

ENUMERATIVE PROBLEMS FOR ARBORESCENCES AND MONOTONE PATHS ON POLYTOPES

CHRISTOS A. ATHANASIADIS, JESÚS A. DE LOERA,
AND ZHENYANG ZHANG

ABSTRACT. Every generic linear functional f on a convex polytope P orients the edges of the graph of P . In this directed graph one can define a notion of f -arborescence and f -monotone path on P . Additionally, a natural notion of adjacency between pairs of f -monotone paths gives us the so called *flip graph* of f -monotone paths. These concepts are of importance in geometric combinatorics and optimization.

We investigate the extreme values of the number of f -arborescences, the number of f -monotone paths, and the diameter of the flip graph of f -monotone paths where P ranges over all convex polytopes of given dimension and number of vertices and f ranges over all generic linear functionals on P .

1. INTRODUCTION AND RESULTS

Consider a d -dimensional convex polytope P in Euclidean space \mathbb{R}^d and a generic linear functional f on P , meaning a linear functional on \mathbb{R}^d which is nonconstant on every edge of P . This paper investigates extremal enumerative problems about f -arborescences and f -monotone paths on P . We first introduce briefly these notions and refer to Section 2 for more information.

The functional f , which we think of as an objective function, induces an orientation on the graph of P which orients every edge in the direction of increasing objective value. Such orientations of polytope graphs are called LP-admissible; they are of great importance in the study of the simplex method for linear optimization (see [11, 16] and the references given there). The resulting directed graph, consisting of all vertices and oriented edges of P and denoted by $\omega(P, f)$, is acyclic and has a unique source and a unique sink on every face of P . An f -monotone path on P is any directed path in $\omega(P, f)$ having as initial and terminal vertex the unique source, say v_{\min} , and the unique sink,

Date: February 5, 2020.

Keywords and phrases. Polytope, arborescence, monotone path, enumeration, simplex method, directed graphs, pivot rule.

say v_{\max} , of $\omega(P, f)$ on P , respectively. An f -arborescence is any (necessarily acyclic) spanning subgraph \mathcal{A} of the directed graph $\omega(P, f)$ such that for every vertex v of P there exists a unique directed path in \mathcal{A} with initial vertex v and terminal vertex v_{\max} . As explained in the sequel, f -arborescences and f -monotone paths are important notions in geometric combinatorics and optimization. When the context is clear, we simply refer to them as arborescences and monotone paths.

The set of all f -monotone paths on P can be given a natural graph structure as follows. We say that two f -monotone paths on P differ by a *polygon flip* (also called *polygon move*, or simply *flip*) across a 2-dimensional face F if they agree on all edges not lying on F but follow the two different f -monotone paths on F , from the unique source to the unique sink of $\omega(P, f)$ on F . The *graph of f -monotone paths* (also called *flip graph*) on P is denoted by $G(P, f)$ and is defined as the simple (undirected) graph which has nodes all f -monotone paths on P and as edges all unordered pairs of such paths which differ by a polygon flip across a 2-dimensional face of P . The graph $G(P, f)$ is connected; its higher connectivity was studied in [3], where it was shown that $G(P, f)$ is 2-connected for every polytope P of dimension $d \geq 3$ and $(d - 1)$ -connected for every simple polytope P of dimension d .

The main questions addressed in this paper ask to determine:

- the minimum and maximum number of f -arborescences on P ,
- the minimum and maximum number of f -monotone paths on P , and
- the minimum and maximum diameter of the graph $G(P, f)$,

where P ranges over all convex polytopes of given dimension and number of vertices and f ranges over all generic linear functionals on P . We will also consider these (or similar) questions when P is restricted to the important class of simple polytopes.

There are good reasons, from both a theoretical and an applied perspective, to study these problems. One motivation comes from the connection of f -arborescences and f -monotone paths to the behavior of the simplex method [21]. The simplex method produces a partial f -monotone path, traversing $\omega(P, f)$ from an initial vertex to the optimal one. The simplex method has to make decisions to choose the improving arcs via a *pivot rule*. It is an open problem to find the longest possible simplex method paths and little is known about bounds (see [9] and references therein). Clearly, the lengths of f -monotone paths are of great interest, as they bound the number of steps in the simplex algorithm. A pivot rule gives a mapping from the set of instances of the algorithm to the set of f -arborescences of $\omega(P, f)$. Two pivot rules

are equivalent if they always produce the same f -arborescence. Therefore, given P and f , there are only finitely many equivalence classes of pivot rules and counting f -arborescences is a proxy for the problem of counting pivot rules.

Another motivation comes from enumerative and polyhedral combinatorics, especially from the theory of fiber polytopes [7]. The flip graph of f -monotone paths on P contains a well behaved subgraph, namely that induced on the set of *coherent* f -monotone paths (these are the monotone paths which come from the shadow vertex pivot rule [10]). This subgraph is isomorphic to the graph of a convex polytope of dimension $d - 1$, where $d = \dim(P)$, which is a fiber polytope known as *monotone path polytope* [7, Section 5] [6]. Monotone paths, monotone path polytopes and flip graphs of polytopes of combinatorial interest often have elegant combinatorial interpretations. For example, the monotone path polytope of a cube is a permutohedron [7, Example 5.4], while the flip graph of the latter encodes reduced decompositions of a certain permutation and the braid relations among them [8, Section 2.4]. More generally, monotone paths on zonotopes [4, 20] correspond to certain galleries of chambers in a central hyperplane arrangement and the problem to estimate the diameter of the flip graph in this important special case has been intensely studied in [12, 20]. The diameter of flip graphs of fiber polytopes has also been studied in [18, 19]. Moreover, certain zonotopes are in fact monotone path polytopes coming from projecting cyclic polytopes [2, Section 3], or polytopes which look like piles of cubes [1]. Monotone path polytopes are also related to fractional power series solutions of algebraic equations [14]. The combinatorial properties of f -monotone paths and flip graphs have thus been studied in comparison to those of coherent f -monotone paths, but also because of their own independent interest.

A special role in our results is played by a distinguished member $X(n)$ of the family of stacked 3-dimensional simplicial polytopes with n vertices. As it turns out, this polytope maximizes the number of both f -arborescences and f -monotone paths, and possibly the diameter of the flip graph too, in this dimension. We refer to Section 2.1 for a discussion of stacked polytopes and the precise definition of $X(n)$, which we always consider endowed with the specific LP-allowable orientation given there. We will typically denote by n (and sometimes by $n + 1$) and m the number of vertices and facets of P , respectively. Let us also denote by

- $\tau(P, f)$ the number of f -arborescences on P ,
- $\mu(P, f)$ the number of f -monotone paths on P ,

- $\text{diam}(G)$ the diameter of the graph $G = G(P, f)$.

Our first two main results provide a fairly complete description of tight bounds for the numbers of f -arborescences and f -monotone paths and the diameter of the graph of f -monotone paths on a 3-dimensional polytope with given number of vertices. The upper bound for the number of f -monotone paths involves the sequence of Tribonacci numbers (sequence A000073 in [22]), defined by the recurrence $T_0 = T_1 = 1$, $T_2 = 2$ and $T_n = T_{n-1} + T_{n-2} + T_{n-3}$ for $n \geq 3$.

Theorem 1.1. *For $n \geq 4$,*

$$(1) \quad 2(n-1) \leq \tau(P, f) \leq 2 \cdot 3^{n-3}$$

$$(2) \quad \left\lceil \frac{n}{2} \right\rceil + 2 \leq \mu(P, f) \leq T_{n-1}$$

for every 3-dimensional polytope P with n vertices and every generic linear functional f on P . The upper bound is achieved by the stacked polytope $X(n)$ in both situations.

The lower bound of (1) can be achieved by pyramids and that of (2) by prisms, when n is even, and by wedges of polygons over a vertex, when n is odd. In particular, prisms minimize the number of f -monotone paths over all simple 3-dimensional polytopes with given number of vertices. Moreover,

$$\tau(P, f) = 3 \cdot 2^{(n-2)/2} = 3 \cdot 2^{m-3}$$

for every 3-dimensional simple polytope P with n vertices and m facets.

Theorem 1.2. *For every $n \geq 4$,*

$$(3) \quad \left\lceil \frac{(n-2)^2}{4} \right\rceil \leq \max \text{diam } G(P, f) \leq (n-2) \left\lfloor \frac{n-1}{2} \right\rfloor,$$

where P ranges over all 3-dimensional polytopes with n vertices and f ranges over all generic linear functionals on P .

Our results are substantially weaker in dimensions $d \geq 4$, where the upper bounds for the number of f -arborescences and the number of f -monotone paths are almost trivial, and leave plenty of room for further research.

