

FINDING SHORT PATHS ON SIMPLE POLYTOPES

ALEXANDER E. BLACK AND RAPHAEL STEINER

ABSTRACT. We prove that computing a shortest monotone path to the optimum of a linear program over a simple polytope is NP-hard, thus resolving a 2022 open question of De Loera, Kafer, and Sanità. As a consequence, finding a shortest sequence of pivots to an optimal basis with the simplex method is NP-hard. In fact, we show this is NP-hard already for fractional knapsack polytopes. By applying an additional polyhedral construction, we show that computing the diameter of a simple polytope is NP-hard, resolving a 2003 open problem by Kaibel and Pfetsch. Finally, on the positive side, we show that every polytope has a small, simple extended formulation for which a linear length path may be found between any pair of vertices in polynomial time building upon a result of Kaibel and Kukhareenko.

1. INTRODUCTION

Understanding the worst-case performance of the simplex method for linear programming across all choices of pivot rules is a longstanding research program established first with Dantzig’s 1947 invention, with foundational contributions made across theoretical computer science, operations research, and combinatorics communities. Breakthroughs on the positive side include the polynomial average case analysis of Borgwardt [7], the polynomial smoothed analysis by Spielman and Teng [41], and polynomial time versions for special families such as Orlin’s network simplex algorithm [37]. In the worst-case, the best known bound in terms of the number of inequalities and number of variables is subexponential originally due to Kalai [29] with follow up work improving the bounds in [23].

On the negative side, essentially all well-studied pivot rules are known to have superpolynomial worst case performance [31, 26, 4, 22, 34, 21, 29, 33, 3, 19, 18, 23, 12, 14, 5, 13]. Pivot rules can even encode hard problems during their execution [15, 17, 1]. Furthermore, the longstanding Hirsch conjecture that the diameter of the vertex-edge graph of a polytope is at most the number of inequalities minus the number of variables was disproven by Santos in [39]. This is a small sample of breakthroughs related to the nearly 80 years of consistent work dedicated to understanding this problem, yet fundamental questions remain open.

Given a polytope $P = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$, defined by a constraint matrix $A \in \mathbb{R}^{m \times d}$ and right-hand side $\mathbf{b} \in \mathbb{R}^m$, it has a set of feasible bases consisting of the set of linearly independent subsets B of rows of A of size d such that $A_B^{-1}\mathbf{b}_B \in P$, where A_B and \mathbf{b}_B denotes the matrix and right hand side restricted to the rows indexed by B . Two feasible bases B and B' are called adjacent if $|B \Delta B'| = 2$, which yields a graph associated to the polytope that we call the **feasible basis graph**. The simplex method solves a linear program by walking from basis to basis along the feasible basis graph. For a linear program $\max_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x}$, the step from a feasible basis B to a new feasible basis $B' = (B \setminus \{i\}) \cup \{j\}$ for some $i \in B, j \notin B$ is called **monotone** if the ray defined by

$$\{\mathbf{x} \in \mathbb{R}^d : A_{B \setminus \{i\}}\mathbf{x} = \mathbf{b}_{B \setminus \{i\}}, A_i\mathbf{x} \leq \mathbf{b}_i\}$$

DEPARTMENT OF MATHEMATICS, BOWDOIN COLLEGE

DEPARTMENT OF MATHEMATICS, ETH ZÜRICH

E-mail addresses: a.black@bowdoin.edu, raphaelmario.steiner@math.ethz.ch.

Research of R.S. supported by SNSF Ambizione Grant No. 216071.

is increasing with respect to \mathbf{c} .

A monotone move along a single edge in the feasible basis exchange graph is called a **pivot**, and the run-time of the simplex method depends on the number of pivots taken to reach an optimum as well as the time to compute each pivot.

There are several different pivot rules for the simplex method that have been studied. One that is particularly fundamental is the “omniscient pivot rule,” which simply chooses a shortest sequence of pivots to the optimum. Despite so many years of study, it is open whether this pivot rule may be computed in polynomial time. That is, given a linear program and a feasible initial basis, can one find a shortest monotone path in the feasible basis graph to an optimal basis in polynomial time? As our first main result, we prove that the answer is no assuming $P \neq NP$. Concretely we show that the following decision problem is NP-hard:

PIVOT-DISTANCE

Input: A linear program $\max_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x}$ defined by an objective vector $\mathbf{c} \in \mathbb{Q}^d$ and a polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$ defined by a matrix $A \in \mathbb{Q}^{m \times d}$ and a vector $\mathbf{b} \in \mathbb{Q}^m$, a feasible basis $B \subseteq [m]$ of P , and a number $k \in \mathbb{N}$.

Decision: Does there exist a monotone sequence of at most k pivots from B to a basis B^* corresponding to an optimal solution of the linear program?

In fact, we show a stronger statement related to another line of research of which the aforementioned hardness result is an immediate consequence (Theorem 1.3 below). Namely, a related graph to the feasible basis graph is the **graph** of the polytope defined by the vertices and edges of the polytope. Originally, in 1994, Frieze and Teng showed [20] that computing the diameter of the graph of a (possibly highly degenerate) input polytope P , called the **combinatorial diameter** and denoted $\text{diam}(P)$, is weakly NP-hard. Then much later in 2018, Sanità showed in [38] that computing the combinatorial diameter of the fractional matching polytope is strongly NP-hard. This result spurred a flurry of other results. For example, Wulf showed that computing the combinatorial diameter is Π_2 -complete [42]. Various hardness results are known in the setting [36, 10, 11, 9]. For special polytopes from algebraic combinatorics, hardness results are known but where the input is no longer the system of inequalities defining the polytope [2, 25]. Similar hardness results have also been shown in generalizations of polytope graphs [11, 6, 8].

However, until very recently, all known hardness results regarding shortest paths and diameters of polytopes with their inequality description as input were for **degenerate** polytopes for which the vertex-edge graph and feasible basis exchange graph do not coincide, since a single vertex may be represented by multiple feasible bases. Polytopes for which these two graphs coincide are called **simple**, and they correspond to polytopes for which every vertex is defined by precisely dimension many tight inequalities. In [11], De Loera, Kafer, and Sanità asked whether there exists a polynomial time algorithm to find shortest (monotone) paths in graphs of simple polytopes. Concurrently with and independently of the work presented in this paper, in a recent breakthrough Dorfer [16] showed that computing distances between pairs of vertices on the associahedron is NP-complete, which implies that computing shortest paths on simple polytopes is NP-hard. As our second main result, we prove the same result through a reduction from a different, arguably significantly simpler, class of polytopes (certain fractional knapsack polytopes, obtained by intersecting a hypercube with a carefully chosen halfspace). Formally, we show that the following decision problem is NP-hard:

k -DISTANCE ON SIMPLE POLYTOPES

Input: A simple polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$ defined by a matrix $A \in \mathbb{Q}^{m \times d}$ and a vector $\mathbf{b} \in \mathbb{Q}^m$, two vertices \mathbf{x}, \mathbf{y} of P and some number $k \in \mathbb{N}$.

Decision: Do \mathbf{x} and \mathbf{y} have distance at most k in the graph of P ?

Theorem 1.1. k -DISTANCE ON SIMPLE POLYTOPES is NP-hard.

While assuming $P \neq NP$, both Dorfer’s result [16] and Theorem 1.1 independently answer the aforementioned question of de Loera, Kafer and Sanità in the negative, there are two further implications of our result which are not implied by that of Dorfer [16]. First, one can easily find a path of length at most $O(\sqrt{m})$ between any pair of vertices on the associahedron in strongly polynomial time (see Lemma 2 of [40]), where m denotes the number of facets. Thus, Dorfer’s result could only imply at most that $O(\sqrt{m})$ -distance is NP-hard. In contrast, our argument shows that checking whether there exists a path of length at most $d - 1$ in a d -dimensional simple polytope with $2d + 1$ facets is NP-hard, so we have the following corollary:

Corollary 1.2. $(m - d - 2)$ -DISTANCE ON SIMPLE POLYTOPES is NP-hard.

In particular, unless $P = NP$, finding a path on a simple polytope shorter than the Hirsch bound $m - d$ by more than 2 cannot be done in polynomial time.

However, the second and most important distinction between our Theorem 1.1 and Dorfer’s work is the fact that our proof extends to the *monotone* setting. A path in the vertex-edge graph is called **monotone** if each step along the path increases the objective function. Under nondegeneracy, monotonicity corresponds exactly to pivoting in the simplex method, and hence this setting is particularly relevant in the optimization context and has been studied in several prior works. As our third main result, we show that the following problem is NP-hard:

k -MONOTONE-DISTANCE ON SIMPLE POLYTOPES

Input: A linear program $\max_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x}$ defined by an objective vector $\mathbf{c} \in \mathbb{Q}^d$ and a simple polytope $\mathbf{x} \in P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$ defined by a matrix $A \in \mathbb{Q}^{m \times d}$ and a vector $\mathbf{b} \in \mathbb{Q}^m$, a vertex \mathbf{x} of P and some number $k \in \mathbb{N}$.

Decision: Is there a monotone path of length at most k from \mathbf{x} to a \mathbf{c} -maximizer?

Theorem 1.3. $(m - d - 2)$ -MONOTONE DISTANCE ON SIMPLE POLYTOPES is NP-hard.

Hence, unlike the results of Dorfer in [16], our result implies the following, which is our first main result mentioned above.

Corollary 1.4. PIVOT-DISTANCE is NP-hard.

The proofs of Theorem 1.1, Corollary 1.2 and Theorem 1.3 will be presented in Section 2.

Our fourth main result concerns a related problem, which appears as Problem 10 in the 2003 survey on polyhedral computation by Kaibel and Pfetsch [28], where they ask for the complexity status of computing the combinatorial *diameter* of a simple polytope. This problem was also reiterated by Sanità [38] and Wulf [42]. By combining our aforementioned distance hardness result for simple polytopes with several additional ideas (that make up most of the technical work of this paper), we show that this problem, too, is NP-hard. Concretely, we address the following decision problem.

DIAMETER OF SIMPLE POLYTOPES

Input: A simple polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$ defined by a matrix $A \in \mathbb{Q}^{m \times d}$ and a vector $\mathbf{b} \in \mathbb{Q}^m$, and a number $k \in \mathbb{N}$.

Decision: Does $\text{diam}(P) \leq k$ hold?

Theorem 1.5. DIAMETER OF SIMPLE POLYTOPES is NP-hard.

Our approach to proving Theorem 1.5 is to reduce k -DISTANCE ON SIMPLE POLYTOPES to DIAMETER OF SIMPLE POLYTOPES. A priori, these are very different problems. To show that such a reduction nevertheless exists, we introduce and carefully analyze an intricate polyhedral construction (dubbed “cyclic siloing”) which can be applied to any simple input polytope P with a pair of vertices \mathbf{u}, \mathbf{v} to efficiently compute a larger simple polytope Q whose diameter can be expressed as the sum of the distance of \mathbf{u} and \mathbf{v} on P and another efficiently computable number K (cf. Theorem 3.11). Given access to an oracle for DIAMETER OF SIMPLE POLYTOPES, one can then efficiently compute the distance of \mathbf{u} and \mathbf{v} on P and solve k -DISTANCE ON SIMPLE POLYTOPES. Our construction takes inspiration from a similar such construction previously analyzed in the context of lower bounds for the shadow simplex method [5]. We believe that the constructions introduced in this paper are of independent interest and will find applications to other problems in computational polytope theory.

At a very high level, our constructions resemble those in the aforementioned work of Frieze and Teng in [20]. In that work, they first construct a simple polytope by taking a linear programming relaxation of a combinatorial optimization problem and show that computing the radius, i.e. the furthest distance away from a given vertex in the graph of that polytope, is NP-hard. They then apply a polyhedral construction to reduce diameter computation to the radius. However, our approach needs to overcome two major technical hurdles that stop Frieze and Teng’s approach from working in our settings. First, we need a different construction in order to show finding shortest paths is NP-hard instead of the radius. Our approach makes use of structural insights coming from understanding the geometric combinatorics of intersecting a hypercube with a halfspace, which was partly inspired by a similar construction in [10]. Second, the polyhedral construction used by Frieze and Teng to go from their hardness result for the radius to a hardness result for the diameter breaks simplicity. In particular, they iteratively cut off a vertex with a hyperplane (a process called **truncation**) and then take the convex hull with a new vertex close to that hyperplane (a process called **stacking**). Doing so repeatedly replaces a vertex with a tower separating that vertex from all of its neighbors. It breaks simplicity, because each vertex in the tower other than the top has more than d neighbors. We instead perform another procedure that preserves simplicity by only applying truncations iteratively. In part, our approach is a refinement of the use of truncations by Holt and Klee in their study of Hirsch-sharp polytopes in [24].

The key idea behind our construction is to add d truncations chosen purposefully to replace a vertex with a new vertex at which exactly the d new added inequalities from the truncation are tight. Furthermore, we choose these truncations to never cut off any other vertex of the original polytope. Then, by construction, the new vertex is always at least d steps away from any other vertex of the polytope before truncation. If we iterate this construction r times the resulting new vertex is $r(d - 1) + 1$ steps away from any of the original vertices. This construction thus mimics the effect of building a tower like Frieze and Teng while preserving simplicity. We call this tower a *cyclic silo*. To reduce distance computation to diameter computation, we replace the pair of vertices \mathbf{u} and \mathbf{v} we want to find a shortest path between with cyclic silos. In the resulting polytope, the pair of vertices at the tops of those towers will have distance precisely $2r(d - 1)$ higher than the distance of \mathbf{u} and \mathbf{v} in the original polytope. One then aims to show that for r sufficiently large, these vertices also attain the diameter. Therefore, computing the diameter of the resulting polytope allows one to find the distance between \mathbf{u} and \mathbf{v} in the original polytope, yielding the desired reduction. While this basic idea is approachable, to implement it in the desired way we navigate several intricate technical challenges. Namely, the choice of truncations, the efficient implementation of the construction, and a *precise* rather than approximate control of the diameter of the resulting polytope turn out to be quite challenging. For the details, we refer to Section 3, where we carefully describe and analyze our constructions and discuss their technical challenges and how we overcome them.

