex subdivision \hat{G}_i by adding a ray incident to v . Denote by k_i the number of unbounded faces in \hat{G}_i for $i = 1, \dots, \ell$. G has ℓ unbounded faces between consecutive subgraphs around v , which are not faces in some \hat{G}_i . On the other hand, each \hat{G}_i has two unbounded faces that are not faces of G . Hence $\sum_{i=1}^{\ell} (k_i - 2) = k - \ell \leq k - 2$. Induction completes the proof as above. \square

Lemma 4. *Let P be a y -monotone polygon with bottom vertex s , top vertex t , and a convex subdivision Π . Let m be the maximum size of a monotone path in $G = G(\Pi)$. Then at most two faces in Π are incident to both s and t ; the remaining faces can be partitioned into at most $m + 1$ sets, each of which is the convex subdivision of a y -monotone polygon whose top or bottom vertex is not in $\{s, t\}$.*

Proof. If a face $f \in \Pi$ is incident to both s and t , then by convexity the closure of f contains the line segment st . Thus, at most two faces in Π are incident to both s and t . We first partition the remaining faces into two subsets, lying on opposite sides of a monotone path α such that the faces in each subset are incident to at most one of s and t . Then we further partition each of the two subsets to form y -monotone polygons. Let H_{st}^- and H_{st}^+ denote the closed halfplanes on the left and right of st , respectively. We distinguish two cases.

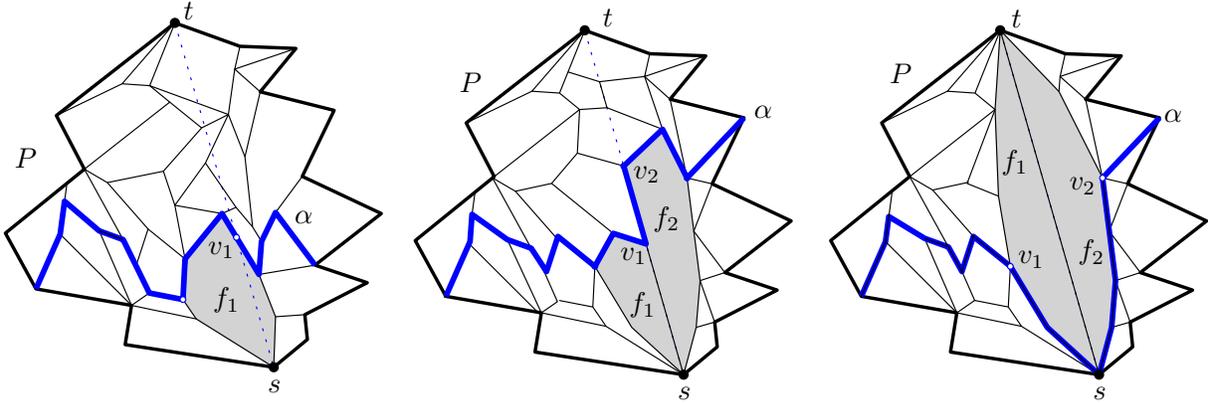


Figure 4: A y -monotone polygon P with a bottom vertex s , a top vertex t , and a convex subdivision. Left: faces f_1 and f_2 are the same, but neither is incident to t . Middle: faces f_1 and f_2 are not incident to t . Right: faces f_1 and f_2 are both incident to t .

Case 1: no face in Π is incident to both s and t . We define two points, v_1 and v_2 , in the relative interior of the segment st . Let $f_1 \in \Pi$ be a face incident to s whose closure contains an initial portion of st , and let the segment sv_1 be the intersection of the closure of f_1 with st . Refer to Fig. 4(left, middle). We may assume, by applying a reflection with respect to the y -axis if necessary, that the interior of f_1 intersects H_{st}^- . Since f_1 is not incident to t , there is an edge in G that contains v_1 and enters the interior of H_{st}^- . If there is some edge in G that contains v_1 and enters the interior of H_{st}^+ or if v_1 is on the boundary of polygon P , then let $v_2 = v_1$. Otherwise, all edges of G incident to v_1 lie in H_{st}^- , hence two consecutive edges are contained in st . Then v_1 is incident to a unique face $f_2 \in \Pi$ on the right of st . Since f_2 is not incident to both s and t , it has a vertex along st which is different from both s and t ; let v_2 be an arbitrary such vertex. In summary: the segment $v_1v_2 \subset st$ is covered by edges of G ; an edge of G contains v_1 and enters the

interior of H_{st}^- ; and v_2 is either on the boundary of P or an edge of G contains v_2 and enters the interior of H_{st}^+ .

Let \vec{u} be a unit vector orthogonal to st . Slightly rotate \vec{u} , if necessary, such that \vec{u} is not parallel or orthogonal to any edge of G and $\vec{u} \cdot \overrightarrow{v_1 v_2} > 0$ if $v_1 \neq v_2$. By Observation 1, there is a $(-\vec{u})$ -monotone path from v_1 to the boundary of P , and a \vec{u} -monotone path from v_2 to the boundary of P . The union of these two paths and the segment $v_1 v_2$ forms a \vec{u} -monotone path, denoted α , between two boundary points of P and passing through v_1 and v_2 . Direct α such that its starting point is in H_{st}^- . Since the endpoints of α are on the boundary of P , every face in Π is either on left or on right side of α . By construction, every face incident to s is on the right of α , and every face incident to t is on the left of α .

By our assumption, α has at most m edges. Construct a path $\gamma(v)$ for every y -extremal interior vertex v of α , in an arbitrary order, as follows. If v is a y -maximal (resp., y -minimal) vertex in α , then construct y -monotone path $\gamma(v)$ starting from v by successively appending edges in direction $\pi/2$ (resp., $-\pi/2$) until the path reaches a vertex on α , the boundary of P or a previously constructed path $\gamma(v')$, $v' \neq v$. These paths $\gamma(v)$ together with α subdivide P into at most $(m-1) + 2 = m + 1$ simple polygons, each of which is y -monotone by Lemma 2.

Case 2: some face in Π is incident to both s and t . Let $f_1 \in \Pi$ be a face incident to both s and t . Refer to Fig. 4(right). We may assume, by applying a reflection with respect to the y -axis if necessary, that the interior of f_1 intersects H_{st}^- . Let \vec{u} be a unit vector orthogonal to st . Slightly rotate \vec{u} , if necessary, such that \vec{u} is not parallel or orthogonal to any edge of G and $\vec{u} \cdot \overrightarrow{st} > 0$. Let v_1 be the \vec{u} -minimal vertex of f_1 . We need to be more careful when defining v_2 . If there is a face f_2 (possibly $f_2 = f_1$) incident to both s and t whose interior intersects H_{st}^+ , then let v_2 be the \vec{u} -maximal vertex of f_2 . Otherwise, if st is an edge of the face f_1 and G has vertices in the relative interior of st , then let v_2 be an arbitrary such vertex. Otherwise, st is an edge on the boundary of P , and then let $v_2 = s$.