Theorem 1.3. (a) *For $n > d \geq 4$,*

$$\tau(P, f) \leq (n-1)!$$

$$\mu(P, f) \leq 2^{n-2}$$

for every d -dimensional polytope P with n vertices and every generic linear functional f on P . These bounds are achieved by any 2-neighborly d -dimensional polytope with n vertices.

(b) For $m > d \geq 4$,

$$d \cdot ((d-1)!)^{m-d} \leq \tau(P, f) \leq \prod_{i=1}^d i^{h_i(m,d)}$$

for every simple d -dimensional polytope P with m facets and every generic linear functional f on P , where $(h_i(m, d))_{0 \leq i \leq d}$ is the h -vector of the d -dimensional cyclic polytope with m vertices. The lower and upper bounds are achieved by the polar duals of stacked simplicial polytopes and the polar duals of neighborly simplicial polytopes, respectively, of dimension d with m vertices.

The proofs of the results on f -arborescences, given in Section 3, rely on the fact that $\tau(P, f)$ is equal to the product of the outdegrees of the vertices of the directed graph $\omega(P, f)$ other than the sink (see Proposition 3.1). This has the curious consequence that $\tau(P, f)$ is independent of f for every simple polytope P . The proofs of the results on f -monotone paths and the diameter of flip graphs, given in Sections 4 and 5, respectively, use ideas from [3, Section 4] [6], reviewed in Section 2.2, to construct $G(P, f)$ as an inverse limit in the category of graphs and simplicial maps. Section 2 contains preliminary material on polytopes, needed to understand the main results and their proofs, defines the stacked polytope $X(n)$ and proves a combinatorial lemma about its diameter (Lemma 2.1) which implies the lower bound in Theorem 1.2. Section 6 concludes with comments about the missing bounds and related open problems.

2. PRELIMINARIES

This section reviews basic background and terminology on convex polytopes and monotone paths and discusses a few constructions and a preliminary result (Lemma 2.1) which will be useful in the sequel. We use the notation $[n] := \{1, 2, \dots, n\}$ for any positive integer n and refer the reader to the book [23] for any undefined concepts and terminology.

2.1. Some special polytopes. Special classes of polytopes play an important role in this paper, since they are optimal solutions of the extremal problems considered. Recall that a polytope is called *simplicial* if all its proper faces are simplices. The *simple polytopes* are the polar duals of simplicial polytopes. A convenient way to encode the numbers of faces of each dimension of a simple or simplicial d -dimensional polytope P is provided by the h -vector, denoted as $h(P) = (h_0(P), h_1(P), \dots, h_d(P))$; see pages 8, 59 and 248 of [23] for details and

more information. The h -vector of a simple polytope P has nonnegative integer coordinates which afford an elegant combinatorial interpretation: $h_k(P)$ equals the number of vertices of P of outdegree k in the directed graph $\omega(P, f)$, discussed in the introduction, for every generic linear functional f on P (see Sections 3.4 and 8.3 and Exercise 8.10 in [23]); in particular, the multiset of such outdegrees is independent of f .

A polytope is called *2-neighborly* if every pair of vertices is connected by an edge. A d -dimensional simplicial polytope is called *neighborly* if any $\lfloor d/2 \rfloor$ or fewer of its vertices form the vertex set of a face. Neighborly polytopes other than simplices (cyclic polytopes being distinguished representatives) exist in dimensions four and higher. Their significance comes from the fact that they maximize the entries of the h -vector among all polytopes with given dimension and number of vertices (see pages 15-16, and 254-257 of [23]); in particular, they maximize the numbers of faces of each dimension among such polytopes.




FIGURE 1. Example of the polytope $X(10)$

A *stacked polytope* is any simplicial polytope which can be obtained from a simplex by repeatedly glueing other simplices of the same dimension along common facets, so as to preserve convexity at each step. Equivalently, the boundary complex of a stacked polytope can be obtained combinatorially from that of a simplex by successive stellar subdivisions on facets. The h -vector of any stacked polytope P of dimension d with n vertices has the simple form $h(P) = (1, n-d, \dots, n-d, 1)$ (see [15]). A fundamental result of Barnette [5] states that among all simplicial polytopes with given dimension and number of vertices, the stacked polytopes have the fewest possible faces of each dimension. Moreover, as a consequence of the generalized lower bound theorem, stacked polytopes minimize the entries of the h -vector among all such polytopes (see [13, 17]).

Many different combinatorial types of stacked polytopes are possible. For each $n \geq 4$, we will consider a 3-dimensional stacked polytope of special type with n vertices, denoted by $X(n)$. This polytope comes together with a linear functional f which linearly orders its vertices as $f(v_1) < f(v_2) < \dots < f(v_n)$. The associated triangulation comprises of all faces of the simplices with vertex sets $\{v_1, v_2, v_3, v_4\}$, $\{v_2, v_3, v_4, v_5\}, \dots, \{v_{n-3}, v_{n-2}, v_{n-1}, v_n\}$, so the dual graph of this triangulation is a path (these dual graphs for general stacked polytopes are trees). The regularity of this triangulation easily implies that such polytope $X(n)$ and linear functional f exist for every $n \geq 4$. Figure 1 shows an example with $n = 10$.

A crucial property of $X(n)$ is that the directed graph $\omega(X(n), f)$ has as arcs the pairs (v_i, v_j) for $i, j \in \{1, 2, \dots, n\}$ with $j \in \{i+1, i+2, i+3\}$. The following combinatorial lemma establishes the lower bound for the diameter of flip graphs, claimed in Theorem 1.2.

Lemma 2.1. *The diameter of the graph of f -monotone paths on $X(n)$ is bounded below by $\lceil (n-2)^2/4 \rceil$ for every $n \geq 4$.*

Proof. Let G be the graph of f -monotone paths on $X(n)$. Denoting f -monotone paths as sequences of vertices, we set

$$\gamma = \begin{cases} (v_1, v_3, v_5, \dots, v_{n-1}, v_n), & \text{if } n \equiv 0 \pmod{2} \\ (v_1, v_2, v_4, \dots, v_{n-3}, v_{n-1}, v_n), & \text{if } n \equiv 1 \pmod{4} \\ (v_1, v_3, v_5, \dots, v_{n-2}, v_n), & \text{if } n \equiv 3 \pmod{4} \end{cases}$$

and $\delta = (v_1, v_2, v_3, \dots, v_n)$. We claim that γ and δ are at a distance of $\lceil (n-2)^2/4 \rceil$ apart in G . Clearly, the lemma follows from the claim.

We only consider the case that n is even, the other two cases being similar. By passing to the complement of the set of vertices appearing on an f -monotone path on $X(n)$, such paths correspond bijectively to the subsets of $\{v_2, v_3, \dots, v_{n-1}\}$ containing no three consecutive elements v_{k-1}, v_k, v_{k+1} . The subset which corresponds to γ , for instance, is $\{v_2, v_4, \dots, v_{n-2}\}$ and the one which corresponds to δ is the empty set. The 2-dimensional faces of $X(n)$ have vertex sets $\{v_1, v_2, v_3\}$, $\{v_{n-2}, v_{n-1}, v_n\}$ and $\{v_{k-1}, v_k, v_{k+2}\}$ and $\{v_{k-1}, v_{k+1}, v_{k+2}\}$ for $2 \leq k \leq n-2$. From these facts it follows that polygon flips on f -monotone paths on $X(n)$ correspond to the following operations on the corresponding subsets:

- removal of v_2 or v_{n-1} , if present,
- inclusion of v_2 , if absent and not both v_3 and v_4 are present,
- inclusion of v_{n-1} , if absent and not both v_{n-2} and v_{n-3} are present,

- removal or inclusion of one of v_k, v_{k+1} , if the other is present but v_{k-1} and v_{k+2} are absent.