Finally, all of these results presented so far are negative and indicate obstacles towards finding polynomial time simplex methods conditional on $P \neq NP$. Our fifth and final contribution is

positive. In a recent work, Kaibel and Kukharenko [27] showed that one can reduce the well-known open problem of solving linear programming in strongly polynomial time (often referred to as *Smale's 9th problem* from his famous problem list for the 21st century) to instances where the feasible region forms a simple polytope with combinatorial diameter bounded linearly in the number of inequalities. To prove this result, they introduce an operation they call a **rock extension**, which creates from a simple d -dimensional polytope with m facets a closely related simple $(d + 1)$ -dimensional polytope with $m + 1$ facets and the remarkable aforementioned property that its diameter is at most $2(m - d)$. Furthermore, these rock extensions have a distinguished vertex $(o, 1)$ known as part of their construction. Their argument implies that there is a path from $(o, 1)$ to any other vertex of length at most $m - d$, certifying the aforementioned diameter bound. In their work, they did not study the complexity of finding such a path. Here we show the following:

Theorem 1.6. *Let Q be a rock extension with m facets in d dimensions. Let \mathbf{u} and \mathbf{v} be vertices of Q . Then one can find a path of length at most $2(m - d)$ from \mathbf{u} to \mathbf{v} in weakly polynomial time. If $(o, 1)$ is taken as part of the input, a path of length at most $2(m - d)$ may be found in strongly polynomial time, and a path from $(o, 1)$ to either vertex of length at most $m - d$ may also be found in strongly polynomial time.*

This theorem follows from a very simple analysis of the beautiful construction of Kaibel and Kukharenko in [27]. In Kukharenko's thesis [32], he showed that the solution of the linear program $\min_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x}$ is determined by the solution to the linear program $\min_{\mathbf{x} \in Q} (\mathbf{c}, c_z)^\top \mathbf{x}$, where c_z may be computed in strongly polynomial time from \mathbf{c} . In that case, the path of length $m - d$ computed from $(o, 1)$ to the optimum of the linear program is monotonically decreasing with respect to (\mathbf{c}, c_z) . Our argument here implies that such a path may be computed in strongly polynomial time assuming the optimum of the linear program is known. More generally, it may be computed in weakly polynomial time by finding the optimum of that linear program.

This gives a weak sense in which there is indeed a weakly polynomial time simplex method. Namely, as a Phase 1 procedure, one implements the strongly polynomial time reduction to compute the rock extension and initializes at a vertex $(o, 1)$. Then a path from $(o, 1)$ to the optimum of (\mathbf{c}, c_z) of length at most $m - d$ may be computed in weakly polynomial time. However, of course, this is somewhat circular, since to compute this path we need to know the optimum, for which one has to appeal to a linear programming solver (however, possibly one quite different from the simplex method). At the same time, this tells us that complexity theory is not the obstruction to a polynomial time version of the simplex method with this Phase 1 procedure. In fact, assuming there is a strongly polynomial time algorithm for linear programming using *any method*, there is a strongly polynomial algorithm to find a monotone path of length at most $m - d$ on a rock extension from $(o, 1)$ to the optimum of (\mathbf{c}, c_z) . In this sense, as a consequence of what we show here, there is a strongly polynomial time algorithm for linear programming if and only if there is a strongly polynomial time simplex method in a wide sense. This is a similar status to that of so-called *circuit augmentation schemes* for linear programming (a generalization of the simplex methods which allow moving along a more general set of directions) due to the very recent breakthrough result of Natura in his proof of the polynomial circuit diameter conjecture in [35]. His result demonstrates that if one can solve linear programming in strongly polynomial time using any method, then one can find a sequence of almost quadratically many circuit augmentations to the optimum of a linear program in strongly polynomial time.

2. SHORTEST PATHS

We prove Theorem 1.1, Corollary 1.2 and Theorem 1.3 by reduction from the following problem.

PARTITION WITH EVEN SUM

Input: A vector $(b_1, b_2, \dots, b_d) \in \mathbb{Z}_{>0}^d$ with $\beta := \sum_{i=1}^d b_i/2 \in \mathbb{Z}$.

Decision: Does there exist a subset $S \subseteq [d]$ such that

$$\beta = \sum_{i \in S} b_i = \sum_{j \in [d] \setminus S} b_j \quad ?$$

Note that PARTITION WITH EVEN SUM is equivalent to the usual Partition problem, as there is trivially no solution to Partition if $\beta \notin \mathbb{Z}$. Thus, it is NP-hard (cf. Problem 20 in [30]).

Given an instance $\mathbf{b} = (b_1, \dots, b_d) \in \mathbb{Z}_{>0}^d$ of PARTITION WITH EVEN SUM, we define an associated polytope $P_{\mathbf{b}}$ as follows, where we set $\beta := \sum_{i=1}^d b_i/2 \in \mathbb{Z}$:

$$P_{\mathbf{b}} := [0, 1]^{d+2} \cap \left\{ \mathbf{x} \in \mathbb{R}^{d+2} \mid \sum_{i=1}^d b_i x_i - \beta x_{d+1} + (\beta + 1/2)x_{d+2} \leq \beta + 1/4 \right\}$$

In what follows, whenever the vector \mathbf{b} is clear from context, we will denote by \mathbf{w} the vector obtained from \mathbf{b} by extending it with entries $-\beta$ and $\beta + 1/2$. That is, we define $\mathbf{w} := (b_1, b_2, \dots, b_n, -\beta, \beta + 1/2)$. Then, in particular,

$$P_{\mathbf{b}} = [0, 1]^{d+2} \cap \{ \mathbf{x} \in \mathbb{R}^{d+2} \mid \mathbf{w}^\top \mathbf{x} \leq \beta + 1/4 \}.$$

In the following, we prove several basic properties about the polytope $P_{\mathbf{b}}$, one of which is that it is a simple polytope. These properties allow us to reduce PARTITION WITH EVEN SUM to the problem of finding shortest paths between two vertices of $P_{\mathbf{b}}$.

Lemma 2.1. *For all $\mathbf{b} \in \mathbb{Z}_{>0}^d$ with $\sum_{i=1}^d b_i$ even, the polytope $P_{\mathbf{b}}$ is $(d+2)$ -dimensional and simple.*

Proof. One can observe directly from the definition that $P_{\mathbf{b}}$ contains $[0, 1/3]^{d+2}$ as a subset and is thus full-dimensional, i.e., of dimension $d+2$.

Since $[0, 1]^{d+2}$ is simple, any vertex of $P_{\mathbf{b}}$ contained in at least $d+3$ defining hyperplanes must be in the hyperplane:

$$H_{\mathbf{b}} = \{ \mathbf{x} \in \mathbb{R}^{d+2} : \sum_{i=1}^n b_i x_i - \beta x_{d+1} + (\beta + 1/2)x_{d+2} = \beta + 1/4 \}.$$

and also be a vertex of $[0, 1]^{d+2}$ and therefore be a $\{0, 1\}$ -vector in that hyperplane. Since $b_i, \beta \in \mathbb{Z}$ for all $i \in [n]$ and $\beta + 1/2 \in \mathbb{Z}[1/2]$, for any $S \subseteq [d+2]$ we have

$$\sum_{i \in S} w_i \in \mathbb{Z}[1/2],$$

where $\mathbb{Z}[1/2]$ denotes the set of rational numbers of the form p/q where $q \in \{1, 2\}$ and $p \in \mathbb{Z}$. Hence, since $\beta + 1/4 \notin \mathbb{Z}[1/2]$, we have

$$\sum_{i \in S} w_i \neq \beta + 1/4.$$

Therefore, $P_{\mathbf{b}}$ is simple. □

Next, we give an explicit combinatorial description of the vertices of $P_{\mathbf{b}}$. This description works in general for intersecting a hypercube with a halfspace, so no special assumptions on the vector \mathbf{w} are used in the proof of the next statement. In what follows, for a subset $S \subseteq [d+2]$, let $\mathbf{e}_S = \sum_{i \in S} \mathbf{e}_i$.




FIGURE 1. Depicted are the four different ways a hyperplane can slice two edges of a 2-face of a hypercube, which gives rise to the notions (c), (d), (e), and (f) of adjacency in Lemma 2.3. Note there are truly six ways this can occur, but the remaining two correspond to swapping i and j for edges of type (d) and (e).

Lemma 2.2. *The graph of $P_{\mathbf{b}}$ has vertex set $V_1 \cup V_2$, where*

$$V_1 = \left\{ \mathbf{e}_S \mid S \subseteq [d+2] \text{ s. t. } \sum_{i \in S} w_i \leq \beta \right\} \text{ and}$$

$$V_2 = \left\{ \mathbf{e}_S + \frac{\beta + 1/4 - \sum_{i \in S} w_i}{w_k} \mathbf{e}_k \mid \sum_{i \in S} w_i < \beta + 1/4 < \sum_{j \in S \cup \{k\}} w_j \text{ or } \sum_{i \in S} w_i > \beta + 1/4 > \sum_{j \in S \cup \{k\}} w_j \right\}.$$

Proof. Every vertex of $[0, 1]^{d+2}$ that is in the halfspace $\mathbf{w}^\top \mathbf{x} \leq \beta + 1/4$ remains a vertex, since

$$P_{\mathbf{b}} = [0, 1]^{d+2} \cap \{ \mathbf{x} \in \mathbb{R}^{d+2} : \mathbf{w}^\top \mathbf{x} \leq \beta + 1/4 \}$$

This encompasses every vertex in V_1 . Every other vertex is given by the intersection of the hyperplane $\{ \mathbf{x} \in \mathbb{R}^{d+2} : \mathbf{w}^\top \mathbf{x} = \beta + 1/4 \}$ with an edge of $[0, 1]^{d+2}$. All edges of the hypercube $[0, 1]^{d+2}$ are spanned between \mathbf{e}_S and $\mathbf{e}_S + \mathbf{e}_k$ for some $S \subseteq [d]$ and $k \in [d] \setminus S$. Then the claimed description of the remaining set of vertices V_2 is obtained by computing the intersection points of such edges with the hyperplane defined by $\mathbf{w}^\top \mathbf{x} = \beta + 1/4$. \square

In what remains, we will encode the vertices of $P_{\mathbf{b}}$ purely combinatorially by identifying vertices in V_1 with their corresponding sets S and vertices in V_2 with the unique pair (S, k) of a set $S \subseteq [d]$ and an element $k \in [d] \setminus S$ satisfying the inequalities in the definition of V_2 . We describe the graph using this terminology.

Lemma 2.3. *Let $\mathbf{b} \in \mathbb{Z}_{>0}^d$ be such that $\sum_{i=1}^d b_i$ is even and let \mathbf{u} and \mathbf{v} be two vertices of $P_{\mathbf{b}}$. Then \mathbf{u} and \mathbf{v} are adjacent on $P_{\mathbf{b}}$ if and only if*

- (a) $\mathbf{u} = S$ and $\mathbf{v} = T$ for some $S, T \subseteq [d+2]$ with $|S \Delta T| = 1$, or
- (b) $\mathbf{u} = S$ and $\mathbf{v} = (S, i)$ (for $i \notin S$) or $\mathbf{v} = (S \setminus \{j\}, j)$ (for $j \in S$) for some $S \subseteq [d+2]$, or
- (c) $\mathbf{u} = (S, i)$ and $\mathbf{v} = (S, j)$ for some $S \subseteq [d+2]$ and distinct $i, j \notin S$, or
- (d) $\mathbf{u} = (S, i)$ and $\mathbf{v} = (S \cup \{i\}, j)$ for some $S \subseteq [d+2]$ and distinct $i, j \notin S$, or
- (e) $\mathbf{u} = (S, i)$ and $\mathbf{v} = (S \cup \{j\}, i)$ for some $S \subseteq [d+2]$ and distinct $i, j \notin S$, or
- (f) $\mathbf{u} = (S \cup \{i\}, j)$ and $\mathbf{v} = (S \cup \{j\}, i)$ for some $S \subseteq [d+2]$ and distinct $i, j \notin S$.

Proof. Case (a) corresponds to adjacency on the hypercube, and two vertices in V_1 will be adjacent if and only if they are adjacent on the hypercube.

Since $P_{\mathbf{b}}$ is simple, the only inequalities which are tight at a vertex $\mathbf{u} \in V_1$ are those coming from the hypercube $[0, 1]^{d+2}$. It follows that \mathbf{u} is adjacent to a vertex $\mathbf{v} \in V_2$ if and only if \mathbf{v} is the intersection of an edge incident to the V_1 vertex on the hypercube $[0, 1]^{d+2}$ with the hyperplane $\mathbf{w}^\top \mathbf{x} = \beta + 1/4$. That is precisely what is captured by Case (b).

All of cases (c), (d), (e), and (f) correspond to adjacency between vertices in V_2 . Since the vertices in V_2 are those obtained as intersections of edges of the hypercube with the hyperplane $\mathbf{w}^\top \mathbf{x} = \beta + 1/4$, the edges between them correspond exactly to the one-dimensional intersections of the two-dimensional faces F of the hypercube $[0, 1]^{d+2}$ with the hyperplane $\mathbf{w}^\top \mathbf{x} = \beta + (1/4)$. Furthermore, the vertices in V_2 connected by such an edge are the intersection points of two of the edges of the hypercube contained in F with the hyperplane $\mathbf{w}^\top \mathbf{x} = \beta + (1/4)$.

Note that the vertex-sets of the two-faces of the hypercube are exactly of the form $S, S \cup \{i\}, S \cup \{j\}, S \cup \{i, j\}$ where $S \subseteq [d+2]$ and $i, j \notin S$, and their four connecting edges are $[S, S \cup \{i\}], [S, S \cup \{j\}], [S \cup \{i\}, S \cup \{i, j\}]$ as well as $[S \cup \{j\}, S \cup \{i, j\}]$. Thus, for a fixed such two-dimensional face F of the hypercube, there are up to $\binom{4}{2} = 6$ pairs of these four edges which could potentially be cut by the hyperplane and lead to adjacent vertices in V_2 on $P_{\mathbf{b}}$. Up to symmetry by swapping i and j , there are really only four types of adjacency that can arise between vertices of V_2 :

Case (c) corresponds to adjacency between the vertices of $P_{\mathbf{b}}$ coming from the edge from S to $S \cup \{i\}$ and the edge from S to $S \cup \{j\}$. Case (d) comes from the edges $[S, S \cup \{i\}]$ and $[S \cup \{i\}, S \cup \{i, j\}]$. Case (e) comes from the edges $[S, S \cup \{i\}]$ and $[S \cup \{j\}, S \cup \{i, j\}]$. Finally case (f) comes from the pair $[S \cup \{i\}, S \cup \{i, j\}]$ and $[S \cup \{j\}, S \cup \{i, j\}]$. See Figure 1 for a visualization of these cases. \square

It turns out that the relevance of this characterization comes down to the following insight. If $S \subseteq T$ and (S, i) and (T, i) are both vertices, then the shortest a path between (S, i) and (T, i) in the graph of $P_{\mathbf{b}}$ could potentially be is $|T| - |S|$ by adding one element of T to S at a time. What we will prove is that if a shortest path of length $|T| - |S|$ exists, then it must be of that form, and that checking if such a path exists is NP-hard by a reduction to PARTITION WITH EVEN SUM.

Lemma 2.4. *Let $\mathbf{b} = (b_1, \dots, b_d) \in \mathbb{Z}_{>0}^d$ such that $\sum_{i=1}^d b_i$ is even. Then*

- $(\emptyset, d+2)$ and $([d+1], d+2)$ are vertices of $P_{\mathbf{b}}$.
- *The shortest path between $(\emptyset, d+2)$ and $([d+1], d+2)$ is of length at most $d+1$ if and only if there exists a solution to PARTITION WITH EVEN SUM with instance \mathbf{b} .*

Proof. Since

$$\mathbf{w}^\top \mathbf{e}_\emptyset = 0 \leq \beta + 1/4 < \beta + 1/2 = \mathbf{w}^\top \mathbf{e}_{d+2},$$

and

$$\mathbf{w}^\top \mathbf{e}_{[d+1]} = \sum_{i=1}^n b_i - \beta = 2\beta - \beta = \beta < \beta + 1/4 < \beta + \beta + 1/2 = \mathbf{w}^\top \mathbf{e}_{[d+2]},$$

$(\emptyset, d+2)$ and $([d+1], d+2)$ are vertices of $P_{\mathbf{b}}$ by the characterization of the vertices in Lemma 2.2. This proves the first item of the lemma.