Similarly to Case 1, construct a $(-\vec{u})$ -monotone path from v_1 to the boundary of P , and a \vec{u} -monotone path from v_2 to the boundary of P . A \vec{u} -monotone path α is formed by the union of these two paths, a path from v_1 to s along the boundary of f_1 , and a path from s to v_2 along the boundary of f_1 or f_2 . Direct α such that its starting point is in H_{st}^- . Since the endpoints of α are on the boundary of P , every face in Π is either on the left or on the right side of α . By construction, every face incident to s other than f_1 and f_2 (if it exists) is on the right of α , and every face incident to t is on the left of α .

By our assumption, α has at most m edges. Construct a path $\gamma(v)$ for every y -extremal interior vertex of α that precedes v_1 or follows v_2 as in Case 1. If v_1 is an interior vertex of α , then let $\gamma(v_1)$ be the y -monotone path from v_1 to t along the boundary of f_1 . Similarly, if v_2 is an interior vertex of α , then let $\gamma(v_2)$ be the y -monotone path from v_2 to t along the boundary of f_1 or f_2 . Note that some of the interior vertices of α between v_1 and v_2 may be on the boundary of P . If exactly k interior vertices of α are on the boundary of P , then α subdivides P into at most $k + 2$ simple polygons. The paths γ (which are not defined for vertices on the boundary of P) further subdivide these polygons into at most $(m-1) + 2 = m + 1$ simple polygons, each of which is y -monotone by Lemma 2.

By construction, one of these polygons is formed by the faces in Π incident to both s and t . Thus, the faces not incident to both s and t are partitioned into at most m sets, each of which is the convex subdivision of a y -monotone polygon whose top or bottom vertex is not in $\{s, t\}$. \square

In Lemma 4, we have partitioned almost all faces of Π into subsets that each form a y -monotone polygon. In the proof of the lower bound in Theorem 2, we will recurse on only one of these polygons.

Corollary 1. *Let P be a y -monotone polygon with bottom vertex s and top vertex t . Let Π be a subdivision of P into n convex faces. Let m be the maximum size of a monotone path in $G = G(\Pi)$. If $n \geq 3$, then there is a subset $\Pi' \subset \Pi$ of at least $n/(m+3)$ faces such that Π' is the convex subdivision of a y -monotone sub-polygon of P whose top or bottom vertex is not in $\{s, t\}$.*

3.2 Proof of Theorem 2

Lower bound proof. Let $3 \leq k \leq n$ and Π be a subdivision of the plane into n convex faces, where k faces are unbounded. By Lemma 3 there exists a subset $\Pi_0 \subset \Pi$ of $n_0 \geq (n-k)/(k-2)$ faces that form a convex subdivision of a monotone polygon P_0 . We may assume, by applying a rotation if necessary, that no edge of $G(\Pi)$ is horizontal and P_0 is y -monotone. We can assume that $n_0 \geq 4$. It is enough to show that $G(\Pi_0)$ contains a monotone path of size at least $c_0 \log n_0 / \log \log n_0$, where $c_0 > 0$ is an absolute constant.

Let m denote the maximum size of a monotone path in $G(\Pi_0)$. We use Corollary 1 to construct a nested sequence $\Pi_0 \supset \Pi_1 \supset \dots \supset \Pi_t$ such that $|\Pi_i| = n_i \geq n_0/(m+3)^i$ and Π_i is a convex subdivision of a y -monotone polygon P_i for $i = 1, 2, \dots, t$. Moreover, the bottom or top vertex of P_i is different from that of P_{i-1} for $i = 1, 2, \dots, t$. If $|\Pi_i| > 2$, then Π_{i+1} can be constructed from Π_i by Corollary 1. We may therefore assume that $n_t \in \{1, 2\}$.

Let β_t be an arbitrary y -monotone path in $G(\Pi_t)$ between the top and bottom vertex of P_t . For $i = t, t-1, \dots, 1$, the path β_i can be extended to a y -monotone path β_{i-1} between the top and bottom vertex of $G(\Pi_{i-1})$ by Observation 1. Note that β_{i-1} is strictly longer than β_i , since at least one of two endpoints of β_i is not the top or bottom vertex of P_{i-1} . Therefore, β_0 is a y -monotone path with at least $t+1$ edges, and $t \leq m-1$ by the definition of m . We have

$$2 \geq n_t \geq \frac{n_0}{(m+3)^t} \geq \frac{n_0}{(m+3)^{m-1}}.$$

hence $n_0 \leq 2(m+3)^{m-1} = 2^{1+(m-1)\log(m+3)}$. This implies $m \geq c_0 \log n_0 / \log \log n_0$, for some absolute constant $c_0 > 0$, as required.

Upper bound construction. For every pair of integers $k, n \in \mathbb{N}$, where $3 \leq k < n$, we subdivide the plane into a set Π of $\Omega(n)$ convex cells, exactly k of which are unbounded, such that every monotone path in $G(\Pi)$ has $O(\log \frac{n}{k} / \log \log \frac{n}{k})$ edges. We first construct the unbounded faces. If $k = 3$, then let Q_1 be a triangle, and subdivide the exterior of Q_1 into 3 convex faces by 3 rays emitted from the vertices of Q_1 . If $k \geq 4$, then subdivide the plane into k unbounded faces by a star graph with $\lfloor k/2 \rfloor$ leaves, $q_1, \dots, q_{\lfloor k/2 \rfloor}$, and 2 or 3 rays emitted from each leaf (similarly to the subdivision in Fig. 1(right)). Then replace each vertex q_i , $1 \leq i \leq \lfloor k/2 \rfloor$ by a small triangle Q_i . Now it is enough to subdivide each triangle Q_i into at least $n_0 = 5n/k$ bounded faces such that every monotone path restricted to Q_i has $O(\log n_0 / \log \log n_0)$ edges. Since a monotone path can visit at most two triangles Q_i , it has $O(\log n_0 / \log \log n_0)$ edges.

Let $m = 2 \cdot \lceil \log n_0 / \log \log n_0 \rceil$. The basic building block of our construction is a plane geometric graph R shown in Fig. 5(left). The outer face of R is a rhombus symmetric with respect to both the x - and the y -axes. Two opposite corners of the rhombus are connected by an x -monotone zig-zag path ξ of $2m$ edges. The edges of ξ have alternately negative and positive slopes, and so the $2m-1$ interior vertices are alternately y -minimal and y -maximal. Denote by Z the y -minimal interior vertices of ξ , so $|Z| = m$. The y -minimal interior vertices of ξ are joined to the bottom vertex s of the rhombus, and the y -maximal ones to the top vertex t . The vertices s and t are sufficiently far below and respectively far above ξ such that all bounded faces of R are convex. The graph R contains a monotone path of size $2m$ in directions close to horizontal, but every monotone

path has at most 3 edges in directions close to vertical. After an appropriate affine transformation, every monotone path has at most 3 edges in all directions except for those in a prescribed interval of length $\frac{\pi}{2m}$ in $[0, 2\pi)$.