To reach the empty set from $\{v_2, v_4, \dots, v_{n-2}\}$ with these operations one needs to remove each of v_2, v_4, \dots, v_{n-2} . A careful consideration shows that at least one flip is needed to remove v_2 , at least three more flips are needed to remove v_{n-2} , at least five more flips are needed to remove v_4 , and so on. For example, to remove v_{n-2} in at most three steps one needs to first include v_{n-1} , then remove v_{n-2} and finally remove v_{n-1} and to remove v_4 in at most five steps one needs to first include v_3 , then remove v_4 , include v_2 , remove v_3 and finally remove v_2 (in particular, v_4 cannot be removed before v_2 , v_{n-4} cannot be removed before v_{n-2} , and so on). This yields a distance of $1+3+5+\dots+(n-3) = (n-2)^2/4$ between γ and δ in G . \square

Remark 2.2. Perhaps it is instructive to visualize the process of flipping γ to δ , described in the previous proof. The two f -monotone paths are shown on Figure 2 for $n = 10$ and the sequence of 2-dimensional faces (recording only vertex indices, for simplicity) across which the flips occur could be $\{1, 2, 3\}$, $\{7, 9, 10\}$, $\{7, 8, 10\}$, $\{8, 9, 10\}$, $\{2, 3, 5\}$, $\{2, 4, 5\}$, $\{1, 2, 4\}$, $\{1, 3, 4\}$, $\{1, 2, 3\}$, $\{5, 7, 8\}$, $\{5, 6, 8\}$, $\{6, 8, 9\}$, $\{6, 7, 9\}$, $\{7, 9, 10\}$, $\{7, 8, 10\}$ and $\{8, 9, 10\}$. \square




FIGURE 2. Two monotone paths on $X(10)$

Finally, we consider prisms and wedges of polygons. Given a $(d-1)$ -dimensional polytope Q , the *prism* over Q is the d -dimensional polytope defined as the Cartesian product $Q \times [0, 1]$. The *wedge* of Q over a face F of Q is the d -dimensional polytope W obtained combinatorially from the prism $Q \times [0, 1]$ by collapsing the face $F \times [0, 1]$ to $F \times 0$. Note that Q becomes a facet of W and that if F is a facet and Q is simple,

then so is W . We will apply the wedge construction in the special cases that Q is a polygon and F is one of its vertices or edges.




FIGURE 3. The wedge of a pentagon over an edge

2.2. The graph of f -monotone paths. Let P be a d -dimensional polytope and f be a generic linear functional on P . We will assume that f does not take the same value on any two distinct vertices of P .

To investigate the graph of f -monotone paths on P , we will describe another way to construct it from simpler graphs, arising in the fibers of the restriction of the projection map f on P . The technical device needed, which we now review, is the *inverse limit* in the category of graphs and simplicial maps. This concept was introduced in [3, Section 4] (with motivation coming from [6]) to study the higher connectivity of $G(P, f)$; it leads to various more general graphs of partial f -monotone paths on P , a useful notion which allows for inductive arguments.

Let us linearly order the vertices v_0, v_1, \dots, v_n of P so that $f(v_0) < f(v_1) < \dots < f(v_n)$. We recall that for every interior point t of the interval $f(P)$, the fiber $P(t) := f^{-1}(t) \cap P$ of the map $f : P \rightarrow \mathbb{R}$ is a $(d - 1)$ -dimensional polytope and thus it has a well defined graph. Setting $t_i = f(v_i)$ for $0 \leq i \leq n$, we may thus consider the graph G_i of $P(t_i)$ for $0 \leq i \leq n$ and the graph $G_{i,i+1}$ of $P(t)$ for some $t_i < t < t_{i+1}$, for $0 \leq i \leq n - 1$ (the precise value of t being irrelevant because, by construction, the other choices of t in the same interval give a normally equivalent fiber $P(t)$); see Figure 4 for an example. Considering these graphs as one-dimensional simplicial complexes, we have a diagram

$$(4) \quad G_{0,1} \xrightarrow{\alpha_1} G_1 \xleftarrow{\beta_1} G_{1,2} \xrightarrow{\alpha_2} G_2 \xleftarrow{\beta_2} \dots \xleftarrow{\beta_{n-2}} G_{n-2,n-1} \xrightarrow{\alpha_{n-1}} G_{n-1} \xleftarrow{\beta_{n-1}} G_{n-1,n}$$

of graphs and simplicial maps for which $\alpha_i : G_{i-1,i} \rightarrow G_i$ and $\beta_i : G_{i,i+1} \rightarrow G_i$ result from the degeneration of the fiber $P(t)$ when t approaches t_i , with $t_{i-1} < t < t_i$ or $t_i < t < t_{i+1}$, respectively (recall that

a simplicial map of one-dimensional simplicial complexes maps vertices to vertices and either maps edges linearly onto edges, or contracts them to vertices; in particular, such a map is determined by its images on vertices).

The inverse limit G of this diagram is defined as follows. The nodes are the sequences

$$(v_{0,1}, v_{1,2}, \dots, v_{n-1,n}),$$

where $v_{i-1,i}$ is a vertex of $G_{i-1,i}$ for every $i \in [n]$ and $\alpha_i(v_{i-1,i}) = \beta_i(v_{i,i+1})$ for every $i \in [n-1]$. Two such sequences $(u_{0,1}, u_{1,2}, \dots, u_{n-1,n})$ and $(v_{0,1}, v_{1,2}, \dots, v_{n-1,n})$ are adjacent nodes in G if there exists a nonempty interval $\mathcal{I} \subseteq [n]$ such that:

- $u_{i-1,i}$ and $v_{i-1,i}$ are adjacent in $G_{i-1,i}$ for $i \in \mathcal{I}$,
- $u_{i-1,i} = v_{i-1,i}$ for $i \in [n] \setminus \mathcal{I}$, and
- the edges $\{u_{i-1,i}, v_{i-1,i}\}$ and $\{u_{i,i+1}, v_{i,i+1}\}$ are mapped homeomorphically onto the same edge of G_i by α_i and β_i , respectively, whenever $i, i+1 \in \mathcal{I}$.

This construction associates an inverse limit graph to any diagram of graphs and simplicial maps (4). As explained in [3, Section 4] (see [3, Proposition 4.1]), the graph G is isomorphic to $G(P, f)$ when the diagram comes from a polytope P and linear functional f , as just described. The inverse limit of a subdiagram of (7) of the form

$$G_{k-1,k} \xrightarrow{\alpha_k} G_k \xleftarrow{\beta_k} G_{k,k+1} \xrightarrow{\alpha_{k+1}} \dots \xleftarrow{\beta_{\ell-1}} G_{\ell-1,\ell} \xrightarrow{\alpha_\ell} G_\ell \xleftarrow{\beta_\ell} G_{\ell,\ell+1},$$

considered in Sections 4 and 5, has nodes which can be viewed as partial f -monotone paths on P , starting at the fiber $P(t)$ with $t_{k-1} < t < t_k$ and ending at $P(t')$ with $t_\ell < t' < t_{\ell+1}$, and adjacency given by a suitable extension of the notion of polygon flip, presented in the introduction.

3. ON THE NUMBER OF ARBORESCENCES

As explained in the introduction, we are interested in counting f -arborescences on a polytope P , meaning oriented trees in the directed graph $\omega(P, f)$ which are rooted at the unique sink v_{\max} . Recall that $\tau(P, f)$ denotes the number of f -arborescences on P . The following statement provides an explicit product formula for this number.

Proposition 3.1. *Given a d -dimensional polytope P and generic linear functional f , let $\text{out}_f(v)$ denote the outdegree of the vertex v of P in the directed graph $\omega(P, f)$. Then,*

$$\tau(P, f) = \prod_{v \neq v_{\max}} \text{out}_f(v),$$




FIGURE 4. A combinatorial cube and some of its fibers

where the product ranges over all vertices of P other than the sink v_{\max} . In particular, if P is simple, then

$$\tau(P, f) = \prod_{i=1}^d i^{h_i(P)}$$

is independent of f .

Proof. Since $\omega(P, f)$ is acyclic, an f -arborescence is uniquely determined by a choice of edge coming out of v for every vertex v of $\omega(P, f)$ other than the sink v_{\max} . Since there are exactly $\text{out}_f(v)$ choices for every such v , the proof of the first formula follows. The second formula follows from the first and the combinatorial interpretation of the h -vector of a simple polytope P , mentioned in Section 2.1. \square

Remark 3.2. Since every edge of $\omega(P, f)$ has a unique initial vertex, the sum of the outdegrees $\text{out}_f(v)$ of the vertices of P in the directed graph $\omega(P, f)$ is equal to the number of edges of P .

Corollary 3.3. For $m > d \geq 4$, the maximum number of f -arborescences over all simple d -dimensional polytopes with m facets is achieved by the

polar duals of neighborly polytopes and is given by the formula

$$\max \tau(P, f) = \prod_{i=1}^d i^{h_i(m,d)},$$

where $(h_i(m,d))_{0 \leq i \leq d}$ is the h -vector of the d -dimensional cyclic polytope with m vertices. Similarly, the minimum number of f -arborescences in this situation is achieved by the polar duals of stacked polytopes and is given by the formula

$$\min \tau(P, f) = d \cdot ((d-1)!)^{m-d}.$$

For 3-dimensional simple polytopes P with m facets, $\tau(P, f) = 3 \cdot 2^{m-3}$.