In the following, we will determine precisely the structure of the paths of length at most $d+1$ from $(\emptyset, d+2)$ to $([d+1], d+2)$ on $P_{\mathbf{b}}$, which will then yield the second item of the lemma.

From the characterization of edges in Lemma 2.3, moving along any edge of $P_{\mathbf{b}}$ can only increase the size of the support of the current vertex by at most 1. Thus, any path between $(\emptyset, d+2)$ and $([d+1], d+2)$ of length at most $d+1$ must in fact be of length *exactly* $d+1$ and each step along the path must increase the size of the support by exactly 1. The only edge types from Lemma 2.3

that increase the size of the support when we move along them starting from a vertex of the form (S, i) are those of type (d) and (e). Since our path starts at $(\emptyset, d+2)$ and since moving along type (d) and (e) edges we stay within vertices of type (S, i) , it follows that any path of length $d+1$ from $(\emptyset, d+2)$ to $([d+1], d+2)$ on $P_{\mathbf{b}}$ must only use type (d) and (e) edges. We now claim that any such path in fact only uses type (e) edges. Indeed, towards a contradiction suppose it uses some type (d) edge and consider the earliest such edge along the path when starting from $(\emptyset, d+2)$. Since edges of type (e) always move from a vertex of the form (S, i) to a vertex of the form (S', i) and hence always preserve the “second coordinate”, and since we start from the vertex $(\emptyset, d+2)$, the first edge of type (d) along the path must then start at a vertex of the form $(S, d+2)$ for some $S \subseteq [d+1]$ and go to $(S \cup \{d+2\}, j)$ for some $j \notin S$ distinct from $d+2$. To have a total length of $d+1$, we would then need to reach $([d+1], d+2)$ from $(S \cup \{d+2\}, j)$ using only type (d) and (e) edges which increase the support. However, this is impossible, since $S \cup \{d+2\}$ contains the element $d+2$ while $[d+1]$ does not, and since any type (d) and (e) edges used after will have to increase the support and hence preserve that $d+2$ is an element of the set in the tuple. Hence, we have reached the desired contradiction, and it follows that indeed all edges used along the path must be support-increasing edges of type (e).

Recalling the definition of type (e) edges, it now follows that every path of length at most $d+1$ from $(\emptyset, d+2)$ to $([d+1], d+2)$ on $P_{\mathbf{b}}$ must be of the form $(S_0, d+2), (S_1, d+2), \dots, (S_{d+1}, d+2)$ where

$$\emptyset = S_0 \subsetneq S_1 \subsetneq S_2, \dots \subsetneq S_{d+1} = [d+1]$$

are such that $(S_i, d+2)$ is a vertex of $P_{\mathbf{b}}$ and $S_i = S_{i-1} \cup \{k\}$ for some $k \in [d+1] \setminus S_{i-1}$ for all $1 \leq i \leq d+1$.

We claim that such a sequence of sets exists (and hence the distance from $(\emptyset, d+2)$ to $([d+1], d+2)$ is at most $d+1$) if and only if there is a solution to PARTITION WITH EVEN SUM. Suppose first that such a sequence of sets exists. Let i be minimal such that $d+1 \in S_i$. Then, since $(S_i, d+2)$ is a vertex of $P_{\mathbf{b}}$,

$$\begin{aligned} -\beta + \sum_{j \in S_{i-1}} b_j &= \mathbf{w}^\top \mathbf{e}_{d+1} + \mathbf{w}^\top \mathbf{e}_{S_{i-1}} \\ &= \mathbf{w}^\top \mathbf{e}_{S_i} \\ &\leq \beta + 1/4 \\ &\leq \mathbf{w}^\top \mathbf{e}_{S_i \cup \{d+2\}} \\ &= \mathbf{w}^\top \mathbf{e}_{d+2} + \mathbf{w}^\top \mathbf{e}_{S_i} \\ &= \beta + 1/2 - \beta + \sum_{j \in S_{i-1}} b_j \\ &= 1/2 + \sum_{j \in S_{i-1}} b_j. \end{aligned}$$

In particular, $\beta + 1/4 \leq 1/2 + \sum_{j \in S_{i-1}} b_j$, so

$$\sum_{j \in S_{i-1}} b_j \geq \beta - 1/4.$$

Similarly, since $(S_{i-1}, d+2)$ is also a vertex of $P_{\mathbf{b}}$ and since S_{i-1} does not contain $d+1$ by definition of i , we have:

$$\sum_{j \in S_{i-1}} b_j = \mathbf{w}^\top \mathbf{e}_{S_{i-1}} \leq \beta + 1/4.$$

It follows that

$$\beta - 1/4 \leq \sum_{j \in S_{i-1}} b_j \leq \beta + 1/4$$

Since $b_i \in \mathbb{Z}$ for all $i \in [n]$ and $\beta \in \mathbb{Z}$, it follows that $\sum_{j \in S_{i-1}} b_j = \beta$. Therefore, in that case, PARTITION WITH EVEN SUM has a solution.

Suppose instead that PARTITION WITH EVEN SUM has a solution. Up to reordering we may without loss of generality assume then that

$$\sum_{i=1}^k b_i = \beta.$$

For each $j \in [d+1]$, define

$$S_j = \begin{cases} \{1, \dots, j\} & \text{if } j \leq k \\ \{1, \dots, j-1\} \cup \{d+1\} & \text{if } j \geq k+1. \end{cases}$$

Then it suffices to show that $(S_j, d+2)$ is a vertex for each $j \in [d+1]$. If $j \leq k$, then

$$\mathbf{w}^\top \mathbf{e}_{S_j} = \sum_{i \in S_j} b_i = \sum_{i=1}^j b_i \leq \sum_{i=1}^k b_i \leq \beta < \beta + 1/4 < \beta + 1/2 + \sum_{i \in S_j} b_i = \mathbf{w}^\top \mathbf{e}_{S_j \cup \{d+2\}}.$$

Hence, $(S_j, d+2)$ is a vertex in that case.

If $j = k+1$, then

$$\mathbf{w}^\top \mathbf{e}_{S_j} = -\beta + \sum_{i=1}^k b_i = 0 < \beta + 1/4 < \beta + 1/2 + 0 = w_{d+2} + \sum_{i \in S_j} w_i = \mathbf{w}^\top \mathbf{e}_{S_j \cup \{d+2\}}.$$

Finally, suppose that $j \geq k+2$. Then $\sum_{i \in S_j} w_i \geq \sum_{i \in S_{k+1}} w_i = 0$, and

$$\mathbf{w}^\top \mathbf{e}_{S_j} = \sum_{i \in S_j} w_i \leq \sum_{i \in [d+1]} w_i = \beta < \beta + 1/4 < \beta + 1/2 \leq \beta + 1/2 + \sum_{i \in S_j} w_i = \mathbf{w}^\top \mathbf{e}_{S_j \cup \{d+2\}}.$$

Hence, in all cases, $(S_j, d+2)$ is a vertex and so there is a path from $(\emptyset, d+2)$ to $([d+1], d+2)$ of length at most $d+1$ of the desired form. This concludes the proof of the equivalence claimed in the second item of the lemma. \square

This lemma yields Theorem 1.1 as an immediate consequence.

Proof of Theorem 1.1 and Corollary 1.2. By Lemma 2.4, one can solve PARTITION WITH EVEN SUM by deciding whether the distance between two specified vertices of $P_{\mathbf{b}}$ is at most $d+1$. Since $P_{\mathbf{b}}$ may be constructed from \mathbf{b} in polynomial time, and its encoding length (in inequality description) is polynomially tied to the encoding length of the input \mathbf{b} of PARTITION WITH EVEN SUM, and since $P_{\mathbf{b}}$ is simple by Lemma 2.1, it follows that k -DISTANCE ON SIMPLE POLYTOPES is NP-hard when setting $k = d+1$. Since this equals $(2d+5) - (d+2) - 2$ which is the number of defining inequalities of $P_{\mathbf{b}}$ minus the dimension of $P_{\mathbf{b}}$ minus two, this also proves Corollary 1.2. \square

We next extend our result to the monotone setting and prove Theorem 1.3. Recall that $\mathbf{e}_S = \sum_{i \in S} \mathbf{e}_i$ for any $S \subseteq [n]$.

Lemma 2.5. *Let $\mathbf{b} = (b_1, \dots, b_d) \in \mathbb{Z}_{>0}^d$ such that $\sum_{i=1}^d b_i = 2\beta$ is even. Let $\varepsilon := \frac{1}{5\beta}$ and $\mathbf{c} = \mathbf{e}_{[d+1]} + \varepsilon \mathbf{e}_{d+2}$. Then*

- $([d+1], d+2)$ is the unique \mathbf{c} -maximum.

- If $S \subsetneq T \subseteq [d+1]$, then $(S, d+2)$ has objective value less than $(T, d+2)$.

Proof. By Lemma 2.4, $([d+1], d+2)$ is a vertex of $P_{\mathbf{b}}$. Furthermore,

$$\mathbf{w}^\top \mathbf{e}_{[d+2]} = \sum_{i=1}^{d+2} \mathbf{w}_i = \sum_{i=1}^n b_i - \beta + (\beta + 1/2) = 2\beta + 1/2 > \beta + 1/4.$$

Hence, $\mathbf{e}_{[d+2]} \notin P_{\mathbf{b}}$. Let \mathbf{v} be the vector corresponding to $([d+1], d+2)$. Then by Lemma 2.2,

$$\mathbf{v} = \mathbf{e}_{[d+1]} + \alpha \mathbf{e}_{d+2}$$

for some $\alpha > 0$. It follows that

$$\mathbf{c}^\top \mathbf{v} = d + 1 + \varepsilon \alpha.$$

Any other vertex is of the form $\mathbf{e}_S + \alpha' \mathbf{e}_i$, where $S \subsetneq [d+2]$, $i \notin S$ and $0 \leq \alpha' < 1$. In particular, by Lemma 2.2, if $\alpha' > 0$, then

$$\alpha' = \frac{\beta + 1/4 - \sum_{j \in S} w_j}{w_i}.$$

Suppose first that $i \neq d+2$. Then $S \cap [d+1]$ is a proper subset of $[d+1]$, and w_i is integral with $|w_i| \leq \beta$. Furthermore, $4w_j \in \mathbb{Z}$ for each $j \in S$. Note that $\alpha' = 0$ or $0 < \alpha' < 1$. In the latter case, we have

$$\alpha' = \frac{\beta + 1/4 - \sum_{j \in S} w_j}{w_i} = \frac{|4\beta + 1 - \sum_{j \in S} 4w_j|}{4|w_i|} \leq 1 - \frac{1}{4|w_i|} < 1 - \frac{1}{5\beta}.$$

Since $\varepsilon = \frac{1}{5\beta}$, it follows that $\alpha' + \varepsilon < 1$ (and this clearly also holds in the case $\alpha' = 0$). Consequently,

$$\mathbf{c}^\top (\mathbf{e}_S + \alpha' \mathbf{e}_i) \leq (|S \cap [d+1]| + \varepsilon) + \alpha' \leq d + \alpha' + \varepsilon < d + 1 + \varepsilon \alpha = \mathbf{c}^\top \mathbf{v}.$$

For the second case, suppose, $i = d+2$ (and hence $S \subseteq [d+1]$). We then obtain

$$\mathbf{c}^\top (\mathbf{e}_S + \alpha' \mathbf{e}_{d+2}) = |S| + \varepsilon \alpha' \leq \max\{d + \varepsilon \alpha', d + 1\} < d + 1 + \varepsilon \alpha = \mathbf{c}^\top \mathbf{v}$$

where in the second step we used that $\mathbf{e}_S + \alpha' \mathbf{e}_i \neq \mathbf{v}$ meaning that $S \neq [d+1]$ or $\alpha' = 0$. Thus, \mathbf{v} is the unique \mathbf{c} -maximizer, as desired. This concludes the proof of the first item of the lemma.

For the second item, consider any $S \subsetneq T \subseteq [d+1]$. Then for any $0 < \alpha < 1$ and $0 < \beta < 1$, $\mathbf{c}^\top (\mathbf{e}_S + \alpha \mathbf{e}_{d+2}) = |S| + \varepsilon \alpha < |S| + 1 \leq |T| < \mathbf{c}^\top (\mathbf{e}_T + \beta \mathbf{e}_{d+2})$. It follows that $(S, d+2)$ has lower objective value than $(T, d+2)$, as desired. \square

This lemma allows us to immediately extend our result to the monotone setting.

Proof of Theorem 1.3. Note that by Lemma 2.4 one can solve PARTITION WITH EVEN SUM by checking whether a path from $(\emptyset, d+2)$ to $([d+1], d+2)$ of length at most $d+1$ exists. By Lemma 2.5, $([d+1], d+2)$ is the optimum of the objective \mathbf{c} from the statement of Lemma 2.5. From the proof of Lemma 2.4, a path of length at most $d+1$ exists if and only if a path exists of the form

$$(S_0, d+2), (S_1, d+2), \dots, (S_{d+1}, d+2),$$

where $S_i \subsetneq S_{i+1}$ for each $i \in [0, d]$. By Lemma 2.5, that path is increasing with respect to \mathbf{c} . Hence, a \mathbf{c} -increasing path of length at most $d+1$ from $(\emptyset, d+2)$ to $([d+1], d+2)$ exists if and only if there is a path of length at most $d+1$ from $(\emptyset, [d+2])$ to $([d+1], d+2)$. This is true if and only if there is a solution to PARTITION WITH EVEN SUM. Hence, the same reduction works, showing that monotone distance on simple polytopes is NP-hard. \square

3. DIAMETERS

Throughout this section, we only consider simple polytopes of dimension $d \geq 3$. Furthermore, we will always assume that we only work with irredundant inequality descriptions of our polytopes. In particular, we assume that every vertex of our polytopes satisfies exactly d of the defining inequalities with equality. Furthermore, we will always assume that the entries of the matrix and the right-hand side defining our polytope have rational entries. This is crucial for some of our statements and lemmas, even though it will not always be explicitly mentioned. We will also throughout use the notation $d_P(\mathbf{u}, \mathbf{v})$ to denote the (combinatorial) distance between two vertices \mathbf{u}, \mathbf{v} in the graph of a polytope P .

Suppose we are given a d -dimensional simple polytope P described by m inequalities and a vertex \mathbf{v} of P . We can then compute the d neighbors $\mathbf{v}_1, \dots, \mathbf{v}_d$ of \mathbf{v} on P and “cut \mathbf{v} off” from each of these neighbors by adding a single new inequality. Namely, we may compute the mid-points $\mathbf{m}_i := \frac{\mathbf{v} + \mathbf{v}_i}{2}, i = 1, \dots, d$ of the incident edges of \mathbf{v} and then compute the unique hyperplane passing through the points $\mathbf{m}_1, \dots, \mathbf{m}_d$. It is easy to see that this hyperplane separates \mathbf{v} from all other vertices of the polytope. Finally, we add a new inequality to P describing the halfspace of this hyperplane which does not contain \mathbf{v} . This operation is called **truncation** and yields a new simple polytope $T(P, \mathbf{v})$. The vertices of $T(P, \mathbf{v})$ are exactly those of P except \mathbf{v} plus the d additional vertices $\mathbf{m}_1, \dots, \mathbf{m}_d$.