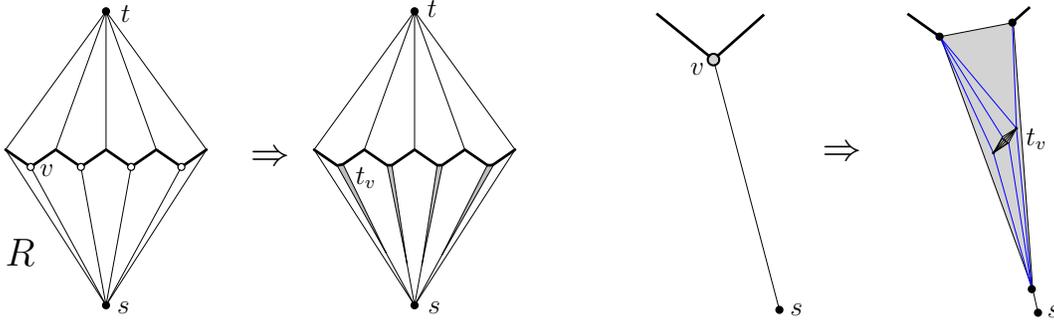


Figure 5: Left: The subgraph R with $m = 4$. The vertices in Z are marked with empty circles. In one phase of our construction, the vertices in Z are replaced by long and skinny triangles. Right: A vertex $v \in Z$ is replaced by a skinny triangle that contains an affine copy of R , and the space between the triangle and R is triangulated.

We construct a subdivision Π of Q in m phases. For $i = 1, \dots, m$, we maintain a convex subdivision Π_i , a set Z_i of special vertices in $G(\Pi_i)$, and a special edge incident to each vertex in Z_i . For constructing Π_1 , consider an affine copy of R , such that any monotone path has at most 3 edges except for directions in the interval $(\frac{\pi}{2m}, \frac{2\pi}{2m})$. This graph, together with four axis-parallel rays from the four corners of the rhombus to the bounding box induce the convex subdivision Π_1 of the plane. Note that $G(\Pi_1) = R$. For each $v \in Z_1 = Z$, let the special edge of v be vs , the edge joining v to the bottom vertex of R .

In phase $i = 2, \dots, m$, we construct Π_i from Π_{i-1} as follows. Replace each vertex $v \in Z_{i-1}$ by a long skinny triangle t_v along the special edge incident to v as in Fig. 5(right). In the interior of t_v , place a small affine copy of R near the midpoint of t_v , such that any monotone path has at most 3 edges in R except for directions in the interval $(\frac{(2i-1)\pi}{2m}, \frac{2i\pi}{2m})$. Denote by r_v the outer boundary of this copy of R . Triangulate the space between r_v and t_v arbitrarily by using $O(1)$ edges; the edges of the triangulation are almost parallel to the special edge sv if t_v is sufficiently skinny and r_v is sufficiently small. Now, any monotone path γ can enter the interior of *at most two* of the m skinny triangles t_v : to see this, note that the direction θ in which γ is monotone is close to perpendicular to at most one skinny triangle, say t_{v_0} . For any other $t_v \neq t_{v_0}$, γ must either enter or leave through the vertex s if it visits the interior. This means that at most two of these triangles can be visited. Moreover, if two of them are visited, then t_{v_0} cannot be entered.

Let Z_i be the union of the vertex sets Z from all affine copies of R created in phase i , and let the special edge of each vertex in Z_i be the edge connecting that vertex to the vertex s of the corresponding copy of R .

We show that $\Pi = \Pi_m$ has $\Omega(n_0)$ faces, and the longest monotone path in $G = G(\Pi)$ has size $O(m) = O(\log n_0 / \log \log n_0)$. Initially, we have $|Z_1| = m$ special vertices. Since $|Z_i| = m|Z_{i-1}|$ for $i = 2, \dots, m$, it follows that $|Z_m| = m^m$. Note that for each special vertex in Z_m , there is an incident quadrilateral face in a copy of R which is not incident to any other vertex in Z_m . Hence Π has at least $|Z_m| = m^m = \Omega(n_0)$ faces.

For any $v \in Z_i$, $i = 1, 2, \dots, m$, our recursive construction did not modify the edges of the triangle t_v and the rhombus r_v (only the interior edges of a copy of R inside r_v are modified in subsequent phases). Let \mathcal{T}_i and \mathcal{R}_i denote the set of triangles t_v and rhombi r_v , respectively, for all $v \in Z_i$. Note that a monotone path enters and exits the interior of a triangle or a rhombus

in G at most once. Let γ be a path in $G(\Pi)$ that is monotone in some direction θ . Assume that $\theta \in \left[\frac{(j-1)\pi}{m}, \frac{j\pi}{m} \right)$, for some $j \in \{1, 2, \dots, m\}$. For $i < j$, the path γ enters at most one triangle of \mathcal{T}_i . For $j = i$, it can visit m triangles of \mathcal{T}_i , that all lie in a common triangle $t \in \mathcal{T}_{i-1}$. However, as we have discussed above, γ can reach the rhombus $r_v \subset t_v$ in at most two triangles $t_v \in \mathcal{T}_i$. For $i > j$, the path γ enters at most two triangles of \mathcal{T}_i , at most one inside each rhombus in \mathcal{R}_{i-1} .

It follows that γ traverses $O(m)$ edges in at most one zig-zag path created in phase j , and it traverses $O(1)$ edges created in any of the other $m - 1$ phases. Consequently, every monotone path in $G(\Pi)$ has at most $O(m)$ edges, as required. \square

We note that the construction uses only $O(m^2)$ distinct slopes of edges: at each level i , the copies of R are rotated into a fixed direction; this gives $O(m)$ slopes for the edges of R . R is embedded inside m distinctly oriented triangles t_v of the higher level; thus we need only m versions of the $O(1)$ edges filling the space between r_v and t_v , for another $O(m)$ slopes. This gives $O(m)$ slopes per level, and $O(m^2)$ slopes in total.

4 Proof of Theorem 3

We construct a polytope P with n vertices with no monotone path longer than $O(\log^2 n)$. The polytope will be built as a hierarchical system of *cells* Δ . The basic building block Δ is the polytope shown in Fig. 6. It is symmetric with respect to the yz -plane. The coordinates of its 10 vertices are given in Table 1. Its *base* is the outer triangle ABB' and it has two *inner triangles* UVW and $U'V'W'$ on which smaller (and transformed) copies of the cell Δ will be recursively placed, as sketched in Fig. 7.

point	(x, y, z)
A	$(0, 0, 0)$
B, B'	$(\mp 1, 1.5, 0)$
C	$(0, 1.4, 1)$
U, U'	$(\mp 0.1, 0.8, 0.55)$
V, V'	$(\mp 0.25, 0.6, 0.25)$
W, W'	$(\mp 0.25, 0.8, 0.39)$

Table 1: The vertex coordinates of the basic polytope Δ .