Proof. The case $d \geq 4$ follows from the last sentence of Proposition 3.1, the upper and lower bound theorems for the h -vector of a simplicial polytope, discussed in Section 2.1, and the formula for the h -vector of a d -dimensional stacked simplicial polytope with m vertices given there. The case $d = 3$ follows again from the second formula of Proposition 3.1, since $h_0(P) = h_3(P) = 1$ and $h_1(P) = h_2(P) = m - 3$ for every 3-dimensional simple polytope P with m facets. \square

The following two statements apply to general polytopes. Combined with Corollary 3.3, they imply the results about f -arborescences stated in the introduction.

Theorem 3.4. *For $n > d \geq 3$, the maximum number of f -arborescences over all d -dimensional polytopes with n vertices is achieved by the stacked polytope $X(n)$ for $d = 3$ and by any 2-neighborly polytope for $d \geq 4$. This number is equal to $2 \cdot 3^{n-3}$ and $(n-1)!$ in the two cases, respectively.*

Proof. Let us order the vertices v_1, v_2, \dots, v_n of the d -dimensional polytope P so that $f(v_1) \leq f(v_2) \leq \dots \leq f(v_n)$, where $v_n = v_{\max}$. Then, arcs of the directed graph $\omega(P, f)$ can only be pairs (v_i, v_j) with $i < j$ and hence $\text{out}_f(v_i) \leq n - i$ for every $i \in [n]$. Thus, in view of Proposition 3.1, we get

$$\tau(P, f) = \prod_{i=1}^{n-1} \text{out}_f(v_i) \leq \prod_{i=1}^{n-1} (n-i) = (n-1)!$$

and equality holds if and only if P is 2-neighborly.

Since no such polytopes other than simplices exist in dimension $d = 3$, this case has to be treated separately. Setting $d_i = \text{out}_f(v_i)$ for $i \in [n-1]$, we have positive integers d_1, d_2, \dots, d_{n-1} such that $d_{n-1} = 1$ and $d_{n-2} \in \{1, 2\}$. Since P can have no more than $3n - 6$ edges, we have $d_1 + d_2 + \dots + d_{n-1} \leq 3n - 6$ by Remark 3.2. It is an elementary

fact that, under these assumptions, the product $\tau(P, f) = d_1 d_2 \cdots d_{n-1}$ is maximized when $d_{n-1} = 1$, $d_{n-2} = 2$ and $d_i = 3$ for every $i \in [n-3]$. Exactly that happens for the stacked polytope $X(n)$ and the proof follows. \square

Theorem 3.5. *For all $n \geq 4$, the minimum number of f -arborescences over all 3-dimensional polytopes with n vertices is equal to $2(n-1)$. This is achieved by any pyramid P and any generic linear functional f which takes its minimum value on P at the apex.*

Proof. As a simple application of Proposition 3.1, we have $\tau(P, f) = 2(n-1)$ for every pyramid P over an $(n-1)$ -gon and every generic functional f which takes its minimum value on P at the apex.

We now consider any 3-dimensional polytope P with n vertices and any generic functional f on P . We need to show that $\tau(P, f) \geq 2(n-1)$. We may linearly order the vertices v_1, v_2, \dots, v_n of P in the order of decreasing outdegree in the directed graph $\omega(P, f)$ and denote by k the number of those vertices which have outdegree larger than one. Then, $k \geq 2$ and the respective outdegrees d_1, d_2, \dots, d_n of v_1, v_2, \dots, v_n satisfy $d_1, d_2, \dots, d_k \geq 2$, $d_n = 0$ and $d_i = 1$ for every other value of i . Letting D_1, D_2, \dots, D_n be the degrees of v_1, v_2, \dots, v_n in the undirected graph of P , respectively, we have $\tau(P, f) = d_1 d_2 \cdots d_k$ and

$$2 \cdot \sum_{i=1}^n d_i = \sum_{i=1}^n D_i$$

by Remark 3.2. Clearly, $D_i = d_i$ for one value of $i \in \{1, 2, \dots, k\}$ (the one corresponding to the source vertex), $D_i \geq d_i + 1$ for every other such value and $D_i \geq 3$ for all $k < i \leq n$. These considerations result in the inequality $d_1 + d_2 + \cdots + d_k \geq n + 1$ and thus, it remains to show that $d_1 d_2 \cdots d_k \geq 2(n-1)$ for every $k \geq 2$ and all $d_1, d_2, \dots, d_k \in \{2, 3, \dots, n-1\}$ summing at least to $n+1$. Indeed, from the inequality $ab > (a-1)(b+1)$ for integers $a \leq b$, applied repeatedly when b is the largest of d_1, d_2, \dots, d_k and a is any other number from these larger than 2, we get

$$d_1 d_2 \cdots d_k \geq (d_1 + d_2 + \cdots + d_k - 2k + 2) \cdot 2^{k-1} \geq (n - 2k + 3) \cdot 2^{k-1}.$$

Applying repeatedly the fact that $2m \geq m + 2$ for $m \geq 2$, we conclude that $d_1 d_2 \cdots d_k \geq 2(n-1)$ and the proof follows. \square

More generally, for any $d \geq 3$, the $(d-2)$ -fold pyramid P over an $(n-d+2)$ -gon has n vertices and dimension d . Moreover, if f is chosen so that every cone vertex has smaller objective value than any of the vertices of the $(n-d+2)$ -gon, then the number of f -arborescences on P is equal to $2(n-1)(n-2) \cdots (n-d+2)$.

Question 3.6. *What is the minimum number of f -arborescences over all d -dimensional polytopes with n vertices, for $d \geq 4$? Does it equal $2(n-1)(n-2)\cdots(n-d+2)$ for all $n > d \geq 4$?*

4. ON THE NUMBER OF MONOTONE PATHS

This section investigates the smallest and largest possible number of f -monotone paths on polytopes. For notational convenience, we let v_0, v_1, \dots, v_n be the vertices of a polytope P , linearly ordered so that $f(v_0) < f(v_1) < \cdots < f(v_n)$, as in Section 2.2. We recall that $\mu(P, f)$ denotes the number of f -monotone paths on P and that we refer to general directed paths in $\omega(P, f)$ as partial f -monotone paths, i.e., they may start or end at vertices other than v_{\min} or v_{\max} .

The following formula is the key to most results in this section.

Lemma 4.1. *The number of f -monotone paths on P can be expressed as*

$$\mu(P, f) = 1 + \sum_{k=0}^{n-1} (d_k - 1)\mu_k(P, f),$$

where $d_k = \text{out}_f(v_k)$ is the outdegree of v_k in $\omega(P, f)$ and $\mu_k(P, f)$ stands for the number of partial f -monotone paths on P with initial vertex v_0 and terminal vertex v_k .

Proof. Let $P(t) = f^{-1}(t) \cap P$ be the fibers of the map $f : P \rightarrow \mathbb{R}$, as in Section 2.2, and $t_i = f(v_i)$ for $0 \leq i \leq n$. For $0 \leq k \leq n-1$ let $\mathcal{H}_k(P, f)$ be the set of partial f -monotone paths on P having initial vertex v_0 and ending in the fiber $P(t)$ with $t_k < t < t_{k+1}$. Formally, these are essentially the nodes of the inverse limit of the part

$$G_{0,1} \xrightarrow{\alpha_1} G_1 \xleftarrow{\beta_1} G_{1,2} \xrightarrow{\alpha_2} G_2 \xleftarrow{\beta_2} \cdots \xrightarrow{\alpha_k} G_k \xleftarrow{\beta_k} G_{k,k+1}$$

of the diagram (4). Let $\eta_k(P, f)$ be the number of these partial f -monotone paths. We claim that

$$(5) \quad \eta_k(P, f) - \eta_{k-1}(P, f) = (d_k - 1)\mu_k(P, f)$$

for every $k \in [n-1]$. Since $\eta_0(P, f) = \text{out}_f(v_0) = d_0$ and $\mu_0(P, f) = 1$, this implies that

$$\eta_k(P, f) = 1 + \sum_{i=0}^k (d_i - 1)\mu_i(P, f)$$

for $0 \leq k \leq n-1$. Since $\eta_{n-1}(P, f) = \mu_n(P, f) = \mu(P, f)$, the desired formula follows as the special case $k = n-1$ of this equation.

To verify (5), let $\varphi_k : \mathcal{H}_k(P, f) \rightarrow \mathcal{H}_{k-1}(P, f)$ be the natural map obtained by restriction of diagrams. More intuitively, $\varphi_k(\gamma)$ is obtained

from $\gamma \in \mathcal{H}_k(P, f)$ by removing its last edge. Paths in $\mathcal{H}_{k-1}(P, f)$ and $\mathcal{H}_k(P, f)$ either pass through vertex v_k or not, depending on whether or not their last edge maps to v_k under the map α_k or β_k , respectively. Clearly, for every $\delta \in \mathcal{H}_{k-1}(P, f)$ which passes through v_k there are exactly d_k paths $\gamma \in \mathcal{H}_k(P, f)$ such that $\varphi_k(\gamma) = \delta$, obtained by choosing an edge of $\omega(P, f)$ coming out of v_k and attaching it to δ , while for every $\delta \in \mathcal{H}_{k-1}(P, f)$ which does not pass through v_k there is a unique path $\gamma \in \mathcal{H}_k(P, f)$ such that $\varphi_k(\gamma) = \delta$. These observations imply directly Equation (5) and the proof follows. \square

Recall that the Tribonacci sequence (T_n) is defined by the recurrence relation $T_0 = T_1 = 1$, $T_2 = 2$ and $T_n = T_{n-1} + T_{n-2} + T_{n-3}$ for $n \geq 3$.