Since we will later need it for our reductions, let us record the following useful statement about computing and encoding repeated truncations of polytopes.

Lemma 3.1. *Suppose P is a simple d -dimensional polytope with rational irredundant inequality description and with bit-encoding length L , and let $r \in \mathbb{N}$. Suppose we are given as input P as well as a sequence of r vertices which are revealed to us during the process one at a time, and each time a new vertex is revealed to us we have to perform a truncation at this vertex. Then an inequality description of the final polytope Q (obtained after performing the sequence of r truncations) with encoding length $\text{poly}(L, r)$ can be computed in time $\text{poly}(L, r)$.*

Proof. Recall that we assume that the coefficients and constants of the inequalities defining P are rational numbers, and hence the same is true for all vertices of P . Since taking midpoints keeps the coordinates of vectors rational, all new vertices constructed during the process are rational. In particular, the vertices of Q are rational.

To start, we bound the encoding-lengths of any vertices appearing in the process polynomially in L and in r . Consider first the vertices of P . Each such vertex is the solution of a linear equation system over a $d \times d$ invertible submatrix of the constraint matrix. Hence, by a standard application of Cramer’s rule and Hadamard’s inequality, the bit-encoding length of each vertex of P is upper-bounded by $O(d^2L)$ each. Next observe that by definition of truncation, each new vertex obtained in one of the r truncations to obtain Q can be written as a convex combination of vertices of P where all coefficients are in $\{0, \frac{1}{2^r}, \dots, \frac{2^r-1}{2^r}, 1\}$. Moreover, note that in this convex combination we only need to consider vertices of P that at some point are either picked as the truncated vertex or a neighbor of it. Since in constructing P , we certainly consider at most $(d+1) \cdot r$ such vertices of P , it follows that each new vertex constructed at some point of the process is a convex combination of at most $(d+1)r$ vertices of P with coefficients in $\{0, \frac{1}{2^r}, \dots, \frac{2^r-1}{2^r}, 1\}$. Recall that each vertex of P has encoding length $O(d^2L)$ and in particular each entry of a vector in P has numerator and denominator at most $2^{O(d^2L)}$. Hence, every entry of any vertex computed in the construction process for Q can be written in the form

$$\sum_{i=1}^{(d+1)r} \frac{\alpha_i}{2^r} \cdot \frac{a_i}{b_i},$$

where $\alpha_i \in \{0, 1, \dots, 2^r\}$ and $|a_i|, |b_i| \leq 2^{O(d^2L)}$ for every i . This equals $\frac{p}{q}$, where

$$|q| = 2^r \prod_{i=1}^{(d+1)r} |b_i| \leq 2^{r+(d+1)r \cdot O(d^2L)} = 2^{O(d^3Lr)},$$

and

$$|p| = \left| \sum_{i=1}^{(d+1)r} \alpha_i a_i \prod_{j \neq i} b_j \right| \leq \sum_{i=1}^{(d+1)r} 2^r \cdot 2^{(d+1)r \cdot O(d^2L)} = 2^{O(d^3Lr)}.$$

Hence, every vertex computed in the construction process for Q has encoding length at most $d \cdot O(d^3Lr) = O(d^4Lr) = O(L^3r)$, where we used that $L \geq d^2$ in the last step.

It remains to argue that we can compute an inequality description of Q with encoding length $\text{poly}(L, r)$ in time $\text{poly}(L, r)$. Let $\mathbf{s}_1, \dots, \mathbf{s}_r$ be the vertices revealed to us one by one during the process, and suppose that for some $1 \leq j \leq r$ we have already computed an inequality description of the polytope P^{j-1} obtained from P after performing truncations at vertices $\mathbf{s}_1, \dots, \mathbf{s}_{j-1}$ in this order, and now a new vertex \mathbf{s}_j of P^{j-1} is revealed to us, at which we are supposed to perform the next truncation. To obtain the inequality description of the next polytope $P^j = T(P^{j-1}, \mathbf{s}_j)$, we proceed as follows:

- We first compute all the d neighbors $\mathbf{u}_1^j, \dots, \mathbf{u}_d^j$ of \mathbf{s}_j on P^{j-1} , in polynomial time in the encoding length of P^{j-1} . Concretely, we can do this by trying out all possible base exchanges at \mathbf{s}_j and thus solving up to $d(m(P^{j-1}) - d) \leq d(m(P) + r - d) \leq d(L + r - d)$ linear equation systems (here $m(P^{j-1}), m(P)$ denote the number of inequalities describing P^j and P , respectively) of size $d \times d$ whose coefficients form a submatrix of the constraint matrix of P^{j-1} . Hence, this can be executed in polynomial time in the encoding length of P^{j-1} .
- We then compute the d midpoints $\frac{\mathbf{u}_1^j + \mathbf{s}_j}{2}, \dots, \frac{\mathbf{u}_d^j + \mathbf{s}_j}{2}$. This can be done in polynomial time in the encoding lengths of $\mathbf{u}_1^j, \dots, \mathbf{u}_d^j$ and \mathbf{s}_j , and hence, by what we argued about the encoding lengths of these vectors above, in time $\text{poly}(L, r)$.
- Finally, we compute the coefficients of the one new inequality to be added to P^{j-1} to obtain P^j , i.e., a vector \mathbf{w} and some $\alpha \in \mathbb{R}$ such that the hyperplane $\mathbf{w}^\top \mathbf{x} = \alpha$ passes through all the midpoints $\frac{\mathbf{u}_1^j + \mathbf{s}_j}{2}, \dots, \frac{\mathbf{u}_d^j + \mathbf{s}_j}{2}$. To do so, it suffices to find the (up to scaling unique) non-trivial solution to the $d \times (d+1)$ -sized homogeneous linear equation system whose row vectors are obtained from $\frac{\mathbf{u}_1^j + \mathbf{s}_j}{2}, \dots, \frac{\mathbf{u}_d^j + \mathbf{s}_j}{2}$ by appending 1-s at the end. Since by what we showed above also the encoding lengths of $\frac{\mathbf{u}_1^j + \mathbf{s}_j}{2}, \dots, \frac{\mathbf{u}_d^j + \mathbf{s}_j}{2}$ are polynomial in L and in r , it follows again by using Cramer's rule and Hadamard's inequality, that we can compute a desired non-trivial solution (\mathbf{w}, α) to this linear system whose encoding length is bounded by $\text{poly}(L, r)$, and of course, it can be also computed in time $\text{poly}(L, r)$ by solving the linear system.

By the last point, we find that each of the r truncation steps increases the encoding length of the polytope by at most $\text{poly}(L, r)$, and hence each of the r polytopes $P^1, P^2, \dots, P^r = Q$ built in the process has encoding length at most $\text{poly}(L, r) \cdot r = \text{poly}(L, r)$. With this knowledge, it follows that each of the three steps above can be executed in time $\text{poly}(L, r)$ for each of the r truncations, and hence computing the inequality description of the final polytope Q also can be done in time $\text{poly}(L, r) \cdot r = \text{poly}(L, r)$. This establishes the desired statements and concludes the proof of the lemma. \square

As an organizational tool to keep track of the impact of truncating repeatedly on the combinatorial structure of the polytope, we use a generating function. Namely, let a d -dimensional simple polytope

P with m inequalities labeled by numbers $1, 2, \dots, m$ be given. Let $\mathcal{B} \subseteq \binom{[m]}{d}$ denote the set of feasible bases of P . Consider the polynomial ring $\mathbb{Z}[x_1, x_2, \dots, x_m]$, and for a subset $S \subseteq [m]$ let us denote $\mathbf{x}^S := \prod_{i \in S} x_i$. We now define the **generating function of feasible bases** of P by

$$f_P(\mathbf{x}) = \sum_{B \in \mathcal{B}} \mathbf{x}^B.$$

Now consider a vertex \mathbf{v}^* of P . Now when we truncate P at \mathbf{v}^* to obtain $T(P, \mathbf{v}^*)$, we will associate a new variable x_{m+1} in the polynomial ring with the added inequality. Our next lemma precisely describes how truncation changes the generating function.

Lemma 3.2. *Let P be a d -dimensional simple polytope with m facets labeled $1, 2, \dots, m$. Let $\mathcal{B} \subseteq \binom{[m]}{d}$ denote the set of feasible bases of P . Let \mathbf{v}^* be a vertex of P and B^* the corresponding feasible basis. Then we have*

$$f_{T(P, \mathbf{v}^*)}(\mathbf{x}) = \sum_{B \in \mathcal{B} \setminus \{B^*\}} \mathbf{x}^B + \sum_{i \in B^*} \mathbf{x}^{B^* \setminus \{i\}} x_{m+1} = f_P(\mathbf{x}) - \mathbf{x}^{B^*} + \sum_{i \in B^*} \mathbf{x}^{B^* \setminus \{i\}} x_{m+1}.$$

Proof. By definition, $T(P, \mathbf{v}^*)$ has one additional new inequality and so has $m + 1$ inequalities. The truncation only removes precisely one vertex, namely \mathbf{v}^* . Thus, each feasible basis of P other than B^* remains a feasible basis of Q . Furthermore, a new feasible basis is also added for each new vertex. There are exactly d new vertices in $T(P, \mathbf{v}^*)$ compared to P , namely those corresponding to the intersection of the d edges of P incident with \mathbf{v} with the hyperplane defining the new inequality we added. Therefore, the feasible bases corresponding to new vertices are precisely of the form $(B^* \setminus \{i\}) \cup \{m+1\}$, where i ranges through the elements of B^* . Putting these facts together yields the desired formula for the generating function of $T(P, \mathbf{v}^*)$. \square

Next, we would like to understand the effect of *repeated* truncation on the generating function. As a first step, it will thus be convenient for us to reformulate the expression in Lemma 3.2 in the case that the polytope P to which we apply the truncation already comes with two classes of inequalities, namely m “old” inequalities associated with variables x_1, \dots, x_m in the generating function, and $k - 1$ “new” inequalities associated with $k - 1$ new variables y_1, y_2, \dots, y_{k-1} , which we think of arising from $k - 1$ previous truncations. Here, we will use the convention that variable y_i corresponds to the inequality added in the i th previous truncation. In particular, in the following y_1, \dots, y_k will take the role of the variables x_{m+1}, \dots, x_{m+k} in the previous formulation of Lemma 3.2. Our next corollary is simply a restatement of Lemma 3.2 in this new set-up. To simplify notation we define, for each $S \subseteq [m]$ and $T \subseteq [k]$, the following shorthand:

$$\mathfrak{C}_k(\mathbf{x}^S \mathbf{y}^T) := \sum_{t \in T} \mathbf{x}^S \mathbf{y}^{(T \setminus \{t\}) \cup \{k\}} + \sum_{s \in S} \mathbf{x}^{S \setminus \{s\}} \mathbf{y}^{T \cup \{k\}}.$$

Corollary 3.3. *Let P be a d -dimensional simple polytope with facets in two classes of size m and $k - 1$, which are labeled $1, 2, \dots, m$ and $1, 2, \dots, (k - 1)$, respectively. Let $\mathcal{B} \subseteq 2^{[m]} \times 2^{[k-1]}$ denote the set of feasible bases of P . Let $B^* = (S, T)$ be a feasible basis, with corresponding vertex \mathbf{v}^* . Then we have*

$$f_{T(P, \mathbf{v}^*)}(\mathbf{x}, \mathbf{y}) = f_P(\mathbf{x}, y_1, \dots, y_{k-1}) - \mathbf{x}^S \mathbf{y}^T + \mathfrak{C}_k(\mathbf{x}^S \mathbf{y}^T).$$

Next, we will use this observation to describe the generating function of a new polytope constructed by a specific sequence of d iterated truncations.

Lemma 3.4. *Let P be a d -dimensional simple polytope with facets labeled $1, 2, \dots, m$ such that $[d]$ is a feasible basis. Then, given P (in inequality description) and this feasible basis as input, we can,*

in polynomial time in the encoding length of P , construct a new simple polytope S with $m + d$ facets, with corresponding variables $x_1, \dots, x_m, y_1, \dots, y_d$ and with feasible basis generating function

$$f_S(\mathbf{x}, \mathbf{y}) = f_P(\mathbf{x}) - \mathbf{x}^{[d]} + \mathbf{y}^{[d]} + \sum_{k=0}^{d-1} \left(\sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+2}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]} \right).$$

Proof. By definition, we have for all $0 \leq k < d$:

$$\begin{aligned} \mathfrak{C}_{k+1}(\mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k]}) &= \sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+1}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]} \\ &= \mathbf{x}^{[d] \setminus [k+1]} \mathbf{y}^{[k+1]} + \sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+2}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]}. \end{aligned}$$

To construct the polytope S , we will construct a sequence T^0, \dots, T^d of polytopes, where we initialize $T^0 := P$, and for $k = 0, \dots, d-1$ we define $T^{k+1} := T(T^k, \mathbf{v}^k)$, where \mathbf{v}^k is defined as the vertex of T^k corresponding to the monomial $\mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k]}$. Note that this is always a well-defined operation, since by the above calculation and by Corollary 3.3, we can see that if the generating function of T^k contains the monomial $\mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k]}$ then the generating function of the resulting polytope T^{k+1} after truncation will contain the monomial $\mathbf{x}^{[d] \setminus [k+1]} \mathbf{y}^{[k+1]}$. Finally, we set $S := T^d$. Then, by Corollary 3.3 and our above computation, the overall sum of the monomials added to the generating function across the whole procedure amounts to

$$\sum_{k=0}^{d-1} \left(\mathbf{x}^{[d] \setminus [k+1]} \mathbf{y}^{[k+1]} + \sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+2}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]} \right).$$

Similarly, the overall sum of the monomials subtracted from the generating function across the whole procedure (cf. Corollary 3.3) equals $\sum_{k=0}^{d-1} \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k]}$. Subtracting that off yields

$$f_S(\mathbf{x}, \mathbf{y}) = f_P(\mathbf{x}) - \mathbf{x}^{[d]} + \mathbf{y}^{[d]} + \sum_{k=0}^{d-1} \left(\sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+2}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]} \right),$$

as desired. Finally, note that since S arises from P by a sequence of d truncations, by Lemma 3.2 we can compute an inequality description of the final polytope S with encoding length polynomial in the encoding length of P , in polynomial time in the encoding length of P . \square

In the remainder of the paper, the construction of a simple polytope S with $m + d$ facets starting from an ordered feasible basis of a simple polytope P with m facets as described in Lemma 3.4 will be referred to as **siloing** due to its interpretation as building a tower to create an isolated vertex as depicted visually in the normal fan in Figure 2. Note that the assumption that the feasible basis corresponding to the vertex where we start truncating equals $[d]$ is not of any essence, and thus, more generally, can be applied to any *ordered* feasible basis of a polytope by relabeling. Thus, given a simple polytope P , a vertex \mathbf{v} of P and a linear ordering $b_1 \prec b_2 \prec \dots \prec b_d$ on the elements of the feasible basis $\{b_1, \dots, b_d\}$ associated with \mathbf{v} , we use the notation $S(P, \mathbf{v}, \prec)$ for the polytope S obtained from Lemma 3.2 applied after relabeling such that b_i is the i th inequality of the polytope P for $i = 1, \dots, d$, and call it the **silo** of P at (\mathbf{v}, \prec) . We then always have that $S(P, \mathbf{v}, \prec)$ can be computed in polynomial time given P, \mathbf{v} and \prec and satisfies a formula for the generating function corresponding to that of Lemma 3.2 after suitable relabeling. As further terminology, when constructing a silo $S(P, \mathbf{v}, \prec)$, we say that the vertex \mathbf{v} is **being siloed**. We also call the final vertex added in the construction process for $S(P, \mathbf{v}, \prec)$ (concretely, the vertex whose feasible basis corresponds to the monomial $\mathbf{y}^{[d]}$) the **peak** of the silo. Overall, the construction effectively replaces




FIGURE 2. Depicted is the silo construction in $d = 3$ dimensions in the normal fan of the polytope. Namely, the outer triangle corresponds to the normal cone of the vertex being cut off. We visualize this as a triangle by slicing the cone with a plane. Then the siloing subdivides that slice. The basis exchange graph corresponds to the dual graph of the triangulation. In this picture it is already visible that two cells may be of distance $d = 3$ away from each other as is the case for the highlighted cells on the right side of the picture.

the vertex \mathbf{v} being siloed with a tower that peaks at the peak of the silo, much like in the reduction of Frieze and Teng’s paper [20].