The *characteristic region* $\chi(\Delta)$ of Δ is the set of directions $\mathbf{c} = (u, v, 1)$ for which there exists a monotone path that starts at a boundary vertex (A , B , or B'), visits at least two vertices of each triangle UVW and $U'V'W'$, and ends at the boundary. The characteristic region for the standard cell is shown in Fig. 8. It has been generated with the help of a computer by enumerating all simple paths from the boundary via UVW and $U'V'W'$ (in either order) back to the boundary according to the definition, and by generating the set of directions for which the path is monotone. This is a (possibly empty) polyhedral set. There are 106 such paths in total; and among these, 38 paths with a nonempty region of directions. In Fig. 8, all these regions are drawn on top of each other, and their union yields the characteristic region. Our restriction to directions $\mathbf{c} = (u, v, w)$ with $w = 1$ excludes the horizontal directions ($w = 0$). These directions appear as infinite rays in the wedges of the figure.

As announced we will recursively place smaller copies of Δ on the *inner triangles* UVW and $U'V'W'$. Before placing them there, we subject them to geometric transformations to make them fit together as a convex polytope, and to ensure certain properties of the characteristic regions. The structure of this construction will be a balanced binary tree of depth k , with $k + 1$ levels and

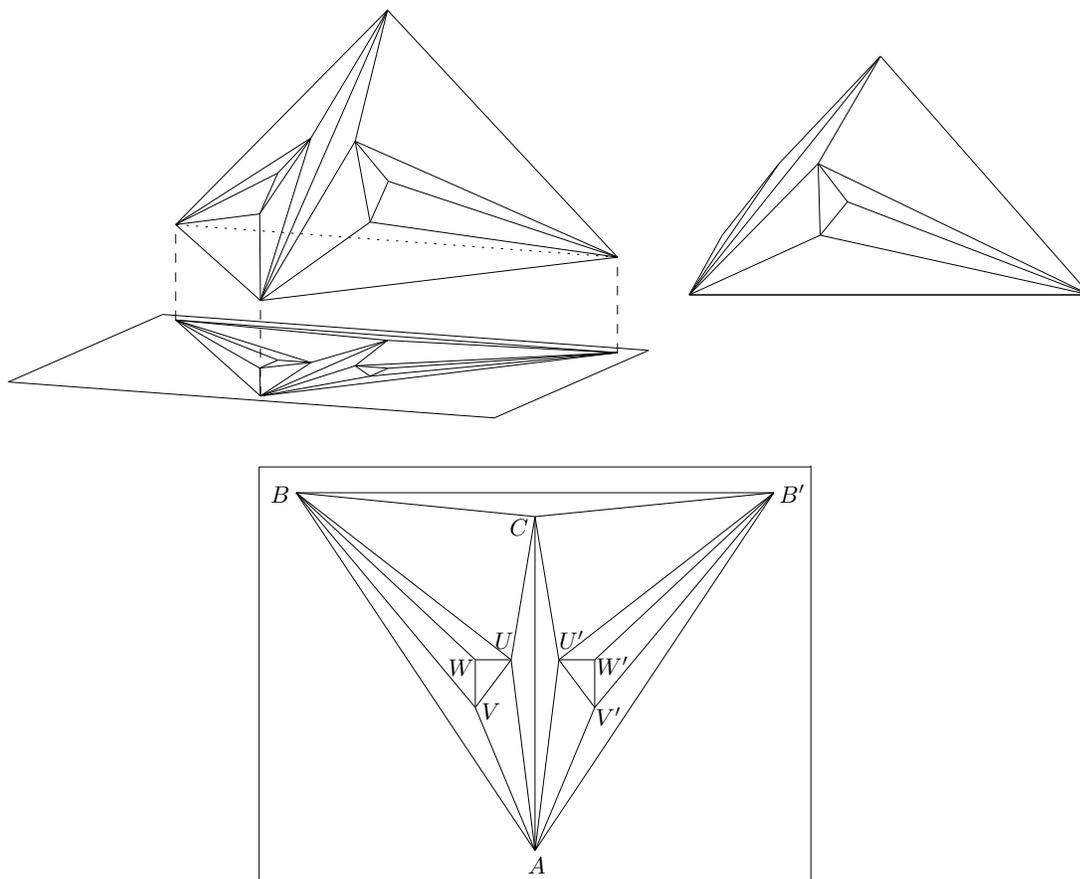


Figure 6: The basic building block Δ . The upper part shows two spatial views, and the lower part shows the vertical projection.

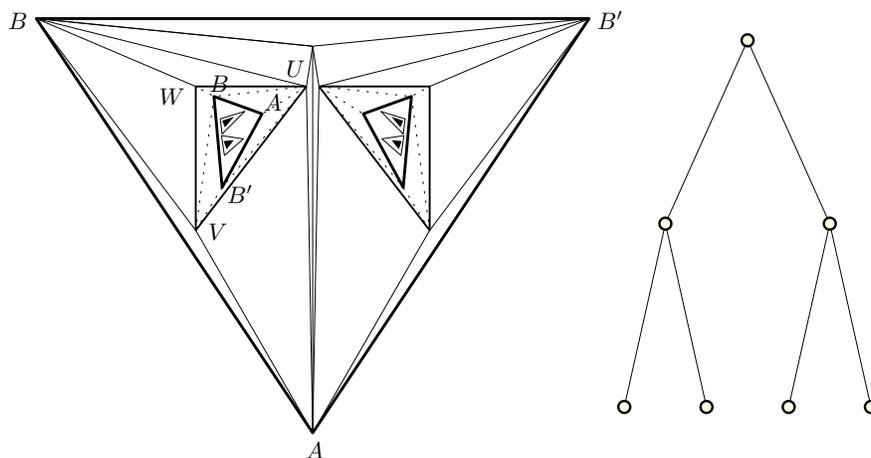


Figure 7: A schematic drawing of two nested levels of cells (not to scale) and the corresponding binary tree.

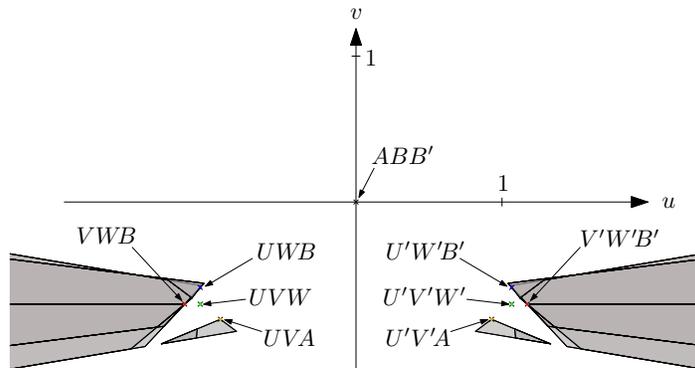


Figure 8: The characteristic region $\chi(\Delta)$. It is symmetric and consists of two unbounded wedge-like regions and two triangles. For orientation, the normal vectors of the triangles ABB' , UVW , and some other triangles are indicated. What is important is that UVW and $U'V'W'$ are outside $\chi(\Delta)$.