Theorem 4.2. *The maximum number of f -monotone paths over all 3-dimensional polytopes with $n + 1$ vertices is equal to the n th Tribonacci number T_n for every $n \geq 3$. This is achieved by the stacked polytope $X(n)$.*

Proof. We proceed by induction on n . The result holds for $n = 3$, since there are exactly $T_3 = 4$ monotone paths on any 3-dimensional simplex. We assume that it holds for integers less than n and consider a 3-dimensional polytope P with $n + 1$ vertices v_0, v_1, \dots, v_n , linearly ordered as in the beginning of this section by a generic functional f .

We wish to apply Lemma 4.1. Since partial f -monotone paths on P with initial vertex v_0 and terminal vertex v_k are f -monotone paths on the convex hull of v_0, v_1, \dots, v_k , we have $\mu_k(P, f) \leq T_k$ for $k \in \{3, 4, \dots, n - 1\}$ by the induction hypothesis. Since this bound holds trivially for $k \in \{0, 1, 2\}$ as well, from Lemma 4.1 we get

$$\mu(P, f) \leq 1 + \sum_{k=0}^{n-1} (d_k - 1)T_k.$$

To bound the right-hand side, we note that

$$d_{n-k} + d_{n-k+1} + \dots + d_{n-1} \leq 3k - 3$$

for $k \in \{2, 3, \dots, n - 1\}$, since $d_{n-k} + d_{n-k+1} + \dots + d_{n-1}$ is equal to the number of edges of P connecting vertices $v_{n-k}, v_{n-k+1}, \dots, v_n$ and hence to the number of edges of a planar simple graph with $k + 1$ vertices. From these inequalities and the trivial one $d_{n-1} \leq 1$, and

setting $T_{-1} := 0$, we get

$$\begin{aligned}
\sum_{k=0}^{n-1} d_k T_k &= \sum_{k=1}^n (d_{n-1} + d_{n-2} + \cdots + d_{n-k})(T_{n-k} - T_{n-k-1}) \\
&\leq (T_{n-1} - T_{n-2}) + (3k - 3) \sum_{k=2}^n (T_{n-k} - T_{n-k-1}) \\
&= T_{n-1} + 2T_{n-2} + 3T_{n-3} + 3T_{n-4} + \cdots + 3T_0 \\
&= \sum_{k=1}^n T_k,
\end{aligned}$$

where the last equality comes from summing the recurrence $T_k = T_{k-1} + T_{k-2} + T_{k-3}$ for $k \in [n]$. We conclude that

$$\mu(P, f) \leq 1 + \sum_{k=0}^{n-1} (d_k - 1)T_k = 1 + \sum_{k=0}^{n-1} d_k T_k - \sum_{k=0}^{n-1} T_k \leq T_n.$$

This completes the induction.

Finally, it is straightforward to verify that the number of f -monotone paths on $X(n+1)$ satisfies the Tribonacci recurrence (or alternatively, that all inequalities hold as equalities in the previous argument) and is thus equal to T_n for every n . \square

Remark 4.3. The number of f -monotone paths on a polytope P with $n+1$ vertices is no larger than the number of subsets of its vertex set containing the source and the sink and hence at most 2^{n-1} . Equality holds exactly when P is 2-neighborly, meaning that the 1-skeleton of P is the complete graph on $n+1$ vertices, since then every such subset is the vertex set of an f -monotone path on P . As a result, the maximum number of f -monotone paths over all d -dimensional polytopes with $n+1$ vertices is equal to 2^{n-1} for all $n \geq d \geq 4$. \square

The following statement completes the proof of the results about the number of f -monotone paths, stated in the introduction.

Theorem 4.4. *The minimum number of f -monotone paths over all 3-dimensional polytopes with n vertices is equal to $\lceil n/2 \rceil + 2$. This is achieved by prisms, when n is even, and by wedges of polygons over a vertex, when n is odd.*

In particular, prisms minimize the number of f -monotone paths over all simple polytopes of dimension three with given number of vertices.

Proof. Applying Lemma 4.1 and noting that $\mu_k(P, f) \geq 1$ for every k , we get

$$\mu(P, f) \geq 1 + \sum_{k=0}^{n-2} (d_k - 1) = \sum_{k=0}^{n-2} d_k - n + 2.$$

Since $\sum_{k=0}^{n-2} d_k$ is equal to the number of edges of P (see Remark 3.2), which is bounded below by $\lceil 3n/2 \rceil$, it follows that $\mu(P, f) \geq \lceil n/2 \rceil + 2$. It is straightforward to verify that prisms achieve the minimum when n is even and wedges of polygons over a vertex (obtained from prisms by identifying two vertices at different levels which are connected by an edge) achieve the minimum when n is odd. \square

The lower bound for the number of f -monotone paths in any dimension, given in the following statement, is not expected to be tight.

Proposition 4.5. *The number of f -monotone paths on any polytope of dimension d with n vertices is bounded below by $\lceil dn/2 \rceil - n + 2$.*

Proof. Once again, this follows from the inequality $\sum_{k=0}^{n-2} d_k \geq \lceil dn/2 \rceil$ and Lemma 4.1. \square

We end this section with a conjecture for the maximum number of monotone paths on simple 3-dimensional polytopes. The proposed maximum can be achieved by wedges of polygons over an edge for which all vertices lie on a monotone path. We recall that the Fibonacci sequence (F_n) is defined by the recurrence $F_1 = F_2 = 1$ and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$.

Conjecture 4.6. *We have $\mu(P, f) \leq F_{n+2} + 1$ for every simple 3-dimensional polytope P with $2n$ vertices.*

The argument in the proof of Theorem 4.2 yields the following weaker result. Let (a_n) be the sequence of numbers defined by the recurrence relation $a_0 = a_1 = 1$, $a_2 = 2$, $a_3 = 4$ and $a_n = a_{n-1} + a_{n-2}$ for $n \geq 4$.

Proposition 4.7. *We have $\mu(P, f) \leq a_n$ for every 3-dimensional simple polytope P with $n + 1$ vertices and every generic linear functional f on P .*

Proof. We mimick the proof of Theorem 4.2. For the inductive step, since P is simple, we have $d_0 = 3$, $d_1, d_2, \dots, d_{n-2} \leq 2$ and $d_{n-1} = 1$ and compute that

$$\begin{aligned} \mu(P, f) &\leq 1 + \sum_{k=0}^{n-1} d_k a_k - \sum_{k=0}^{n-1} a_k \leq 1 + a_{n-2} + a_{n-3} + \dots + a_1 + 2a_0 \\ &\leq a_{n-1} + a_{n-2} = a_n, \end{aligned}$$




FIGURE 5. An example of a polytope on 8 vertices conjectured to be the maximizer of the number of monotone paths among simple polytopes.

since $a_{n-1} = 1 + a_{n-3} + \cdots + a_1 + 2a_0$. □

5. ON THE DIAMETER OF MONOTONE PATH GRAPHS

The main goal of this section is to prove Theorem 1.2.

The lower bound of (3) for the maximum diameter follows from Lemma 2.1. The upper bound will be deduced from the following result. Clearly, given a polytope P and a generic linear functional f , every f -monotone path on P meets each of the fibers $f^{-1}(t) \cap P$, where $t \in f(P)$, in a unique point. For f -monotone paths γ and γ' on P , let us denote by $\nu(\gamma, \gamma')$ the number of connected components of the set of values $t \in f(P)$ for which γ and γ' disagree on $f^{-1}(t) \cap P$. For example, for the two monotone paths, say γ and γ' , shown on Figure 2 we have $\nu(\gamma, \gamma') = 4$. Note that $\nu(\gamma, \gamma') = 0 \Leftrightarrow \gamma = \gamma'$.

Theorem 5.1. *Let P be a 3-dimensional polytope and f be a generic linear functional on P . The distance between any two f -monotone paths γ and γ' in the graph $G = G(P, f)$ satisfies*

$$(6) \quad d_G(\gamma, \gamma') \leq \frac{\nu(\gamma, \gamma')}{2} \cdot f_2(P),$$

where $f_2(P)$ is the number of 2-dimensional faces of P .