Siloing is almost enough to achieve our goal of reducing the shortest path problem on simple polytopes to the problem of computing the diameter of simple polytopes: By design of the construction (and as will be formally verified later), given some vertex \mathbf{u} of the original polytope P distinct from the siloed vertex \mathbf{v} , the distance from \mathbf{u} to the peak of $S(P, \mathbf{v}, \prec)$ is precisely the distance from \mathbf{u} to \mathbf{v} in P plus $d - 1$.

This property naturally leads to the following idea for our reduction: Namely, to apply the silo construction repeatedly. Concretely, suppose we are given as input a simple polytope P and a pair of vertices \mathbf{u}, \mathbf{v} between which we want to solve k -DISTANCE ON SIMPLE POLYTOPES. Then we silo \mathbf{u} , silo at the peak of that silo, and keep siloing at peaks repeatedly r times (for some suitably chosen, large enough, parameter r), such that the last peak \mathbf{u}' we created will have distance at least $r(d - 1)$ from any vertex of the original polytope. Then we repeat the same process at \mathbf{v} , yielding another “last” peak vertex \mathbf{v}' with the same property. One can then check that in the graph of the resulting polytope Q , we will have found a new pair \mathbf{u}', \mathbf{v}' of vertices satisfying $d_Q(\mathbf{u}', \mathbf{v}') = d_P(\mathbf{u}, \mathbf{v}) + 2r(d - 1)$. Furthermore, one can check that for any pair of vertices in the original polytope, even after this repeated siloing their distance will have changed by an additive constant of at most 6. Hence, (provided r was chosen large enough) the distance between any two original vertices of P in the final polytope Q will be much smaller than that of \mathbf{u}' and \mathbf{v}' . The hope would thus be to show that $d_P(\mathbf{u}', \mathbf{v}')$ equals the diameter of the new polytope Q , such that we could reduce the problem of computing/bounding the distance between \mathbf{u}, \mathbf{v} in the graph of P to the problem of computing the diameter of Q , providing the desired hardness result claimed by Theorem 1.5.

However, this idea narrowly fails to work, at least in the simple form that we now described. On the one hand, one can check that for any two vertices in the same so-called tower of silos, their distance is at most $r(d - 1) + 1$ and hence these vertices will be closer to each other than \mathbf{u}' and \mathbf{v}' , as desired. However, the problem is that it may happen that the distance between two vertices in different towers may be $2r(d - 1) + d_P(\mathbf{u}, \mathbf{v}) + 2$ in the worst case and thus (slightly) bigger than $d_Q(\mathbf{u}', \mathbf{v}')$. Hence, in such a case all we may conclude is that the diameter of Q lies somewhere between $d(\mathbf{u}, \mathbf{v}) + 2r(d - 1)$ and $d(\mathbf{u}, \mathbf{v}) + 2r(d - 1) + 2$. This, unfortunately, is not quite enough to determine the shortest path distance between \mathbf{u} and \mathbf{v} exactly. One would need an APX-hardness




FIGURE 3. Depicted is the graph G_d for $d = 5$. Vertices of the same height (i.e., second coordinate) are pairwise adjacent. Otherwise, there is an edge from a vertex to the first vertex above it and below it that is in the graph.

result here, which we however have no access to. In fact, we leave finding such an APX-hardness result as an open problem in Section 5.

Due to this subtle technical difficulty, we must be careful with how exactly we perform the described sequence of repeated siloings. To do so, we need to delve down into the combinatorics of the polytopes resulting from siloing and identify precisely which types of vertices can lead to the aforementioned increased distances. Based on this deeper understanding of the construction, we can construct the towers of silos such that the aforementioned bad situation never arises.

To start making these high-level ideas more concrete, we first need to better understand the adjacencies between the new vertices after siloing. To do so, for a natural number d , we define a graph G_d as having vertex set

$$V(G_d) := \{(a, b) \in [d]^2 \mid a \neq b\}$$

and where two distinct vertices (a, b) and (a', b') with $b \leq b'$ are adjacent if and only if one of the following holds:

- $b = b'$, or
- $b' = b + 1$, $a = a'$ and $b \neq a - 1$, or
- $b' = b + 2$, $a = a'$, and $b = a - 1$.

We call G_d the d -th silo graph. See Figure 3 for a visual illustration for $d = 5$. The following lemma describes the adjacencies between new vertices of the silo of a d -dimensional polytope precisely in terms of the graph G_d .

Lemma 3.5. *Let P be a d -dimensional simple polytope described by m irredundant inequalities. Let \mathbf{v} be a vertex of P and let \prec be a linear order on the elements $b_1 \prec b_2 \prec \dots \prec b_d$ of the feasible basis defining \mathbf{v} . Let H be the subgraph of the graph of $S(P, \mathbf{v}, \prec)$ induced by the vertices in $S(P, \mathbf{v}, \prec)$ that are not vertices of P and distinct from the peak. Then H and G_d are isomorphic. Moreover, assuming that we label the inequalities as $1, \dots, m$ such that b_i receives label i for $i = 1, \dots, d$, an isomorphism from G_d to H is given by mapping a vertex (a, b) of G_d to the vertex of $S(P, \mathbf{v}, \prec)$ associated with the monomial $\mathbf{x}^{[d] \setminus ([b-1] \cup \{a\})} \mathbf{y}^{[b]}$ if $a > b$ and $\mathbf{x}^{[d] \setminus [b-1]} \mathbf{y}^{[b] \setminus \{a\}}$ if $a < b$.*

Proof. Notice that by Lemma 3.4 the vertices of $S(P, \mathbf{v}, \prec)$ that are considered in H are exactly those whose monomials in the generating function appear in the sum

$$\sum_{k=0}^{d-1} \left(\sum_{i=1}^k \mathbf{x}^{[d] \setminus [k]} \mathbf{y}^{[k+1] \setminus \{i\}} + \sum_{j=k+2}^d \mathbf{x}^{[d] \setminus ([k] \cup \{j\})} \mathbf{y}^{[k+1]} \right).$$

These are exactly the monomials of the form $\mathbf{x}^{[d] \setminus ([b-1] \cup \{a\})} \mathbf{y}^{[b]}$ for some $1 \leq b < a \leq d$ and $\mathbf{x}^{[d] \setminus [b-1]} \mathbf{y}^{[b] \setminus \{a\}}$ for some $1 \leq a < b \leq d$, giving the desired bijection between vertices of H and G_d .

Now consider any distinct vertices $(a, b), (a', b')$ of G_d with $b \leq b'$. Let \mathbf{v}, \mathbf{v}' be their associated distinct vertices in H . To prove the statement of the lemma, we have to show that \mathbf{v} and \mathbf{v}' are adjacent in H if and only if $b = b'$; or $b' = b + 1, a = a'$ and $b \neq a - 1$; or $b' = b + 2, a = a'$ and $b = a - 1$. We start by showing sufficiency and split this into cases.

Case 1. Suppose first that $b = b'$. Then the monomials associated with \mathbf{v} and \mathbf{v}' are both obtained from $\mathbf{x}^{[d] \setminus [b-1]} \mathbf{y}^{[b]}$ by omitting exactly one variable. Hence, their bases have a symmetric difference of at most two and so \mathbf{v} and \mathbf{v}' are adjacent in the graph of $S(P, \mathbf{v}, \prec)$ and hence also in H , as desired.

Case 2. Suppose next that $b' = b + 1, a = a'$ and $b \neq a - 1$. Then then we can obtain the monomial of \mathbf{v}' from that of \mathbf{v} by replacing the variable x_b with the variable y_{b+1} (note that since $a = a' \notin \{b, b + 1\}$, the variable x_b indeed always occurs in the monomial representing \mathbf{v} , and the variable y_{b+1} indeed occurs in the monomial representing \mathbf{v}'). Hence, again the corresponding bases have a symmetric difference of size at most two and so \mathbf{v} and \mathbf{v}' are adjacent in H .

Case 3. Finally suppose that $b' = b + 2, a = a'$ and $b = a - 1$. Then \mathbf{v} and \mathbf{v}' are represented by $\mathbf{x}^{[d] \setminus ([b-1] \cup \{b+1\})} \mathbf{y}^{[b]}$ and $\mathbf{x}^{[d] \setminus [b+1]} \mathbf{y}^{[b] \cup \{b+2\}}$, respectively. Since the latter can be obtained from the first by exchanging the variable x_b for the variable y_{b+2} , indeed \mathbf{v} and \mathbf{v}' are adjacent also in this last case.

It remains to show necessity of the conditions. So suppose that \mathbf{v} and \mathbf{v}' are adjacent in H , i.e. their corresponding feasible bases of $S(P, \mathbf{v}, \prec)$ have symmetric difference of size two, and let us prove that at least one of the three conditions for adjacency in G_d is satisfied. Since the first condition holds if $b = b'$, in what follows we may and will assume $b' > b$.

Let us denote by $M, M' \subseteq \{x_1, \dots, x_d, y_1, \dots, y_d\}$ the sets of variables occurring in the monomials representing \mathbf{v} and \mathbf{v}' , respectively. We then have

$$M = \{y_1, \dots, y_b, x_b, \dots, x_d\} \setminus \{s\}, M' = \{y_1, \dots, y_{b'}, x_{b'}, \dots, x_d\} \setminus \{s'\},$$

where $s \in \{x_a, y_a\} \cap \{y_1, \dots, y_b, x_b, \dots, x_d\}$ and $s' \in \{x_{a'}, y_{a'}\} \cap \{y_1, \dots, y_{b'}, x_{b'}, \dots, x_d\}$. Since \mathbf{v} and \mathbf{v}' are adjacent, the sets M and M' must have symmetric difference exactly two. We therefore find

$$\begin{aligned} 2 &= |M \Delta M'| = |(\{y_1, \dots, y_b, x_b, \dots, x_d\} \Delta \{s\}) \Delta (\{y_1, \dots, y_{b'}, x_{b'}, \dots, x_d\} \Delta \{s'\})| \\ &= |\{x_b, \dots, x_{b'-1}, y_{b+1}, \dots, y_{b'}\} \Delta (\{s\} \Delta \{s'\})| \\ &= 2(b' - b) + |\{s\} \Delta \{s'\}| - 2|\{x_b, \dots, x_{b'-1}, y_{b+1}, \dots, y_{b'}\} \cap (\{s\} \Delta \{s'\})|. \end{aligned}$$

This immediately implies that either $s = s'$ and $b' = b + 1$ or $s \neq s'$ and

$$b' - b = |\{x_b, \dots, x_{b'-1}, y_{b+1}, \dots, y_{b'}\} \cap \{s, s'\}| \leq 2.$$

In the first case, since $s \in \{x_a, y_a\}$ and $s' \in \{x_{a'}, y_{a'}\}$, we must have $a = a'$, and hence $b \neq a - 1$, for otherwise $a' = a = b + 1 = b'$. However, this means that the second condition on (a, b) and (a', b') is satisfied, as desired.

So moving on suppose that the second case holds, i.e. $s \neq s', b' \in \{b + 1, b + 2\}$ and $b' - b = |\{x_b, \dots, x_{b'-1}, y_{b+1}, \dots, y_{b'}\} \cap \{s, s'\}|$.

Suppose first that $b' = b + 1$. Then $\{x_b, y_{b+1}\}$ shares exactly one element with $\{s, s'\}$. Since $s \in \{y_1, \dots, y_b, x_b, \dots, x_d\}$ and $s' \in \{y_1, \dots, y_{b+1}, x_{b+1}, \dots, x_d\}$, it follows that either $s = x_b$ or $s' = y_{b+1}$. Recall that $s \in \{x_a, y_a\}$ and $s' \in \{x_{a'}, y_{a'}\}$. Hence, in the first case, we obtain $a = b$, a contradiction, and in the second we obtain $a' = b + 1 = b'$, also a contradiction. It follows that this case is impossible and we may move on with the case $b' = b + 2$. We then have that $\{x_b, x_{b+1}, y_{b+1}, y_{b+2}\}$ shares exactly two common elements with $\{s, s'\}$, i.e., we have that s, s' are distinct elements of

$\{x_b, x_{b+1}, y_{b+1}, y_{b+2}\}$. Since $s \in \{y_1, \dots, y_b, x_b, \dots, x_d\}$ and $s' \in \{y_1, \dots, y_{b+2}, x_{b+2}, \dots, x_d\}$, we conclude that in fact $s \in \{x_b, x_{b+1}\}$ and $s' \in \{y_{b+1}, y_{b+2}\}$. Recalling further that $a \neq b$ and $a' \neq b' = b + 2$, we find that necessarily $s = x_{b+1}$ and $s' = y_{b+1}$. Hence, we have $a = a' = b + 1$ and so $b = a - 1$. Thus the third of the three conditions for adjacency in G_d is satisfied. This concludes the proof. \square

Note that for any simple polytope P , the peak of the silo $S(P, \mathbf{v}, \prec)$ is of distance at least d from any of the original vertices in P distinct from \mathbf{v} , since its feasible basis is disjoint from all feasible bases of original vertices in P , and hence the symmetric difference with these bases is of size $2d$.

Under the graph isomorphism in Lemma 3.5, the new vertices of the silo whose corresponding monomial has \mathbf{x} -support of size $d - 1$ are associated to the vertices $(1, 2)$ and $(i, 1)$ for $2 \leq i \leq d$ of G_d . These are also exactly the vertices that have some neighbor in the original polytope. At the same time, the new vertices of the silo whose corresponding monomial has \mathbf{y} -support of size $d - 1$ are associated to the vertices (i, d) for $1 \leq i \leq d - 1$ and $(d, d - 1)$ of G_d . These are also exactly the neighbors of the peak in the silo.