$2^{k+2} - 1$ nodes. Each vertex of P will belong to only one copy of Δ . Thus, there will be $10(2^{k+2} - 1)$ vertices in total. Each copy of Δ (each node of the tree) will have its characteristic region.

The root is the cell Δ in its original form. We describe how to place two transformed copies of a construction with k levels on Δ . Inductively, we maintain the following properties.

- (i) No three characteristic regions intersect.
- (ii) Two characteristic regions can only intersect in a disk of radius $R = 3$ around the origin.
- (iii) The characteristic regions lie inside the region $|v| \leq 2|u| + 1/2$.

For the basic building block itself (the induction basis $k = 0$), (i) and (ii) are vacuous, and (iii) can be checked by inspection, see Fig. 9.

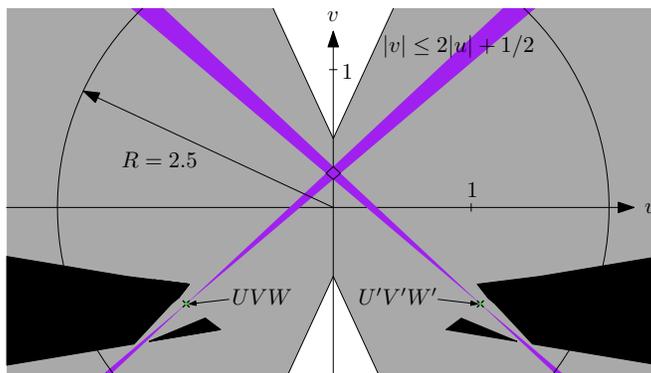


Figure 9: The recursive construction in terms of characteristic regions.

Geometric Transformations. Table 2 lists some simple affine transformations that can be applied to some polytope P , and the effect that they have on the characteristic region χ of some cell Δ that belongs to P . Here, $\varepsilon > 0$ is a small constant, a, b, α are arbitrary constants.

Suppose we have a polytope P with $k - 1$ levels that fulfills the properties (i)–(iii). We first *squeeze* it to make the wedge $|v| \leq 2|u| + 1/2$ narrow enough (see Fig. 10). Then we *turn* it, in order to be able to fit it between the two components of the characteristic region, as shown in Fig. 9. We

operation	coordinate transformation	effect on $\chi(\Delta)$
squash	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} x \\ y \\ \varepsilon \cdot z \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \varepsilon \cdot \begin{pmatrix} u \\ v \end{pmatrix}$
squeeze	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} \varepsilon \cdot x \\ y \\ \varepsilon \cdot z \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} u \\ \varepsilon \cdot v \end{pmatrix}$
scale	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \varepsilon \cdot \begin{pmatrix} x \\ y \\ z \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} u \\ v \end{pmatrix}$
turn	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}$
tilt	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} x \\ y \\ z + ax + by \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} u \\ v \end{pmatrix} - \begin{pmatrix} a \\ b \end{pmatrix}$
translate	$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} a \\ b \\ c \end{pmatrix}$	$\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} u \\ v \end{pmatrix}$

Table 2: The affine transformations to which the cells are subjected.

also have to *tilt* it, to make the base parallel to the triangle UVW . In terms of uv -coordinates, this corresponds to translating the origin to the normal vector of UVW that is shown in Fig. 9. We want to place the tilted base slightly above the triangle UVW . Thus, before tilting it, we *scale* it to make it small enough to fit inside UVW , and we *squash* it in the z -direction. Squashing achieves two purposes.

(a) The polytope P is made very flat (almost parallel to its base). Thus, when P is *translated* close enough over UVW so that it does not see any face beyond UVW , the triangle UVW in turn does not see any vertex of P except the base. Thus the convex hull of the union will have the combined set faces of Δ and P , except UWV and the base of P , plus a *collar* or *ring* of six triangle faces between UVW and the base of P . (The ring is shown with dotted edges in Fig. 7. The faces are guaranteed to be triangles except when the turning operation has created parallel edges between the base and UVW .)

(b) The second effect of squashing, in terms of characteristic regions, is to scale the uv -coordinates down. This brings any intersections between characteristic regions of P (which were initially within a radius of $R = 2.5$) as close as we want to the point representing the normal of UVW . We ensure that they are within distance $1/2$ of UVW .

A symmetric procedure is applied to another copy P' of P which is placed on $U'V'W'$.

Property (iii) can be checked visually from Fig. 9. Any intersections between two critical regions of P lie within distance $1/2$ of UVW the point, and thus within radius $R = 2.5$ of the origin. By construction there are no intersections between critical regions of P (they are confined in the squeezed wedge) and $\chi(\Delta)$. There can be intersections between critical regions of P and critical regions of P' . As Fig. 9 shows, these new intersections are confined to the intersections between the two corresponding squeezed wedges, which is a small quadrilateral area symmetric to the v -axis, close to the origin. So these new intersections are certainly within radius $R = 2.5$ about the origin, establishing property (ii). Finally, there cannot be any triple intersections, since the regions of P are disjoint when they are more than 0.5 away from UVW , and similarly for P' .

It is clear that the desired effects can be obtained by choosing the constants ε in the different steps small enough. One could work out particular values of ε (and α , a , and b) and get an explicit

affine transformation for placing P on top of UVW and $U'V'W'$.