We will first state a technical result (see Proposition 5.2) which constructs a walk in $G(P, f)$ between two monotone paths γ and γ' with the required properties from walks on the fibers, assuming that the latter satisfy certain necessary compatibility conditions. To allow for all possible ways that γ and γ' may intersect each other, we consider the following general situation. Let \mathcal{F} be a connected polygonal complex in

\mathbb{R}^d having vertices v_0, v_1, \dots, v_n and $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a linear functional such that $f(v_0) < f(v_1) < \dots < f(v_n)$. The graph of f -monotone paths on \mathcal{F} , denoted by $G(\mathcal{F}, f)$, having initial vertex v_0 and terminal vertex v_n , can be defined with adjacency given by polygon flips just as in the special case in which \mathcal{F} is the 2-skeleton of a convex polytope (see Section 2.2). Alternatively, and in order to relate it to the graphs of the fibers of f , we may view $G(\mathcal{F}, f)$ as the inverse limit associated to a diagram

$$(7) \quad G_{0,1} \xrightarrow{\alpha_1} G_1 \xleftarrow{\beta_1} G_{1,2} \xrightarrow{\alpha_2} G_2 \xleftarrow{\beta_2} \dots \xleftarrow{\beta_{n-2}} G_{n-2,n-1} \xrightarrow{\alpha_{n-1}} G_{n-1} \xleftarrow{\beta_{n-1}} G_{n-1,n}$$

of graphs and simplicial maps. This is defined as in Section 2.2 provided the fiber $f^{-1}(t) \cap P$ is replaced with $f^{-1}(t) \cap \|\mathcal{F}\|$, where $\|\mathcal{F}\|$ is the polyhedron (union of faces) of \mathcal{F} . Thus, the G_i and $G_{i,i+1}$ are graphs of (one-dimensional) fibers $f^{-1}(t) \cap \|\mathcal{F}\|$ and the α_i and β_i are natural degeneration maps.

Given an f -monotone path γ on \mathcal{F} and $i \in [n]$, let us denote by $\pi_i(\gamma)$ the node of $G_{i-1,i}$ in which the union of the edges of γ intersects the corresponding fiber $f^{-1}(t) \cap \|\mathcal{F}\|$. Then, $\pi_i : G(\mathcal{F}, f) \rightarrow G_{i-1,i}$ is a simplicial map. Given a walk \mathcal{P} in a graph G , thought of as a sequence of edges, and a simplicial map $\varphi : G \rightarrow H$ of graphs, let us denote by $\varphi(\mathcal{P})$ the walk in H which is formed by the images of the edges of \mathcal{P} under φ , disregarding those edges of \mathcal{P} which are contracted to a node by φ .

Proposition 5.2. *Let γ and δ be f -monotone paths on \mathcal{F} . Suppose that for every $i \in [n]$ there exists a walk \mathcal{P}_i in $G_{i-1,i}$ with initial node $\pi_i(\gamma)$ and terminal node $\pi_i(\delta)$ which traverses each edge in $G_{i-1,i}$ exactly once, so that*

$$(8) \quad \alpha_i(\mathcal{P}_i) = \beta_i(\mathcal{P}_{i+1})$$

for every $i \in [n-1]$. Then, there exists a walk \mathcal{P} in $G(\mathcal{F}, f)$ with initial node γ and terminal node δ which traverses each 2-dimensional face of \mathcal{F} exactly once, such that $\pi_i(\mathcal{P}) = \mathcal{P}_i$ for every $i \in [n]$.

We first illustrate the proposition with an important special case and then use it to prove Theorem 5.1.

Example 5.3. To motivate the proof of Theorem 5.1, consider the special case in which the monotone paths γ and γ' do not have common vertices, other than those on which f attains its minimum and maximum value on P . Then, $\nu(\gamma, \gamma') = 1$ and the edges of γ and

γ' form a simple cycle C which divides the boundary of P into two closed balls, say B^+ and B^- , having common boundary C . Let \mathcal{F}^+ and \mathcal{F}^- be the two subcomplexes of the boundary complex of P which correspond to these balls. We wish to show that for each $\varepsilon \in \{+, -\}$, there exists a walk in $G(P, f)$ joining γ and γ' which traverses each 2-dimensional face of \mathcal{F}^ε exactly once. This would imply the desired bound for $d_G(\gamma, \gamma')$. Such a walk must traverse every edge of each fiber $f^{-1}(t) \cap B^\varepsilon$ exactly once and thus induce walks on these fibers with the same property.

Let us consider the diagram (7) for the polygonal complex \mathcal{F}^ε . Clearly, the fiber $f^{-1}(t) \cap \partial P$ is the boundary of a polygon for every interior point t of the interval $f(P)$, where ∂P denotes the boundary of P . Since, by the f -monotonicity of γ and γ' , this fiber intersects the cycle C , which is the boundary of the ball B^ε , in exactly two points, its intersection with B^ε must be homeomorphic to a line segment. Thus, all graphs appearing in the diagram (7) for \mathcal{F}^ε are path graphs, where $G_{i-1,i}$ has endpoints $\pi_i(\gamma)$ and $\pi_i(\gamma')$ for every $i \in [n]$. As a result, there are unique walks \mathcal{P}_i , as in the statement of Proposition 5.2, where condition (8) holds by the degeneration of fibers. Thus, Proposition 5.2 implies the existence of a walk in $G(\mathcal{F}^\varepsilon, f)$ with initial node γ and terminal node γ' which traverses each 2-dimensional face of \mathcal{F}^ε exactly once. \square

Proof of Theorem 5.1. We first observe that it suffices to prove the special case $\nu(\gamma, \gamma') = 1$. Indeed, given f -monotone paths γ and γ' on P and setting $\nu = \nu(\gamma, \gamma')$, it is straightforward to define f -monotone paths $\gamma = \gamma_0, \gamma_1, \dots, \gamma_\nu = \gamma'$ on P satisfying $\nu(\gamma_{i-1}, \gamma_i) = 1$ for every $i \in [\nu - 1]$. Then, the triangle inequality and the special case imply that

$$d_G(\gamma, \gamma') \leq \sum_{i=1}^{\nu} d_G(\gamma_{i-1}, \gamma_i) \leq \nu \cdot \frac{f_2(P)}{2},$$

as claimed by (6).

So let γ, γ' be f -monotone paths on P such that $\nu(\gamma, \gamma') = 1$. Let u and v be their unique common vertices, satisfying $f(u) < f(v)$, for which γ and γ' disagree on each fiber $f^{-1}(t) \cap P$ with $f(u) < t < f(v)$ and agree on the other fibers; in the special case of Example 5.3, u and v are the unique vertices v_{\min} and v_{\max} on which f attains its minimum and maximum value on P , respectively. As in that special case, the edges of γ and γ' joining u and v form a simple cycle C which divides the 2-dimensional sphere ∂P into two closed 2-dimensional balls B^+ and B^- having common boundary C . Moreover, the f -monotonicity of γ

and γ' implies that for each $\varepsilon \in \{+, -\}$ and every interior point t of the interval $f(B^\varepsilon)$, the fiber $f^{-1}(t) \cap B^\varepsilon$ is homeomorphic to a line segment or a circle. We wish to apply Proposition 5.2 to the subcomplex \mathcal{F}^ε of the boundary complex of P which corresponds to B^ε .

We claim that there exist unique walks \mathcal{P}_i satisfying the assumptions of the proposition. Indeed, according to our previous discussion, every graph $G_{i-1,i}$ appearing in the diagram (7) for \mathcal{F}^ε is either a path graph, with endpoints $\pi_i(\gamma)$ and $\pi_i(\gamma')$, or a cycle. As a result, there exists a unique walk \mathcal{P}_i in $G_{i-1,i}$ with initial node $\pi_i(\gamma)$ and terminal node $\pi_i(\gamma')$ which traverses each edge in $G_{i-1,i}$ exactly once, if the latter is a path graph, and exactly two such walks, corresponding to the two possible orientations of $G_{i-1,i}$, if the latter is a cycle. There are the following cases, illustrated in Example 5.4, to consider:

Case 1: The relative interior of B^ε contains neither v_{\min} nor v_{\max} . Then, all the $G_{i-1,i}$ are path graphs and conditions (8) hold by degeneration of fibers, as in the special case $u = v_{\min}$ and $v = v_{\max}$ of Example 5.3.

Case 2: The relative interior of B^ε contains exactly one of v_{\min} and v_{\max} , say v_{\min} . Then, the $G_{i-1,i}$ associated to fibers $f^{-1}(t) \cap B^\varepsilon$ with $t < f(u)$ are cycles and all others are path graphs which degenerate to cycles as the value of f approaches $f(u)$. Clearly, the cycles can be uniquely oriented, so that the resulting walks \mathcal{P}_i satisfy conditions (8).