With our next lemma below, we will bound the distances between pairs of new vertices in a silo. To do so, by Lemma 3.5 it suffices to bound the distances between the associated vertices in the d -th silo graph G_d . Parts (a) and (b) of the following lemma show that any new vertex of a silo $S(P, \mathbf{v}, \prec)$ adjacent to an original vertex of P has a path of length at most $d - 2$ to a vertex of the silo adjacent to the peak. In fact, part (a) shows that there is a path of length at most $d - 2$ from the vertex with associated monomial $\mathbf{x}^{[d] \setminus \{1\}} y_2$ to all but one of the d neighbors of the peak in the silo. This flexibility of endpoints of paths starting from the vertex represented by $\mathbf{x}^{[d] \setminus \{1\}} y_2$ in conjunction with a rotation action will later allow us to effectively analyze a specific variant of the ‘‘repeated siloing’’ construction mentioned further above. Finally, part (c) of the lemma shows that all new vertices of $S(P, \mathbf{v}, \prec)$ can reach a vertex with a neighbor in the original polytope in at most $d - 2$ steps, and part (d) guarantees that any two vertices adjacent to original vertices of P are close to each other.

Lemma 3.6. *Let G_d be the d -th silo graph. Then*

- (a) *There is a path of length $d - 2$ from $(1, 2)$ to (i, d) for all $i \in [d - 1] \setminus \{2\}$ and to $(d, d - 1)$.*
- (b) *For each $2 \leq i \leq d - 1$, there is a path of length $d - 2$ from $(i, 1)$ to (i, d) and from $(d, 1)$ to $(d, d - 1)$.*
- (c) *Every vertex of G_d can reach some vertex in $\{(1, 2)\} \cup \{(i, 1) | 2 \leq i \leq d\}$ in at most $d - 2$ steps.*
- (d) *Any two vertices in the set $\{(1, 2)\} \cup \{(i, 1) | 2 \leq i \leq d\}$ have distance at most 3 from each other.*

Proof. For (a), if $i = 1$, simply increase the second coordinate until reaching $(1, d)$, and this takes at most $d - 2$ steps. If $3 \leq i \leq d - 1$, take one step to move from $(1, 2)$ to $(i, 2)$. Then increase the second coordinate until reaching (i, d) . This takes at most $d - 3$ steps, since each step increases the second index by 1 except for the step from $(i, i - 1)$ to $(i, i + 1)$, which increases it by 2. Thus, it takes $d - 2$ steps overall to reach (i, d) for $3 \leq i \leq d - 1$. Finally, for moving to $(d, d - 1)$, first take 1 step to move from $(1, 2)$ to $(d, 2)$ and then increase the second coordinate $d - 3$ times to reach $(d, d - 1)$ in $d - 2$ steps.

For (b), increase the second coordinate iteratively. This takes $d - 2$ steps for moving from $(i, 1)$ to (i, d) for $i \leq d - 1$, because all except one of the steps increase the second coordinate by 1, and as in the justification for part (a), the second coordinate increases by 2 from $(i, i - 1)$ to $(i, i + 1)$. For moving from $(d, 1)$ to $(d, d - 1)$, increasing the second coordinate straightforwardly takes $d - 2$ steps.

For (c), consider any vertex of (a, b) of G_d . Note that the vertices of G_d in $\{(x, b) | x \in [d] \setminus \{b\}\}$ induce a path on $d - 1$ vertices (and hence of length $d - 2$) as a subgraph of G_d , whose vertex with the lowest second coordinate is $(2, 1)$ for $b = 1$ and $(1, b)$ otherwise. Hence, by moving along this path we can always connect (a, b) to a vertex in $\{(1, 2)\} \cup \{(i, 1) | 2 \leq i \leq d\}$ in at most $d - 2$ steps, as desired.

For (d), it suffices to note that by definition of the graph G_d , the set of vertices $\{(i, 1) | 2 \leq i \leq d\}$ form a clique, and hence have pairwise distance one. Furthermore, since $(1, 2)$ is adjacent to $(3, 2)$, which is adjacent to $(3, 1)$ in G_d , the vertex $(1, 2)$ has distance at most 2 from some vertex in this clique. Hence, it has distance at most 3 from any vertex in this clique, proving the desired statement. \square

The following corollary records some further observations and consequences of Lemma 3.5 and Lemma 3.6 in a somewhat different language, which shall become useful later.

Corollary 3.7. *Let \mathbf{v} be a vertex of a simple polytope P and let \prec be a linear order on the elements $b_1 \prec \dots \prec b_d$ of the corresponding feasible basis. Let H be the subgraph of the graph of $S(P, \mathbf{v}, \prec)$ induced by the vertices not in P and distinct from the peak, and let $\phi : V(G_d) \rightarrow V(H)$ denote the graph isomorphism from Lemma 3.5. Let $\mathbf{u}_1 := \phi(1, 2)$, $\mathbf{u}_i := \phi(i, 1)$ for $2 \leq i \leq d$, $\mathbf{v}_i := \phi(i, d)$ for $1 \leq i \leq d - 1$ and $\mathbf{v}_d := \phi(d, d - 1)$. Furthermore, for $1 \leq i \leq d$ let \mathbf{s}_i denote the unique neighbor of \mathbf{v} on P whose feasible basis contains the elements $\{b_1, \dots, b_d\} \setminus \{b_i\}$. Then the following hold:*

- \mathbf{s}_i and \mathbf{u}_i are adjacent on $S(P, \mathbf{v}, \prec)$ for every $i \in [d]$.
- For each $i \in [d]$, there exists a path of length at most $d - 2$ from \mathbf{u}_i to \mathbf{v}_i in H , as well as a path of length at most $d - 2$ from \mathbf{u}_1 to \mathbf{v}_i for each $i \neq 2$.
- $\mathbf{v}_1, \dots, \mathbf{v}_d$ are the neighbors of the peak of $S(P, \mathbf{v}, \prec)$.
- Every vertex in H has distance at most $d - 2$ from some vertex in $\{\mathbf{u}_1, \dots, \mathbf{u}_d\}$.
- Any two of $\mathbf{u}_1, \dots, \mathbf{u}_d$ have distance at most 3 from each other in H .

Proof. The first and the third items can easily be checked by inspecting the monomials representing the respective vertices and observing that they only differ by exchanging a single variable.

The second item follows immediately by combining Lemmas 3.5 and 3.6. Finally, the last two items follow directly from parts (c) and (d) of Lemma 3.6, since $\{\mathbf{u}_1, \dots, \mathbf{u}_d\}$ is the image of $\{(1, 2)\} \cup \{(i, 1) | 2 \leq i \leq d\}$ under the graph isomorphism ϕ . \square

Guided by these structural observations about short paths between the new vertices in a silo, we will now introduce the previously announced construction which involves repeated siloing in a cyclic manner, which we call **cyclic siloing**. In the following, we will repeatedly use the following notation: For an integer $z \in \mathbb{Z}$, we denote by \bar{z} the unique member of $[d]$ which is congruent to z modulo d .

To explain this construction, suppose we are given as input a simple d -dimensional polytope P described by m inequalities, and a vertex \mathbf{v} . Suppose further we are given as input some positive integer r . We will then construct, in polynomial time in the encoding length of P and in r , a sequence of simple d -dimensional polytopes C_0, C_1, \dots, C_{rd} , and for each $0 \leq j \leq rd$ a special vertex \mathbf{v}_j on C_j , an enumeration $\mathbf{v}_{1,j}, \dots, \mathbf{v}_{d,j}$ of the d neighbors of \mathbf{v}_j on C_j , and for every $j \in [rd]$ another sequence $\mathbf{u}_{1,j}, \dots, \mathbf{u}_{d,j}$ of special vertices on C_j , as follows.

- To initialize, for $j = 0$ we set $C_0 := P$, $\mathbf{v}_0 := \mathbf{v}$. Finally, we fix some arbitrary enumeration $\mathbf{v}_{1,0}, \dots, \mathbf{v}_{d,0}$ of the d neighbors of \mathbf{v} on P .
- Next let $j \geq 1$, and suppose we already computed C_{j-1} , \mathbf{v}_{j-1} and enumerated the d neighbors of \mathbf{v}_{j-1} on C_{j-1} as $\mathbf{v}_{1,j-1}, \dots, \mathbf{v}_{d,j-1}$. Now compute the unique labeling $b_{1,j-1}, \dots, b_{d,j-1}$ of the elements of the feasible basis of C_{j-1} corresponding to \mathbf{v}_{j-1} such that the tight inequalities shared by $\mathbf{v}_{i,j-1}$ and \mathbf{v}_{j-1} are exactly $\{b_{1,j-1}, \dots, b_{d,j-1}\} \setminus \{b_{i,j-1}\}$. Let \prec_{j-1} be

the linear order on the elements of the feasible basis of \mathbf{v}_{j-1} in C_{j-1} defined by

$$b_{\bar{j},j-1} \prec_{j-1} b_{\bar{j+1},j-1} \prec_{j-1} \cdots \prec_{j-1} b_{\bar{j+d-1},j-1}.$$

We then set $C_j := S(C_{j-1}, \mathbf{v}_{j-1}, \prec_{j-1})$. Next, we set \mathbf{v}_j to be the peak of the silo C_j . Finally, we consider the isomorphism ϕ_j from the d -th silo graph G_d to the subgraph H_j of the graph of C_j induced by all vertices not in C_{j-1} and distinct from the peak \mathbf{v}_j , as described in Lemma 3.5. We then set $\mathbf{u}_{i,j} := \phi_j(\overline{i-j+1}, 1)$ for all $i \in [d] \setminus \{\bar{j}\}$ and $\mathbf{u}_{\bar{j},j} := \phi_j(1, 2)$. Furthermore, we set $\mathbf{v}_{i,j} := \phi_j(\overline{i-j+1}, d)$ for all $i \in [d] \setminus \{\bar{j}-1\}$ and $\mathbf{v}_{\bar{j}-1,j} := \phi_j(d, d-1)$. Note that by Corollary 3.7, the so-defined vertices $\mathbf{v}_{1,j}, \dots, \mathbf{v}_{d,j}$ indeed form the neighbors of the peak \mathbf{v}_j of C_j .

Finally, the last polytope in our construction, C_{rd} , is a simple d -dimensional polytope, which we call the **r -cyclic siloing** of P and denote by $C^r(P, \mathbf{v})$. Note that by following the steps described above, and by Lemma 3.1, given as input P , \mathbf{v} , \prec and r we can compute in polynomial time in the encoding length of P and in r an inequality description of $C^r(P, \mathbf{v})$ whose encoding length is polynomial in the encoding length of P and in r .

We call the subgraph of the graph of the r -cyclic siloing of a polytope P induced by all vertices not in the original polytope P the **cyclic silo**. We also refer to the set of vertices $\{\mathbf{u}_{1,1}, \dots, \mathbf{u}_{d,1}\}$ of the cyclic silo as the **ground layer** of the cyclic silo.

Before proceeding to analyze the distances between vertices on the r -cyclic siloing of a polytope in more detail, we record the following observation which follows directly from our previous lemmas and the construction described above.

Remark 3.8.

- For every $(i, j) \in [d] \times [rd]$ we have that $\mathbf{u}_{i,j}$ and $\mathbf{v}_{i,j}$ are vertices in the cyclic silo of $C^r(P, \mathbf{v})$.
- For every $(i, j) \in [d] \times [rd]$, we have that $\mathbf{v}_{i,j-1}$ and $\mathbf{u}_{i,j}$ are adjacent vertices of the polytope $C^r(P, \mathbf{v})$.
- For every $(i, j) \in [d] \times [rd]$, we have that $\mathbf{u}_{i,j}$ and $\mathbf{v}_{i,j}$ have distance at most $d-2$ in the cyclic silo. Additionally, if $i \neq \bar{j}+1$, then $\mathbf{u}_{\bar{j},j}$ and $\mathbf{v}_{i,j}$ have distance at most $d-2$ in the cyclic silo.
- Any two vertices in the ground layer have distance at most 3 in the cyclic silo.

Proof.

- Consider any $(i, j) \in [d] \times [rd]$. Then by definition of the process for building the r -cyclic siloing of P , we have that $\mathbf{u}_{i,j}$ and $\mathbf{v}_{i,j}$ are vertices of the silo C_j constructed during the process distinct from the peak \mathbf{v}_j of C_j . In the process, the final polytope $C^r(P, \mathbf{v}) = C_{rd}$ arises from C_j by repeatedly siloing vertices, starting with \mathbf{v}_j and then continuing always by siloing *new* vertices created in the previous siloing step. In particular, throughout the rest of the process, no vertex of C_j except \mathbf{v}_j gets siloed, and hence all vertices of C_j distinct from \mathbf{v}_j remain vertices of $C^r(P, \mathbf{v})$. This includes $\mathbf{u}_{i,j}$ and $\mathbf{v}_{i,j}$, confirming the statement claimed in the first item.
- By the same argument we can observe that for every $j \in \{0, \dots, rd\}$ all adjacencies between vertices of C_j distinct from \mathbf{v}_j remain intact in the final polytope $C^r(P, \mathbf{v})$. Hence, for the second item it suffices to verify that $\mathbf{v}_{i,j-1}$ and $\mathbf{u}_{i,j}$ are adjacent on the polytope $C_j = S(C_{j-1}, \mathbf{v}_{j-1}, \prec_{j-1})$. By our description of the process, we have that $\mathbf{v}_{i,j-1}$ is the unique neighbor of \mathbf{v}_{j-1} on C_{j-1} whose feasible basis includes $\{b_{1,j-1}, \dots, b_{d,j-1}\} \setminus \{b_{i,j-1}\}$. By definition of \prec_{j-1} , this means that $\mathbf{v}_{i,j-1}$ corresponds to the vertex $\mathbf{s}_{\overline{i-j+1}}$ in the notation of Corollary 3.7 (applied to C_{j-1} and \prec_{j-1} instead of P and \prec). Hence, it follows from the first item of Corollary 3.7 that $\mathbf{v}_{i,j-1}$ is adjacent to $\phi_j(\overline{i-j+1}, 1)$ on C_j if $i \neq \bar{j}$ and to

$\phi_j(1, 2)$ on C_j if $i = \bar{j}$. By definition of $\mathbf{u}_{i,j}$ in the process, it follows that indeed $\mathbf{v}_{i,j-1}$ is adjacent to $\mathbf{u}_{i,j}$ on C_j , as desired.

- Let $(i, j) \in [d] \times [rd]$. Recall that $C_j = S(C_{j-1}, \mathbf{v}_{j-1}, \prec_{j-1})$, where \prec_j is the linear order on the elements of the feasible basis representing \mathbf{v}_{j-1} in C_{j-1} , defined as

$$b_{\bar{j},j-1} \prec_{j-1} b_{\overline{j+1},j-1} \prec_{j-1} \cdots \prec_{j-1} b_{\overline{j+d-1},j-1}.$$

Recall that ϕ_j denotes the isomorphism from G_d to H_j as given by Lemma 3.5, and that $\mathbf{u}_{i,j} = \phi_j(\overline{i-j+1}, 1)$ for $i \in [d] \setminus \{\bar{j}\}$ as well as $\mathbf{u}_{\bar{j},j} = \phi_j(1, 2)$ by definition in the process. Similarly, we have $\mathbf{v}_{i,j} = \phi_j(\overline{i-j+1}, d)$ for all $i \in [d] \setminus \{\bar{j-1}\}$ and $\mathbf{v}_{\overline{j-1},j} = \phi_j(d, d-1)$.