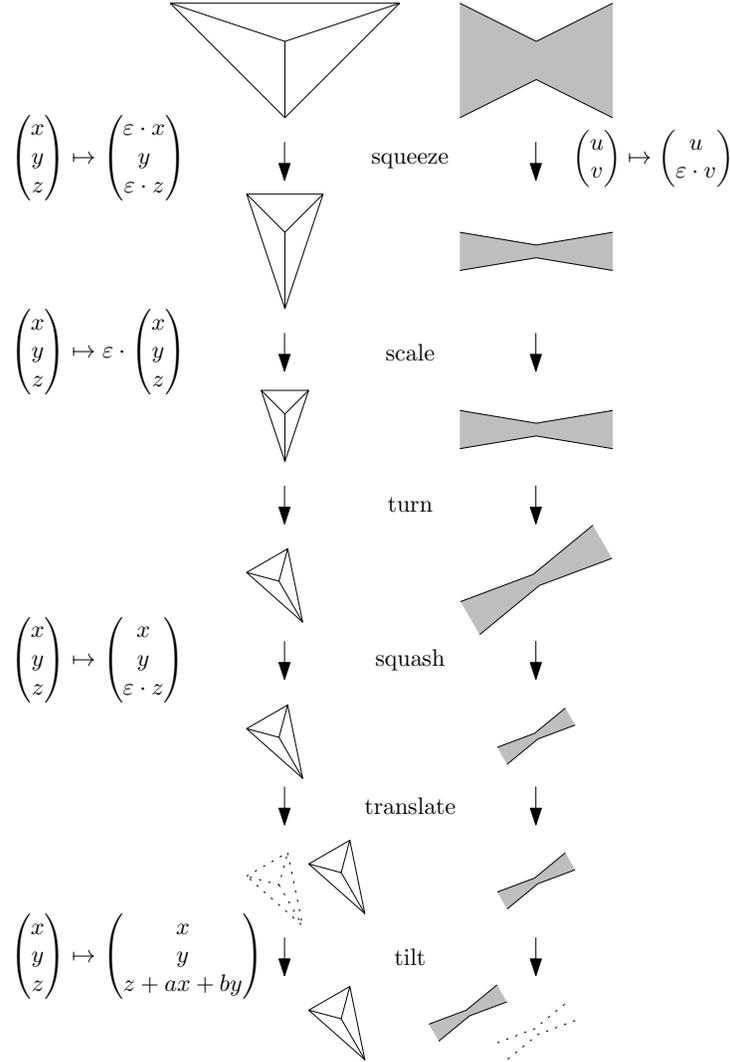


Figure 10: Schematic drawing of the sequence of transformations on Δ and their effects on $\chi(\Delta)$.

Let us now prove that there is no long monotone path in P . We say that a cell Δ is *visited* if the path contains at least two vertices of the cell. These might be just two boundary vertices like A and B , or also vertices inside Δ . From the definition of characteristic regions, we can conclude the following. Let Δ be a cell with children Δ_1 and Δ_2 . Then a path P that is monotone in direction \mathbf{c} can visit *both* Δ_1 and Δ_2 only if

- (a) P starts or ends inside Δ , or
- (b) \mathbf{c} lies in the characteristic region of Δ .

We now analyze the set S of nodes that correspond to the visited cells. By definition, if a cell is visited, its parent is also visited. Thus, S is a subtree of T containing the root.

A node Δ in S can have both its children in S only if conditions (a) and (b) are fulfilled. Therefore, there can be at most $2k + 1$ nodes with two children: $2k - 1$ because of condition (a), plus 2 because of condition (b), since by property (i) there cannot be more than two cells whose characteristic set \mathbf{c} lies in.

It follows that the tree S is the union of at most $2k + 2$ paths, each of length at most $k + 1$. So at most $2(k + 2)^2$ nodes are visited. The monotone path can have at most 12 vertices per visited node: 10 for the corresponding cell Δ , and at most one vertex from each child node. In total, the length is at most $24(k + 2)^2$.

The polytope P we have constructed has $n = (2^{k+1} - 1) \cdot 10 = 2^k \cdot 20 - 10$ vertices and $2n - 4 = 2^k \cdot 40 - 4$ triangular faces. To get an example P' with any given number of vertices, one can add (flat) pyramids on a subset of the faces. Between any two new vertices of P' , a path must visit at least one “old” vertex of P . Hence the length of monotone paths can at most double. \square

We can also analyze the size of the coordinates or the *spread* (ratio between the diameter and the smallest distance between two points) of this construction. The sequence of geometric transformations which are applied to place a small copy of the level- k construction on UVW can be composed into one fixed affine transformation (and its “mirror image”, for $U'V'W'$). This transformation is iterated k times. It follows that the size of coordinates, or the spread, grows like c^k for some constant $c > 1$; in terms of n , this is a polynomial.

Remark. We note that P is even a *stacked* polytope: such a polytope can be generated from an initial tetrahedron by successively gluing tetrahedra on triangular faces. (Equivalently, the graph $G(P)$ can be generated by successively inserting new degree-3 vertices into triangular faces.) It is easy to check that the basic building block Δ is a stacked polytope. The only place where one has to be careful are the connecting edges between successive levels of P (the dashed edges of Fig. 7). This connecting polytope is the convex hull of two triangles that lie in parallel planes. It is not hard to see that this such a polytope a stacked polytope if and only if the edge directions of the two triangles, UVW and the small copy of ABB' , do not alternate when going around the triangles. (Otherwise it has the combinatorial structure of an octahedron.) One can check that in our construction, this alternation does not happen: the main reason for this is that (contrary to the appearance in Fig. 7), the triangle ABB' is squeezed so much that the sides AB and AB' are almost parallel.

5 Proof of Theorem 4

The problem with the construction of the previous section is that it has to make the pieces very “flat” as the recursion proceeds. Therefore we could not control the entrance into sibling subtrees for the nodes in which the monotone path starts or ends.

In this section, we construct another polytope Q with a recursive structure. The “principal directions” of the parts will be chosen in advance, and they are not subject to transformations. We have to pay for this advantage by having a universal vertex A that is connected to all other vertices. In terms of a degree bound, our example is therefore as bad as it can be.

We start by picking $2^k + 1$ equally spaced *initial directions* on the cone K with the y -axis as the cone-axis and opening angle 45° , all pointing upwards such that all angles made by these directions with the xy -plane are bounded by 45° , see Fig. 11.

As in the previous section, we visualize directions by the (u, v) -component of the direction vectors $(u, v, 1)$. For an initial direction ℓ , the directions that are perpendicular to ℓ form a line in uv -space, called the *characteristic line* χ_ℓ of ℓ . The characteristic lines of the initial directions form an arrangement of lines in uv -space, as shown in the right part of Fig. 11. The perpendicular planes of the cone of initial directions are tangent to the polar cone K^* of K , which is in this case also a cone with opening angle 45° . The intersections of these planes with the plane $z = 1$, which are the characteristic lines that we see in Fig. 11, are therefore tangent to the hyperbola which is

the intersection of K^* with the plane $z = 1$. Since these lines are tangent to a common hyperbola, there are no triple intersections.

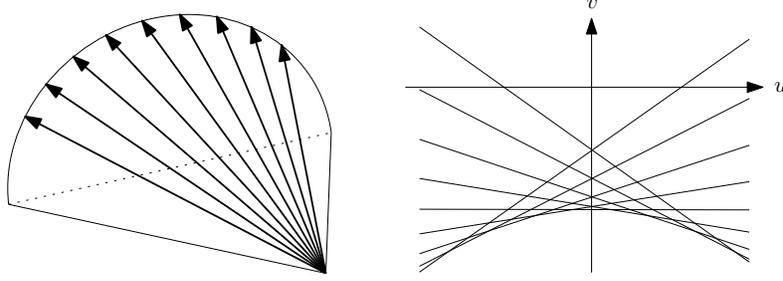


Figure 11: $2^k + 1 = 9$ initial directions, and the characteristic lines of their normals (for $k = 3$).

Every edge of our polytope Q (with one exception, the edge BB') will have a direction very close to one of the initial directions. It will therefore be difficult to traverse two such edges that are “almost parallel” in opposite directions. The following lemma quantifies this intuition.