Case 3: The relative interior of B^ε contains both v_{\min} and v_{\max} . Then, the $G_{i-1,i}$ associated to fibers $f^{-1}(t) \cap B^\varepsilon$ with $f(u) < t < f(v)$ are path graphs and the rest are cycles which can be uniquely oriented, so that the resulting walks \mathcal{P}_i satisfy conditions (8).

Thus, Proposition 5.2 applies in all cases and we may conclude that $d_G(\gamma, \gamma') \leq f_2(\mathcal{F}^\varepsilon)$ for each $\varepsilon \in \{+, -\}$. Hence,

$$d_G(\gamma, \gamma') \leq \frac{f_2(\mathcal{F}^+) + f_2(\mathcal{F}^-)}{2} = f_2(P)/2$$

and the proof follows. \square

Example 5.4. Let $P = X(10)$ be the stacked polytope shown in Figure 1. The following two situations illustrate the three cases within the proof of Theorem 5.1.

(a) Consider the f -monotone paths on P

$$\begin{aligned} \gamma &= (v_1, v_3, v_6, v_9, v_{10}), \\ \gamma' &= (v_1, v_3, v_5, v_8, v_9, v_{10}), \end{aligned}$$

presented as sequences of vertices. Then, the cycle C has edges with vertex sets $\{v_3, v_5\}$, $\{v_5, v_8\}$, $\{v_8, v_9\}$, $\{v_6, v_9\}$ and $\{v_3, v_6\}$, and one

of the \mathcal{F}^ε consists of the faces of the facets of P with vertex sets $\{v_3, v_5, v_6\}$, $\{v_5, v_6, v_8\}$ and $\{v_6, v_8, v_9\}$ and falls in the first case of the proof, while the other consists of the faces of the remaining thirteen facets of P and falls in the third case. Three flips are needed to reach γ' from γ across \mathcal{F}^ε in the former case, and thirteen flips in the latter. (b) Consider also the f -monotone paths

$$\begin{aligned}\gamma &= (v_1, v_3, v_6, v_9, v_{10}), \\ \gamma'' &= (v_1, v_3, v_4, v_5, v_8, v_9, v_{10}).\end{aligned}$$

Now C has six edges with vertex sets $\{v_3, v_4\}$, $\{v_4, v_5\}$, $\{v_5, v_8\}$, $\{v_8, v_9\}$, $\{v_6, v_9\}$ and $\{v_3, v_6\}$, and one of the \mathcal{F}^ε consists of the faces of the facets of P with vertex sets $\{v_1, v_2, v_3\}$, $\{v_1, v_2, v_4\}$, $\{v_1, v_3, v_4\}$, $\{v_2, v_3, v_5\}$, $\{v_2, v_4, v_5\}$, $\{v_3, v_5, v_6\}$, $\{v_5, v_6, v_8\}$ and $\{v_6, v_8, v_9\}$, while the other consists of the faces of the remaining eight facets of P . Both fall in the second case of the proof. The fibers $f^{-1}(t) \cap B^\varepsilon$ are path graphs for $f(v_3) < t < f(v_9)$ in either case, and cycles for $t \leq f(v_3)$ or $t \geq f(v_9)$ in the two cases, respectively. \square

Proof of Theorem 1.2. As we have already mentioned, the lower bound of (3) follows from Lemma 2.1. The upper bound follows from Theorem 5.1 and the obvious inequalities $\nu(\gamma, \gamma') \leq \lfloor (n-1)/2 \rfloor$ and $f_2(P) \leq 2n - 4$. \square

Question 5.5. *What is the exact value of the maximum diameter in Theorem 1.2? In particular, is it equal to the lower bound given there for every n ?*

Proof of Proposition 5.2. Consider indices $0 < k \leq m \leq \ell < n$ and denote by K and L the graphs of partial f -monotone paths on \mathcal{F} which arise as inverse limits of the subdiagrams

$$(9) \quad G_{k-1,k} \xrightarrow{\alpha_k} G_k \xleftarrow{\beta_k} G_{k,k+1} \xrightarrow{\alpha_{k+1}} \cdots \xrightarrow{\alpha_{m-1}} G_{m-1} \xleftarrow{\beta_{m-1}} G_{m-1,m}$$

and

$$(10) \quad G_{m,m+1} \xrightarrow{\alpha_{m+1}} G_{m+1} \xleftarrow{\beta_{m+1}} \cdots \xleftarrow{\beta_{\ell-1}} G_{\ell-1,\ell} \xrightarrow{\alpha_\ell} G_\ell \xleftarrow{\beta_\ell} G_{\ell,\ell+1}$$

of (7), respectively. Let us call *polygon* any 2-dimensional face of \mathcal{F} which intersects the fiber $f^{-1}(t) \cap \|\mathcal{F}\|$ for some $t_{k-1} < t < t_m$ in the case of (9) and any 2-dimensional face of \mathcal{F} which intersects the fiber $f^{-1}(t) \cap \|\mathcal{F}\|$ for some $t_m < t < t_{\ell+1}$ in the case of (10). Thus, the polygons are exactly the 2-dimensional faces of \mathcal{F} in the case of the entire diagram (7) and are in one-to-one correspondence with the

edges of $G_{m-1,m}$ in the special case $k = m$ of (9). Define similarly the graph H of partial f -monotone paths on \mathcal{F} and its polygons from the subdiagram

$$(11) \quad G_{k-1,k} \xrightarrow{\alpha_k} \cdots \xleftarrow{\beta_{m-1}} G_{m-1,m} \xrightarrow{\alpha_m} G_m \xleftarrow{\beta_m} G_{m,m+1} \xrightarrow{\alpha_{m+1}} \cdots \xleftarrow{\beta_\ell} G_{\ell,\ell+1}$$

of (7) and note that there are natural restriction maps $\pi_K : G(\mathcal{F}, f) \rightarrow K$, $\pi_L : G(\mathcal{F}, f) \rightarrow L$ and $\pi_H : G(\mathcal{F}, f) \rightarrow H$.

Assuming that there exist a walk \mathcal{Q} in K with initial node $\pi_K(\gamma)$ and terminal node $\pi_K(\delta)$ which traverses each polygon of (9) exactly once and a walk \mathcal{R} in L with initial node $\pi_L(\gamma)$ and terminal node $\pi_L(\delta)$ which traverses each polygon of (10) exactly once, such that $\pi_i(\mathcal{Q}) = \mathcal{P}_i$ for $k \leq i \leq m$ and $\pi_i(\mathcal{R}) = \mathcal{P}_i$ for $m < i \leq \ell + 1$, we claim that there exists a walk \mathcal{P} in H with initial node $\pi_H(\gamma)$ and terminal node $\pi_H(\delta)$ which traverses each polygon of (11) exactly once, such that $\pi_i(\mathcal{P}) = \mathcal{P}_i$ for $k \leq i \leq \ell + 1$. The proposition then follows by applying the claim several times, for instance when $k = 1$ and $m = \ell$, for $m \in [n - 1]$.

To prove the claim, we only need to patch \mathcal{Q} and \mathcal{R} along the walk $\alpha_m(\mathcal{P}_m) = \beta_m(\mathcal{P}_{m+1})$ in G_m . Any two nodes ζ of K and η of L produce by concatenation a node $\zeta * \eta$ of H , provided that the terminal edge of ζ and the initial edge of η have equal images under α_m and β_m , respectively. Let $\zeta_0, \zeta_1, \dots, \zeta_q$ be the successive nodes of \mathcal{Q} and $\eta_0, \eta_1, \dots, \eta_r$ be the successive nodes of \mathcal{R} . By our assumptions, we have $\zeta_0 * \eta_0 = \pi_K(\gamma) * \pi_L(\gamma) = \pi_H(\gamma)$ and $\zeta_q * \eta_r = \pi_K(\delta) * \pi_L(\delta) = \pi_H(\delta)$. We define \mathcal{P} to have nodes of the form $\zeta_i * \eta_j$, starting with $\zeta_0 * \eta_0$, so that the node immediately following $\zeta_i * \eta_j$ is

$$(12) \quad \begin{cases} \zeta_{i+1} * \eta_j, & \text{if well defined,} \\ \zeta_i * \eta_{j+1}, & \text{if well defined but } \zeta_{i+1} * \eta_j \text{ is not,} \\ \zeta_{i+1} * \eta_{j+1}, & \text{otherwise.} \end{cases}$$