Hence, for each $i \in [d]$, the vertices $\mathbf{u}_{i,j}$ and $\mathbf{v}_{i,j}$ correspond exactly to the vertices $\mathbf{u}_{\overline{i-j+1}}$ and $\mathbf{v}_{\overline{i-j+1}}$ in the notation of Corollary 3.7, applied to $C_j = S(C_{j-1}, \mathbf{v}_{j-1}, \prec_{j-1})$ in place of $S(P, \mathbf{v}, \prec)$. It thus follows by the second item of Corollary 3.7 that for each $i \in [d]$ there exists a path of length at most $d-2$ on C_j from $\mathbf{u}_{i,j}$ to $\mathbf{v}_{i,j}$, and from $\mathbf{u}_{\bar{j},j}$ to $\mathbf{v}_{i,j}$, provided that $\overline{i-j+1} \neq 2$, i.e., $i \neq \bar{j+1}$. Moreover, all vertices of these paths are distinct from the peak of C_j . As we argued in the first and second item of this proof, this means that these paths also exist in the cyclic silo of $C^r(P, \mathbf{v})$. This establishes the statement claimed in the third item of the remark and concludes its proof.

- For the fourth statement claimed in the remark, note that in the language of Corollary 3.7 $\mathbf{u}_{1,1}, \dots, \mathbf{u}_{1,d}$ correspond exactly to the vertices $\mathbf{u}_1, \dots, \mathbf{u}_d$ in the first silo $C_1 = S(C_0, \mathbf{v}_0, \prec_0)$ constructed in the process. Hence, by the fifth item of that corollary, we have that any two of $\mathbf{u}_{1,1}, \dots, \mathbf{u}_{1,d}$ can be connected by a path of length at most 3 in the graph of C_1 which does not use the peak \mathbf{v}_1 . By what we observed before, any such path also forms a path in the cyclic silo of $C^r(P, \mathbf{v})$, and so we obtain the claimed statement.

□

Our next lemma is a key technical step towards the proof of our Theorem 1.5. It gives an upper bound on the distances between certain pairs of vertices in the cyclic silo.

Lemma 3.9. *Let P be a simple d -dimensional polytope, \mathbf{v} a vertex of P and let $r \geq 3$ be an integer. Then for every pair of vertices \mathbf{s}, \mathbf{t} of the cyclic silo, at least one of which lies in the ground layer, we have that their distance in the cyclic silo is at most $rd(d-1)$. In particular, the diameter of the cyclic silo is at most $2rd(d-1)$. Furthermore, the distance on $C^r(P, \mathbf{v})$ from the peak \mathbf{v}_{rd} to any vertex of P distinct from \mathbf{v} is at least $rd(d-1) + 1$.*

Proof. Throughout this proof, we will use the same notation for vertices and polytopes as defined in the description of the process for constructing the r -cyclic siloing $C^r(P, \mathbf{v})$.

In the following, let us define an auxiliary graph Γ on vertex-set $\{\mathbf{u}_{i,j} | (i, j) \in [d] \times [rd]\}$, where we make $\mathbf{u}_{i,j}$ and $\mathbf{u}_{i',j'}$ adjacent if and only if there is a path of length at most $d-1$ between them in the cyclic silo. It follows from the first three items of Remark 3.8 that there is an edge in Γ from $\mathbf{u}_{i,j}$ to $\mathbf{u}_{i,j+1}$ for all $(i, j) \in [d] \times [rd-1]$. Additionally, Remark 3.8 implies that $\mathbf{u}_{\bar{j},j}$ is adjacent to $\mathbf{u}_{i,j+1}$ in Γ provided $i \neq \overline{j+1}$. It will be useful to first prove the following claim.

Claim 1. Let $i, i' \in [d]$ and $j \in [rd]$ be such that $j \geq d+3$. Then $\mathbf{u}_{i,1}$ and $\mathbf{u}_{i',j}$ have distance at most $(d-1)(j-1)$ in the cyclic silo.

Proof. By definition of the graph Γ , it clearly suffices to show that the distance between $\mathbf{u}_{i,1}$ and $\mathbf{u}_{i',j}$ in Γ is at most $j-1$. To do so, we will construct a path from $\mathbf{u}_{i,1}$ to $\mathbf{u}_{i',j}$ in Γ where whenever we move along an edge of the path, the second index of the current vertex increases by exactly 1. Clearly, such a path will always have the desired length of $j-1$.

Suppose first that $i' \neq \overline{i+1}$. Starting at $\mathbf{u}_{i,1}$, we can then move in Γ to $\mathbf{u}_{i,i}$ in $i-1$ steps by successively increasing the second index. Next, we can move to $\mathbf{u}_{i',i+1}$ in one step. Finally, we again increase the second index successively until arriving at the desired vertex $\mathbf{u}_{i',j}$ along a path of the desired in type in Γ . This is possible since $i+1 \leq d+1 \leq j$.

Next, suppose $i' = \overline{i+1}$. As before, starting at $\mathbf{u}_{i,1}$ we first successively increase the second index to reach $\mathbf{u}_{i,i}$. We then move from $\mathbf{u}_{i,i}$ to $\mathbf{u}_{\overline{i+2},i+1}$ in one step and then to $\mathbf{u}_{\overline{i+2},i+2}$ in another step. Since $\overline{i+1} \neq \overline{i+3}$ (here we use $d \geq 3$), we can next move to $\mathbf{u}_{\overline{i+1},i+3}$. Finally we successively increase the second coordinate until reaching $\mathbf{u}_{\overline{i+1},j} = \mathbf{u}_{i',j}$ via a path in Γ of the desired form. This is possible since $i+3 \leq d+3 \leq j$.

Having found the desired path in Γ of length $j-1$ in both cases, we may conclude the proof. \blacksquare

Next we want to show the desired upper bound on the distance between pairs \mathbf{s}, \mathbf{t} of vertices in the cyclic silo at least one of which belongs to the ground layer.

We first prepare some useful setup. Recall that for $j \in [rd]$ we denote by H_j the subgraph of the graph of C_j induced by all vertices distinct from the peak \mathbf{v}_j of C_j that are not vertices of C_{j-1} . Pause to note that each H_j is in fact a subgraph of the cyclic silo, and that the vertices of the cyclic silo partition into the disjoint sets of vertices $V_1 := V(H_1), \dots, V_{rd} := V(H_{rd})$ and $V_{rd+1} := \{\mathbf{v}_{rd}\}$. Furthermore, note that each graph H_j is isomorphic to the d -th silo graph by Lemma 3.5. Additionally, it follows straightforwardly from the fourth item of Corollary 3.7 and from the definition of the vertices $\mathbf{u}_{i,j}$ in the process of constructing the r -cyclic siloing that for each $j \in [rd]$, every vertex \mathbf{w} in H_j has distance at most $d-2$ from *some* vertex in $\{\mathbf{u}_{1,j}, \dots, \mathbf{u}_{d,j}\}$.

So let now a pair \mathbf{s}, \mathbf{t} of vertices of the cyclic silo be given to us such that \mathbf{s} belongs to the ground layer. Let $i \in [d]$ be such that $\mathbf{s} = \mathbf{u}_{i,1}$ and let $j \in [rd+1]$ be the unique index such that $\mathbf{t} \in V_j$.

Suppose first that $j \leq rd$. Then $\mathbf{t} \in V(H_j)$. By our above remark, there then exists some $i' \in [d]$ such that \mathbf{t} has distance at most $d-2$ from $\mathbf{u}_{i',j}$ in H_j , and hence in the cyclic silo.

If $j \geq d+3$, then by Claim 1 we have that $\mathbf{s} = \mathbf{u}_{i,1}$ and $\mathbf{u}_{i',j}$ have distance at most $(d-1)(j-1)$ in the cyclic silo. By the triangle inequality, we then find that \mathbf{s} and \mathbf{t} have distance at most $(d-1)(j-1) + d-2 = j(d-1) - 1 < rd(d-1)$ in the cyclic silo, as desired.

On the other hand, if $j \leq d+2$, then by successively decreasing the second index we can see that $\mathbf{u}_{i',j}$ has a path of length at most $j-1$ in Γ to $\mathbf{u}_{i',1}$. In particular, the distance between $\mathbf{u}_{i',j}$ and $\mathbf{u}_{i',1}$ in the cyclic silo is at most $(d-1)(j-1)$. Furthermore, by the last item of Remark 3.8, we have that $\mathbf{s} = \mathbf{u}_{i,1}$ and $\mathbf{u}_{i',1}$ have distance at most 3 in the cyclic silo. Altogether, it follows from the triangle inequality that \mathbf{s} and \mathbf{t} have distance at most $3+(d-1)(j-1)+(d-1) = j(d-1)+3 \leq (d+2)(d-1)+3 \leq rd(d-1)$, where we used our assumption $r \geq 3$ as well as our general assumption that $d \geq 3$, in the last step.

Hence, it only remains to consider the case that $j = rd+1$, i.e., $\mathbf{t} = \mathbf{v}_{rd}$. Then, since by definition \mathbf{v}_{rd} is the peak of the silo $C_{rd} = S(C_{rd-1}, \mathbf{v}_{rd-1}, \prec_{rd-1})$, it follows by the third item of Corollary 3.7 that each of the vertices $\mathbf{v}_{rd,1}, \dots, \mathbf{v}_{rd,d}$ are the neighbors of $\mathbf{t} = \mathbf{v}_{rd}$ on C_{rd} and hence in the cyclic silo. Further, by the third item of Remark 3.8 we have that $\mathbf{u}_{i,rd}$ has distance at most $d-2$ from $\mathbf{v}_{i,rd}$ in the cyclic silo, and hence distance at most $d-1$ from $\mathbf{v}_{rd} = \mathbf{t}$ in the cyclic silo. By Claim 1 or simply by walking in the graph Γ , we can also see that $\mathbf{s} = \mathbf{u}_{i,1}$ has distance at most $(d-1)(rd-1)$ from $\mathbf{u}_{i,1} = \mathbf{s}$ in the cyclic silo. Hence, by the triangle inequality, we find that \mathbf{s} and \mathbf{t} have distance at most $(d-1)(rd-1) + (d-1) = rd(d-1)$ in the cyclic silo. This concludes the proof of the first part of the lemma.

It remains to prove that the distance on $C_{rd} = C^r(P, \mathbf{v})$ from the peak \mathbf{v}_{rd} to any vertex of P distinct from \mathbf{v} is at least $rd(d-1) + 1$. We will do this by proving the following, more general, statement by induction.

Claim 2. For every $j \in \{0, 1, \dots, rd\}$ and every $i \in [d]$, the distance on C_j between \mathbf{v}_j and any vertex of P distinct from \mathbf{v} is at least $j(d-1) + 1$.

Proof. The induction basis $j = 0$ is obvious, since $C_0 = P$ and $\mathbf{v}_0 = \mathbf{v}$. So suppose that for some $j \in [rd]$ we already proved that the distance on C_{j-1} from \mathbf{v}_{j-1} to any vertex of P distinct from \mathbf{v} is at least $(j-1)(d-1) + 1$, and let us prove the analogous claim with $j-1$ replaced by j . Let R denote a shortest path from \mathbf{v}_j to some vertex \mathbf{w} of P distinct from \mathbf{v} on C_j . Recall that $C_j = S(C_{j-1}, \mathbf{v}_{j-1}, \prec_{j-1})$, and pause to note that by definition of the siloing operation, the members of the set $N := \{\mathbf{v}_{1,j-1}, \dots, \mathbf{v}_{d,j-1}\}$ are the only vertices of C_{j-1} who have a neighbor on C_j not in C_{j-1} . Hence, N forms a separator between \mathbf{w} and \mathbf{v}_j in the graph of C_j . In particular, R must contain at least one vertex in N . Let \mathbf{n} denote the vertex in N which we meet first when traversing R from \mathbf{w} to \mathbf{v}_j . Note that all vertices of the segment of R from \mathbf{w} to \mathbf{n} must be vertices of C_{j-1} , for otherwise there would be a vertex in N along the path that is closer to \mathbf{w} than \mathbf{n} . Hence, we find that the length of the segment of R from \mathbf{w} to \mathbf{n} is at least $d_{C_{j-1}}(\mathbf{w}, \mathbf{n}) \geq d_{C_{j-1}}(\mathbf{w}, \mathbf{v}_{j-1}) - d_{C_{j-1}}(\mathbf{n}, \mathbf{v}_{j-1}) = d_{C_{j-1}}(\mathbf{w}, \mathbf{v}_{j-1}) - 1 \geq (j-1)(d-1)$, where we used the fact that $\mathbf{n} \in N$ is a neighbor of \mathbf{v}_{j-1} on C_{j-1} in the second to last step, and the inductive assumption in the last step. Now, let us consider the segment of R from \mathbf{n} to \mathbf{v}_j . Since \mathbf{n} is a vertex of C_{j-1} and \mathbf{v}_j the peak of a siloing performed on C_{j-1} , we have that the feasible basis representing \mathbf{v}_j on C_j is disjoint from that of \mathbf{n} (compare also our analogous remark directly after the proof of Lemma 3.5). Hence, the distance between \mathbf{n} and \mathbf{v}_j on C_j and thus the length of the segment of R from \mathbf{n} to \mathbf{v}_j must be at least d . It now follows that $d_{C_j}(\mathbf{w}, \mathbf{v}_j) = |R| \geq (j-1)(d-1) + d = j(d-1) + 1$, as desired. This established the inductive claim and concludes the proof. \blacksquare

□

With Lemma 3.9, the key technical point of our argument has now been established. The construction we use for our hardness reduction from k -DISTANCE ON SIMPLE POLYTOPES to DIAMETER OF SIMPLE POLYTOPES will look as follows. We are given as input a d -dimensional simple polytope P and the two vertices \mathbf{u}, \mathbf{v} of a simple input polytope P for which we want to decide if their distance on P is at most some input number $k \in \mathbb{N}$. We will then apply, for some large enough choice of $r \in \mathbb{N}$, the r -cyclic siloing operation both at \mathbf{u} and then at \mathbf{v} . Finally, our goal in the following will be to show that the diameter of the arising simple polytope Q equals $d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, and hence we will be able to determine the distance between \mathbf{u} and \mathbf{v} on P by computing the diameter of Q , establishing the desired reduction.

As a first step we prove the following lemma, bounding the effect that r -cyclic siloing can have on the distances between the original vertices of the polytope we apply it to.

Lemma 3.10. *Let P be a simple polytope and \mathbf{v} a vertex of P . Let \mathbf{s}, \mathbf{t} be vertices of $C^r(P, \mathbf{v})$ not contained in the cyclic silo. In particular, \mathbf{s} and \mathbf{t} are also vertices of P . Then*

$$d_{C^r(P, \mathbf{v})}(\mathbf{s}, \mathbf{t}) \leq d_P(\mathbf{s}, \mathbf{t}) + 3.$$

Proof. Suppose first that some shortest path from \mathbf{s} to \mathbf{t} on P does not pass through \mathbf{v} . Then it is still a path in $C^r(P, \mathbf{v})$, so $d_{C^r(P, \mathbf{v})}(\mathbf{s}, \mathbf{t}) \leq d_P(\mathbf{s}, \mathbf{t})$.