Lemma 5. *Consider a line ℓ that makes an angle at most $\pi/4$ with the xy -plane, and two directed edges \mathbf{c}_1 and \mathbf{c}_2 which make an angle at most $\alpha \leq 1/10$ with ℓ but point in opposite directions. If a monotone path passes through both edges, then any direction $(u, v, 1)$ in which it is monotone lies inside the double-wedge of opening angle 2α and width 8α , centered about the characteristic line χ_ℓ of ℓ , as shown in Fig. 12.*

We call this region the *sharp-turn region* of ℓ .

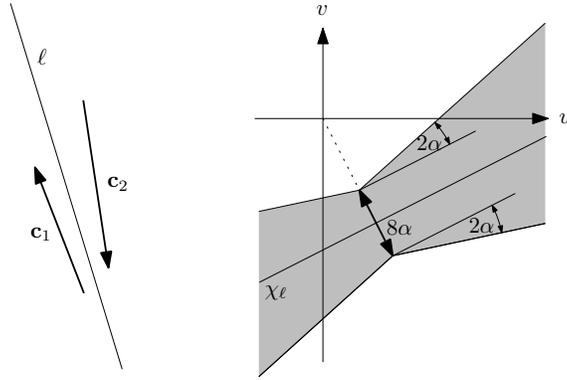


Figure 12: The sharp-turn region of two almost opposite directions.

Proof. This is an elementary calculation. We assume w.l.o.g. that ℓ has unit direction vector $(0, b, c)$. Since the angle with the horizontal plane is at most 45° , we have $c \leq b$, and since $b^2 + c^2 = 1$, we also get $b \geq 1/\sqrt{2}$, and

$$b - \alpha \geq 1/\sqrt{2} - 1/10 \geq 1/2. \quad (2)$$

The characteristic line χ_ℓ is then given by the solution of $0 \cdot u + b \cdot v + c \cdot 1 = 0$, that is, $v = -c/b$. (In Fig. 12 this would correspond to a rotated coordinate system where χ_ℓ becomes horizontal.)

We write the unit direction vector of \mathbf{c}_1 as $(0, b, c) + (\Delta x, \Delta y, \Delta z)$, with $\sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2} \leq 2 \sin \frac{\alpha}{2} \leq \alpha$. The normal plane to \mathbf{c}_1 appears in the uv -plane as given by the equation

$$\Delta x \cdot u + (b + \Delta y) \cdot v + (c + \Delta z) \cdot 1 = 0,$$

which yields

$$v = -\frac{c + \Delta z}{b + \Delta y} - \frac{\Delta x}{b + \Delta y} \cdot u.$$

We claim that (i) the slope $\frac{-\Delta x}{b + \Delta y}$ of this line is bounded by $\tan(2\alpha)$ in absolute value and (ii) the intercept $-\frac{c + \Delta z}{b + \Delta y}$ deviates from $-c/b$ by at most 4α . This means that the line lies within the claimed sharp-turn region.

The normal plane to \mathbf{c}_2 appears in the uv -plane as another line within this wedge, and the directions which are monotone in \mathbf{c}_1 and \mathbf{c}_2 must lie above one line and below the other, i.e., they are confined to some wedge within the sharp-turn region.

To finish the proof of the lemma, we verify the claims (i) and (ii).

(i) The absolute slope is bounded by $|\frac{-\Delta x}{b + \Delta y}| \leq \frac{|\Delta x|}{b - |\Delta y|} \leq \frac{\alpha}{b - \alpha} \leq \frac{\alpha}{1/2} = 2\alpha \leq \tan(2\alpha)$, by (2).

(ii) Similarly, the difference of intercepts is

$$\begin{aligned} \left| \frac{c}{b} - \frac{c + \Delta z}{b + \Delta y} \right| &= \left| \frac{c\Delta y - b\Delta z}{b(b + \Delta y)} \right| \leq \frac{c}{b} \left| \frac{\Delta y}{b + \Delta y} \right| + \left| \frac{\Delta z}{b + \Delta y} \right| \\ &\leq \left| \frac{\Delta y}{b + \Delta y} \right| + \left| \frac{\Delta z}{b + \Delta y} \right| = \frac{|\Delta y| + |\Delta z|}{|b + \Delta y|} \\ &\leq \frac{|\Delta y| + |\Delta z|}{b - |\Delta y|} \leq \frac{\alpha + \alpha}{b - \alpha} \leq \frac{2\alpha}{1/2} = 4\alpha. \quad \square \end{aligned}$$

By choosing α small enough, we can therefore ensure that sharp-turn regions of all initial directions remain so close to their characteristic lines that no triple intersections of sharp-turn regions appear. This implies that a monotone path can make at most two such “sharp turns” as specified in Lemma 5.

The polytope Q has an outer triangle ABB' and k interior layers $1, 2, \dots, k$. Layer i has 2^{i-1} vertices, and thus there are $2^k + 2$ vertices in total. A schematic drawing of the polytope with $k = 4$ inner layers is given in Fig. 13. From vertex A , there is an edge to each other vertex, and these edges have the chosen initial directions. We construct the polytope layer by layer. We start with an isosceles triangle ABB' , and the vertices B and B' form layer 0. The vertices of layer i are drawn interspersed (in angular order) with the vertices of the previous layers. When a vertex is first generated, it has two edges to previous layers and an edge to A . Later it collects two edges to each successive layer.

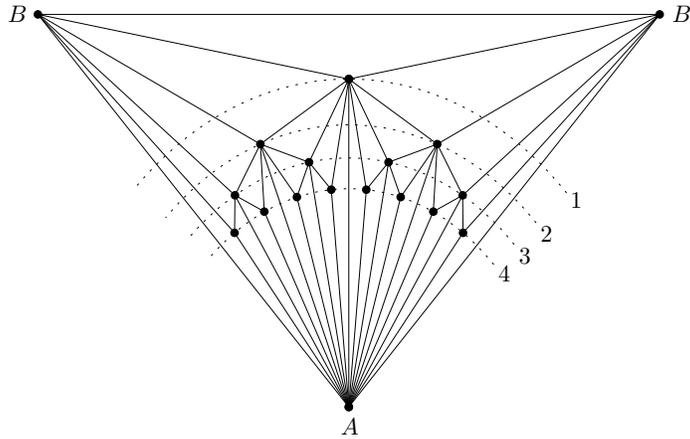


Figure 13: Schematic drawing of the polytope Q with $k = 4$ inner layers.