We leave to the reader to verify that, because $\alpha_m(\mathcal{P}_m) = \beta_m(\mathcal{P}_{m+1})$, this is a well defined walk in H with initial node $\zeta_0 * \eta_0 = \pi_H(\gamma)$ and terminal node $\zeta_q * \eta_r = \pi_H(\delta)$. By construction, we have $\pi_i(\mathcal{P}) = \pi_i(\mathcal{Q})$ for $k \leq i \leq m$ and $\pi_i(\mathcal{P}) = \pi_i(\mathcal{R})$ for $m < i \leq \ell + 1$, and hence $\pi_i(\mathcal{P}) = \mathcal{P}_i$ for $k \leq i \leq \ell + 1$. Finally, we note that \mathcal{P} traverses the polygons traversed by \mathcal{Q} or \mathcal{R} which do not intersect the fiber $f^{-1}(t_m) \cap \|\mathcal{F}\|$ by steps which move $\zeta_i * \eta_j$ to the first two paths shown in (12), respectively, each exactly once by our assumptions on \mathcal{Q} and

\mathcal{R} , and the 2-dimensional faces of \mathcal{F} which intersect $f^{-1}(t_m) \cap \|\mathcal{F}\|$ by steps which move $\zeta_i * \eta_j$ to the third path shown in (12), each exactly once by our assumptions on \mathcal{P}_m and \mathcal{P}_{m+1} , and that these are precisely the polygons of (11). \square

6. CONCLUSIONS

The following tables summarize our results on monotone paths and arborescences and indicate the problems which remain open.

We have no reason to doubt that Question 3.6 on the minimum number of f -arborescences in dimensions $d \geq 4$ and Question 5.5 on the maximum diameter of flip graphs in dimension 3 have positive answers. For the minimum diameter of flip graphs, we expect that the diameter of $G(P, f)$ is bounded below by the integral part of half the number of facets for every 3-dimensional polytope P . In particular, we expect that the following conjecture is true.

Conjecture 6.1. *The minimum diameter of $G(P, f)$, when P ranges over all 3-dimensional polytopes with n vertices and f ranges over all generic linear functionals on P , is equal to $\lfloor (n+5)/4 \rfloor$ for every $n \geq 4$. This can be achieved by simple polytopes for every even n .*

# of arborescences		all polytopes	simple polytopes
$d = 3$	upper bound	Theorem 3.4	Corollary 3.3
	lower bound	Theorem 3.5	
$d \geq 4$	upper bound	Theorem 3.4	Corollary 3.3
	lower bound	Question 3.6	Corollary 3.3

TABLE 1. Summary for f -arborescences

# of monotone paths		all polytopes	simple polytopes
$d = 3$	upper bound	Theorem 4.2	Conjecture 4.6, Proposition 4.7
	lower bound	Theorem 4.4	
$d \geq 4$	upper bound	Remark 4.3	open
	lower bound	Proposition 4.5	

TABLE 2. Summary for f -monotone paths

diameter of flip graph		all polytopes	simple polytopes
$d = 3$	upper bound	Theorem 1.2, Question 5.5	open
	lower bound	Conjecture 6.1	
$d \geq 4$	upper bound	open	open
	lower bound	open	open

TABLE 3. Summary for the diameter of flip graphs

The outdegrees of the vertices of $\omega(P, f)$ play an important role in the proofs of Theorems 1.1 and 1.3. It seems a very interesting problem to characterize, or at least obtain significant information about, the possible multisets of these outdegrees when P ranges over all polytopes of given dimension and number of vertices and f ranges over all generic linear functionals on P . Finally, it would be interesting to address the questions raised in this paper for *coherent* f -monotone paths as well. Their number typically grows much slower than the total number of f -monotone paths [2].

Acknowledgements: The first author warmly thanks the Department of Mathematics at UC Davis for hosting him and providing an excellent academic atmosphere during the Fall of 2019, when this paper was prepared. The second and third authors are grateful for the support received through NSF grants DMS-1818969 and HDR TRIPODS.

REFERENCES

- [1] C.A. Athanasiadis. Piles of cubes, monotone path polytopes, and hyperplane arrangements. *Discrete Comput. Geom.*, 21(1):117–130, 1999. URL: <https://doi.org/10.1007/PL00009404>, doi:10.1007/PL00009404.
- [2] C.A. Athanasiadis, J.A. De Loera, V. Reiner, and F. Santos. Fiber polytopes for the projections between cyclic polytopes. *European J. Combin.*, 21:19–47, 2000.
- [3] C.A. Athanasiadis, P.H. Edelman, and V. Reiner. Monotone paths on polytopes. *Math. Z.*, 235:549–555, 2000.
- [4] C.A. Athanasiadis and F. Santos. Monotone paths on zonotopes and oriented matroids. *Canad. J. Math.*, 53(6):1121–1140, 2001. URL: <https://doi.org/10.4153/CJM-2001-042-3>, doi:10.4153/CJM-2001-042-3.
- [5] D. Barnette. A proof of the lower bound conjecture for convex polytopes. *Pacific J. Math.*, 46:349–354, 1973. URL: <http://projecteuclid.org/euclid.pjm/1102946311>.
- [6] L.J. Billera, M. M. Kapranov, and B. Sturmfels. Cellular strings on polytopes. *Proc. Am. Math. Soc.*, 122(2):549–555, 1994.
- [7] L.J. Billera and B. Sturmfels. Fiber polytopes. *Ann. of Math.*, 135(3):527–549, 1992.

- [8] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, second edition, 1999. doi:10.1017/CB09780511586507.
- [9] M. Blanchard, J.A. De Loera, and Q. Louveaux. On the length of monotone paths in polyhedra. In preparation.
- [10] D. Dadush and N. Hähnle. On the shadow simplex method for curved polyhedra. *Discrete Comput. Geom.*, 56(4):882–909, 2016. URL: <https://doi.org/10.1007/s00454-016-9793-3>, doi:10.1007/s00454-016-9793-3.
- [11] M. Develin. LP-orientations of cubes and crosspolytopes. *Adv. Geom.*, 4(4):459–468, 2004. URL: <https://doi.org/10.1515/advg.2004.4.4.459>, doi:10.1515/advg.2004.4.4.459.
- [12] R.B. Edman. *Diameter and Coherence of Monotone Path Graphs in Low Corank*. PhD thesis, UNIVERSITY OF MINNESOTA, May 2015. URL: <http://www-users.math.umn.edu/~reiner/edman-thesis.pdf>.
- [13] G. Kalai. Rigidity and the lower bound theorem. I. *Invent. Math.*, 88(1):125–151, 1987. URL: <https://doi.org/10.1007/BF01405094>, doi:10.1007/BF01405094.
- [14] J. McDonald. Fiber polytopes and fractional power series. *J. Pure Appl. Algebra*, 104(2):213–233, 1995. URL: [https://doi.org/10.1016/0022-4049\(94\)00129-5](https://doi.org/10.1016/0022-4049(94)00129-5).
- [15] P. McMullen. Triangulations of simplicial polytopes. *Beiträge Algebra Geom.*, 45(1):37–46, 2004.
- [16] J. Mihalisin and V. Klee. Convex and linear orientations of polytopal graphs. *Discrete Comput. Geom.*, 24(2-3):421–435, 2000. URL: <https://doi.org/10.1007/s004540010046>, doi:10.1007/s004540010046.
- [17] S. Murai and E. Nevo. On the generalized lower bound conjecture for polytopes and spheres. *Acta Math.*, 210(1):185–202, 2013. URL: <https://doi.org/10.1007/s11511-013-0093-y>, doi:10.1007/s11511-013-0093-y.
- [18] L. Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, Jul 2014. URL: <http://dx.doi.org/10.1016/j.aim.2014.02.035>, doi:10.1016/j.aim.2014.02.035.
- [19] L. Pournin. The asymptotic diameter of cyclohedra. *Israel J. Math.*, 219(2):609–635, Apr 2017. URL: <http://dx.doi.org/10.1007/s11856-017-1492-0>, doi:10.1007/s11856-017-1492-0.
- [20] V. Reiner and Y. Roichman. Diameter of graphs or reduced words and galleries. *Trans. Amer. Math. Soc.*, 365(5):2779–2802, 2013.
- [21] A. Schrijver. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons Ltd., 1986. A Wiley-Interscience Publication.
- [22] N.J.A. Sloane. The on-line encyclopedia of integer sequences. URL: <https://oeis.org/>.
- [23] G.M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. doi:10.1007/978-1-4613-8431-1.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ATHENS, 15784 ATHENS,
GREECE

E-mail address, C.A. Athanasiadis: `caath@math.uoa.gr`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, DAVIS, CA
95616, USA

E-mail address, J.A. De Loera: `deloera@math.ucdavis.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, DAVIS, CA
95616, USA

E-mail address, Z. Zhang: `zhenyangz@math.ucdavis.edu`