So far we have described the combinatorial structure of Q . The geometry is already constrained insofar as each vertex lies on one of the initial direction rays through A . It remains to specify the distances from A . All vertices of a level i have the same distance d_i from A . When constructing layer i , we choose the distance d_i so much smaller than all previous distances that the new points appear extremely close to A . More precisely, we want that an edge UV between a vertex U on layer $j < i$ and a vertex V on layer i has an angle at most α with the edge UA . By a continuity argument, such a value d_i always exists. For instance, it suffices to choose $d_i \leq \alpha d_{i-1}$. In addition, we have to ensure that the polytope has the right combinatorial structure, i.e., the convex hull edges are the edges shown in Fig. 13. By construction, all edges emanating from A lie on the convex hull. We need to make sure that the insertion of a layer does not destroy any previously inserted edge; the new layer will just add small triangular pyramids on top of some triangular faces. This can again be achieved by choosing d_i small enough, by a continuity argument. More explicitly, one can argue that is sufficient to set $d_i \leq 2d_{i-1}^2/d_0$: this ensures that the new vertex cannot even see any other vertex from previous layers in the plane projection (Fig. 13), without crossing edges incident to A . (In fact, it would not matter if edges were flipped; the only property we actually need is that no edge crosses between vertices of the same level.)

Now, consider a monotone path in the 1-skeleton of Q . If it goes through A , B or B' , we cut it at these vertices, into at most 4 pieces. To show that any monotone path on the boundary of Q has length $O(k) = O(\log n)$, it is therefore sufficient to consider monotone paths that do not pass through A , B or B' . Each edge in the remaining graph connects vertices from different levels. An edge on the monotone path can either ascend (into a smaller level) or descend (into a larger level). When the path changes at a vertex U from ascending to descending, it must traverse two edges that are almost parallel to UA , in opposite orientations. As we have observed, this can only happen at most twice. \square

Remark. By the way in which it is constructed, Q , just as P in the previous section, is a *stacked* polytope. As for the size of coordinates and the spread, we need to fulfill the conditions $d_i \leq \alpha d_{i-1}$ and $d_i \leq 2d_{i-1}^2/d_0$. Setting $d_0 = 1$, and $d_i = \alpha^{2^{i-1}}$ for $i \geq 1$ would work. The angle α can be chosen as a rational function involving a polynomial in n . But the whole construction leads then to coordinates that are exponential in n (but require only polynomially many bits, in a suitable fixed-precision approximation). As we have mentioned, if we do not insist on Q having the given combinatorial structure, we can ignore the second condition, and setting $d_i = \alpha^i$ for $i \geq 0$ leads to coordinates that are polynomial in n . But then we can no longer guarantee that Q is a stacked polytope.

6 Conclusion

Open problems. The proofs of Lemmas 1, 2, and 3 crucially depend on the planarity of the subdivisions. Extending Theorems 1 and 2 for convex subdivisions of \mathbb{R}^d , $d \geq 3$, remain as open problems.

It also left for future research to close the gaps between: (i) the upper bound $O(\log^2 n)$ in Theorem 3 and the trivial lower bound $\Omega(\log n)$ for n -vertex polytopes with bounded vertex degrees; (ii) the upper bound $O(\log n)$ in Theorem 4 and the lower bound $\Omega(\log n / \log \log n)$ for n -vertex polytopes with arbitrary vertex degrees.

We have shown (Theorem 2) that in any connected subdivision with n faces, $k \geq 3$ of which are unbounded, there exists a monotone path of length $\Omega(\log \frac{n}{k} / \log \log \frac{n}{k})$. Fig. 2 shows that this lower bound does not hold for (strictly) monotone paths starting from an arbitrary vertex. Deciding

whether there exists a *weakly* monotone path of length $\Omega(\log \frac{n}{k} / \log \log \frac{n}{k})$ starting from *every* vertex in a convex subdivision with k unbounded faces remains an open problem.

Acknowledgment. The authors thank János Pach for insistently asking the question to which Theorem 1 gives the answer. A recursive construction involving a hierarchy of squeezed zigzags, similar to the upper-bound construction for Theorem 2 in Section 3.2, has been suggested by Boris Bukh.

References

- [1] E. M. Arkin, R. Connelly, and J. S. B. Mitchell, On monotone paths among obstacles with applications to planning assemblies, in *Proc. 1st Sympos. on Comput. Geom.*, ACM Press, 1989, pp. 334–343.
- [2] F. Aurenhammer, A criterion for the affine equality of cell complexes in \mathbb{R}^d and convex polyhedra in \mathbb{R}^{d+1} , *Discrete & Comput. Geom.* **2** (1987), 49–64.
- [3] J. Balogh, O. Regev, C. Smyth, W. Steiger, and M. Szegedy, Long monotone paths in line arrangements, *Discrete & Comput. Geom.* **32** (2004), 167–176.
- [4] B. Chazelle, H. Edelsbrunner, and L. J. Guibas, The complexity of cutting complexes, *Discrete & Comput. Geom.* **4** (1989), 139–181.
- [5] A. Dumitrescu, On some monotone path problems in line arrangements, *Comput. Geom.* **32**(1) (2005), 13–25.
- [6] A. Dumitrescu and Cs. D. Tóth, Monotone paths in planar convex subdivisions, *Abstracts of the 21st Fall Workshop on Comput. Geom.*, 2011, New York.
- [7] A. Dumitrescu, G. Rote, and Cs. D. Tóth, Monotone paths in planar convex subdivisions, in: *Computing and Combinatorics. Proc. 18th Ann. Int. Computing and Combinatorics Conf. (COCOON 2012)*, ed. J. Gudmundsson, J. Mestre, and T. Viglas, Lect. Notes Comp. Sci., vol. **7434**, pp. 240–251, Springer-Verlag, 2012.
- [8] V. Klee, Heights of convex polytopes, *J. Math. Anal. Appl.* **11** (1965), 176–190.
- [9] P. McMullen, The maximum numbers of faces of a convex polytope, *Mathematika* **17** (1971) 179–184.
- [10] J. Pfeifle, Long monotone paths on simple 4-polytopes, *Israel J. Math.* **150**(1) (2005), 333–355.
- [11] J. Pfeifle and G. M. Ziegler, On the monotone upper bound problem, *Exper. Math.* **13** (2004), 1–11.
- [12] R. Radoičić and G. Tóth, Monotone paths in line arrangements, *Comput. Geom.* **24** (2003), 129–134.
- [13] G. Rote, Partial least-squares point matching under translations, *Abstracts of the 26th European Workshop on Comput. Geom.*, 2010, Dortmund, pp. 249–251.
- [14] G. Rote, Long monotone paths in convex subdivisions, in *Abstracts of the 27th European Workshop on Comput. Geom.*, 2011, Morschach, pp. 183–184.

- [15] F. Santos, A counterexample to the Hirsch conjecture, *Annals of Math. (2)* **176** (2012), 383–412. [arXiv:1006.2814](#).
- [16] M. J. Todd, The monotonic bounded Hirsch conjecture is false for dimension at least 4, *Math. Oper. Res.* **5**(4) (1980), 599–601.
- [17] Cs. D. Tóth, Stabbing numbers of convex subdivisions, *Period. Math. Hung.* **57**(2) (2008), 217–225.
- [18] G. M. Ziegler, *Lectures on Polytopes*, vol. 152 of GTM, Springer, 1994, pp. 83–93.