Next suppose that every shortest path between \mathbf{s} and \mathbf{t} on P does go through \mathbf{v} . Pick one such shortest path R , and note that it uses precisely two edges incident to \mathbf{v} . Recall that in the process for constructing the r -cyclic siloing $C^r(P, \mathbf{v})$ of P at \mathbf{v} , the vertices $\mathbf{v}_{1,0}, \dots, \mathbf{v}_{d,0}$ denote the neighbors of \mathbf{v} on P . Hence, R must use vertices $\mathbf{v}_{i,0}, \mathbf{v}, \mathbf{v}_{i',0}$ in this order, for some distinct $i, j \in [d]$. By the second item of Remark 3.8, we have that on the polytope $C^r(P, \mathbf{v})$, the vertices $\mathbf{v}_{i,0}, \mathbf{u}_{i,1}$ and $\mathbf{v}_{i',0}, \mathbf{u}_{i',1}$ are adjacent. Furthermore, by the fourth item of Remark 3.8, we have that there exists a path R' on $C^r(P, \mathbf{v})$ between $\mathbf{u}_{i,1}$ and $\mathbf{u}_{i',1}$ of length at most 3, all whose vertices are in the cyclic

silo of $C^r(P, \mathbf{v})$ (and hence, in particular R' shares no vertices with R). We can now see that by replacing the subpath $\mathbf{v}_{i,0}, \mathbf{v}, \mathbf{v}'_{i,0}$ of R of length two by the path of length at most 5 in $C^r(P, \mathbf{v})$ obtained as the union of the edges $\mathbf{v}_{i,0}\mathbf{u}_{i,1}$, $\mathbf{v}'_{i,0}\mathbf{u}'_{i,1}$ and the path R' defined above, we obtain a path between \mathbf{s} and \mathbf{t} on $C^r(P, \mathbf{v})$. Hence, we have $d_{C^r(P, \mathbf{v})}(\mathbf{s}, \mathbf{t}) \leq d_P(\mathbf{s}, \mathbf{t}) - 2 + 5 = d_P(\mathbf{s}, \mathbf{t}) + 3$, as desired. This concludes the proof. \square

With this auxiliary statement at hand, we are now finally in the position to prove the desired formula for the diameter of the polytope Q obtained after r -cyclic siloing at two given vertices \mathbf{u}, \mathbf{v} of a polytope P , as we mentioned above.

Theorem 3.11. *Let P be a simple polytope with vertices $\mathbf{u} \neq \mathbf{v}$. Let Q be the result of applying an r -cyclic siloing at both u and v , where $r \geq \max(\text{diam}(P), 6)$. Formally, $Q := C^r(C^r(P, \mathbf{u}), \mathbf{v})$. Then*

$$\text{diam}(Q) = d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1).$$

Proof. In the following, let us denote by $W_{\mathbf{u}}, W_{\mathbf{v}}$ the sets of vertices in the cyclic silos corresponding to \mathbf{u} and \mathbf{v} , respectively.

Let us first show that $\text{diam}(Q) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, i.e. that every given pair \mathbf{s}, \mathbf{t} of vertices of Q satisfies $d_Q(\mathbf{s}, \mathbf{t}) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$.

Suppose first that none of \mathbf{s}, \mathbf{t} lie in $W_{\mathbf{u}} \cup W_{\mathbf{v}}$. Then by Lemma 3.10, we have

$$d_Q(\mathbf{s}, \mathbf{t}) \leq d_P(\mathbf{s}, \mathbf{t}) + 6 \leq \text{diam}(P) + 6 \leq 2r \leq 2rd(d-1) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1),$$

as desired.

Next, suppose that exactly one of \mathbf{s}, \mathbf{t} lie in $W_{\mathbf{u}} \cup W_{\mathbf{v}}$. W.l.o.g. (possibly after relabeling), we may then assume $\mathbf{s} \notin W_{\mathbf{u}} \cup W_{\mathbf{v}}$, $\mathbf{t} \in W_{\mathbf{v}}$.

By Lemma 3.10, we then have that $d_{C^r(P, \mathbf{u})}(\mathbf{s}, \mathbf{v}) \leq d_P(\mathbf{s}, \mathbf{v}) + 3 \leq \text{diam}(P) + 3$. In particular, this implies that there exists a path R (not traversing \mathbf{v}) of length at most $\text{diam}(P) + 2$ from \mathbf{s} to a neighbor \mathbf{v}' of \mathbf{v} on $C^r(P, \mathbf{u})$. By the second item of Remark 3.8, applied with $j = 1$, it now follows that \mathbf{v}' has a neighbor in the ground layer of $W_{\mathbf{v}}$. Since also $\mathbf{t} \in W_{\mathbf{v}}$, Lemma 3.9 implies that there exists a path of length at most $rd(d-1)$ between this neighbor of \mathbf{v}' and \mathbf{t} on Q , all whose vertices lie in $W_{\mathbf{v}}$. It now follows that $d_Q(\mathbf{s}, \mathbf{t}) \leq |R| + 1 + rd(d-1) \leq \text{diam}(P) + 3 + rd(d-1) \leq 2r + rd(d-1) \leq 2rd(d-1) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, as desired.

Finally, suppose that both of \mathbf{s}, \mathbf{t} lie in $W_{\mathbf{u}} \cup W_{\mathbf{v}}$. If both lie in the same cyclic silo, then by Lemma 3.9 we have $d_Q(\mathbf{u}, \mathbf{v}) \leq 2rd(d-1) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, as desired. Hence, moving on we may assume that they lie in different cyclic silos, so we may w.l.o.g. assume $\mathbf{s} \in W_{\mathbf{u}}$, $\mathbf{t} \in W_{\mathbf{v}}$. Now, let R be a path of length $d_P(\mathbf{u}, \mathbf{v}) - 2$ on P connecting a neighbor \mathbf{u}' of \mathbf{u} on P to a neighbor \mathbf{v}' of \mathbf{v} on P , and note that all vertices of R are distinct from \mathbf{u} and \mathbf{v} . In particular, R is also a path on Q . Again by the second item of Remark 3.8, applied with $j = 1$, we have that \mathbf{u}' is adjacent on Q to a vertex \mathbf{u}'' in the ground layer of $W_{\mathbf{u}}$, and that \mathbf{v}' is adjacent on Q to a vertex \mathbf{v}'' in the ground layer of $W_{\mathbf{v}}$. Since $\mathbf{s} \in W_{\mathbf{u}}$ and $\mathbf{t} \in W_{\mathbf{v}}$, Lemma 3.9 implies that \mathbf{u}'' and \mathbf{s} as well as \mathbf{v}'' and \mathbf{t} have distance at most $rd(d-1)$ on Q . By the triangle inequality, we may thus conclude that $d_Q(\mathbf{s}, \mathbf{t}) \leq rd(d-1) + 1 + |R| + 1 + rd(d-1) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, as desired. Having covered all the cases, this concludes the proof of the upper bound $\text{diam}(Q) \leq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$.

For the lower bound, let us denote by $\mathbf{p}_{\mathbf{u}}, \mathbf{p}_{\mathbf{v}}$ the final peaks of the cyclic silos $W_{\mathbf{u}}, W_{\mathbf{v}}$, respectively. We will show that $d_Q(\mathbf{p}_{\mathbf{u}}, \mathbf{p}_{\mathbf{v}}) \geq d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$, which will clearly establish the desired lower bound on $\text{diam}(Q)$. By construction of the r -cyclic siloing, we have that the set $N_{\mathbf{u}}$ of neighbors of \mathbf{u} on P separates $W_{\mathbf{u}}$ from the rest of the graph of Q , and similarly the set $N_{\mathbf{v}}$ of neighbors of \mathbf{v} on P separates $W_{\mathbf{v}}$ from the rest of the graph of Q . Now consider a shortest path from $\mathbf{p}_{\mathbf{u}}$ to $\mathbf{p}_{\mathbf{v}}$ on Q . Let \mathbf{u}' denote the last vertex of the path in $N_{\mathbf{u}}$ when traversing it from $\mathbf{p}_{\mathbf{u}}$ to

\mathbf{p}_v . Furthermore, let \mathbf{v}' denote the first vertex of N_v we meet when traversing the segment of the path from \mathbf{u}' to \mathbf{p}_v . Note that all vertices in the segment of the path between \mathbf{u}' and \mathbf{v}' must also be vertices of P , and hence this segment forms a path in P and has length at least $d_P(\mathbf{u}', \mathbf{v}') \geq d_P(\mathbf{u}, \mathbf{v}) - d_P(\mathbf{u}', \mathbf{u}) - d_P(\mathbf{v}', \mathbf{v}) = d_P(\mathbf{u}, \mathbf{v}) - 2$.

Furthermore, it follows directly from Lemma 3.9 that the segment of the path from \mathbf{p}_u to $\mathbf{u}' \in N_u$, as well as the segment of the path from \mathbf{p}_v to $\mathbf{v}' \in N_v$, both must have length at least $rd(d-1) + 1$.

Altogether, this implies that the total length $d_Q(\mathbf{p}_u, \mathbf{p}_v)$ of the shortest path we considered is at least $(rd(d-1) + 1) + (d_P(\mathbf{u}, \mathbf{v}) - 2) + (rd(d-1) + 1) = d_P(\mathbf{u}, \mathbf{v}) + 2rd(d-1)$. This is what we wanted to prove, and hence we may conclude the proof of the theorem. \square

The following result summarizes our observations made and the auxiliary results we proved so far.

Theorem 3.12. *Given as input a simple polytope P in inequality description, a pair of vertices \mathbf{u}, \mathbf{v} of P and a number $r \in \mathbb{N}$ such that $r \geq \max(\text{diam}(P), 6)$, one can compute another simple polytope Q (also in inequality description) as well as a constant K for which the diameter of Q equals $d_P(\mathbf{u}, \mathbf{v}) + K$, in time bounded polynomially in the encoding length of P and in r .*

Proof. This follows directly by combining Lemma 3.1 and Theorem 3.11. \square

There is one last issue to consider before proving our main theorem: Since the polynomial Hirsch conjecture remains open, a priori the combinatorial diameter of the input polytope P for k -DISTANCE ON SIMPLE POLYTOPES in our attempted hardness reduction may not be polynomially bounded, potentially rendering the time needed to construct the polytope Q in Theorem 3.12 superpolynomial. Hence, for the desired reduction to be polynomial, we must ensure that k -DISTANCE ON SIMPLE POLYTOPES when restricted to instances P with polynomially bounded diameter still remains NP-hard. However, this directly follows from the following lemma, showing that the fractional knapsack polytopes we used in our hardness proof for k -DISTANCE ON SIMPLE POLYTOPES in the first half of this paper do in fact have linear diameter.

Lemma 3.13. *For any choice of $\mathbf{b} \in \mathbb{Z}_{>0}^d$, $P_{\mathbf{b}}$ has combinatorial diameter at most $2(d+2)$.*

Proof. We will show that every vertex in the graph of $P_{\mathbf{b}}$ has distance at most $d+2$ to the vertex \emptyset , which will clearly imply the desired bound on the diameter.

Let us first consider a vertex of the form (S, i) where $S \subseteq [d+2]$ and $i \in [d+2] \setminus S$. Suppose first that $\sum_{i \in S} w_i \leq \beta + 1/4$. Then S is a neighbor of (S, i) via an edge of type (b) in Lemma 2.3. Removing all elements of S other than $d+1$ and then finally $d+1$ will lead to the vertex \emptyset after at most $|S| + 1 \leq (d+1) + 1 = d+2$ many steps using type (a) moves.

Suppose instead that $\sum_{i \in S} w_i \geq \beta + 1/4$. Then move to $S \cup \{i\}$ via a type (b) move and apply the same argument. Notice that since $[d+2]$ is not a vertex, we have $|S| + 1 = |S \cup \{i\}| \leq d+1$ in this situation. Hence, it takes a total of at most $|S| + 2 \leq (d+1) + 1 = d+2$ many steps to reach \emptyset from (S, i) , as desired.

For any vertex of the form $T \subseteq [d+2]$, the same decrementing procedure works and takes at most $|T| \leq d+2$ many steps to reach vertex \emptyset . Since $|S| \leq d+2$ and $|T| \leq d+2$. This concludes the proof. \square

Corollary 3.14. *k -DISTANCE ON SIMPLE POLYTOPES, restricted to input instances P of diameter at most $2 \dim(P) + 4$, is NP-hard.*

Proof. This follows directly from Lemma 3.13 and from our proof of Theorem 1.1. \square

Applying this theorem yields the proof of our second main result:

Proof of Theorem 1.5. By Corollary 3.14, k -DISTANCE ON SIMPLE POLYTOPES restricted to input instances P with diameter at most $2\dim(P) + 4$ is NP-hard. Given an input instance $P, \mathbf{u}, \mathbf{v}, k$ of this problem, we then compute $r := \max(2\dim(P) + 4, 6)$ and apply Theorem 3.12, using which we can construct, in polynomial time in the encoding length of P and in r (and hence simply in polynomial time in the encoding length of P) a simple polytope Q in inequality description and a number K such that $d_P(\mathbf{u}, \mathbf{v}) = \text{diam}(Q) - K$. Hence, given access to an oracle for DIAMETER OF SIMPLE POLYTOPES we can compute the distance between \mathbf{u} and \mathbf{v} on P and hence decide whether $d_P(\mathbf{u}, \mathbf{v}) \leq k$, in polynomial time in the encoding length of P and k plus the time needed for executing the oracle. Hence, we have found a Turing reduction from k -DISTANCE ON SIMPLE POLYTOPES restricted to input instances P with diameter at most $2\dim(P) + 4$ to DIAMETER OF SIMPLE POLYTOPES. Since the former is NP-hard, so is the latter, concluding the proof. \square

4. ROCK EXTENSIONS

In [27], Kaibel and Kukhareenko made the stunning observation that linear programming may be reduced in strongly polynomial time to the case of linear programs over a special family of simple polytopes called **rock extensions**, which have linear diameters. In the degenerate setting, this is trivial as one can simply take a pyramid over the original polytope, and the resulting polytope will have diameter 2. Hence, the notable feature of these polytopes is that they are simple. For understanding whether there exists a strongly polynomial time algorithm for linear programming, it suffices to study the case of linear programs over rock extensions.

For our purposes, the candidate algorithm we would be interested in is a path following algorithm like the simplex method that traverses the graph of the polytope. Hence, we ask the following question: Can one find a polynomial length path between any pair of vertices of a rock extension in strongly polynomial time? It turns out the answer is yes. However, one needs to be careful with the setup. A rock extension Q is built from a simple polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$, where A is $m \times d$ satisfying strong nondegeneracy assumptions and such that we know a strictly feasible point $o \in P$. The rock extension Q is a $(d+1)$ -dimensional simple extended formulation for P with $m+1$ facets and a distinguished vertex $(o, 1)$.

Our argument is essentially a corollary of their proof in [27]. However, the result was important enough context for ours that we include it here together with a proof, and it is not said in their paper. For brevity, we do not completely rewrite their construction here and instead only include the details of it relevant for the consequence we are interested in. We refer the reader to [27] for further, more explicit details about the construction.

Proof of Theorem 1.6. To prove this, we need to unpack the proof of Theorem 2.7 of [27] for constructing rock extensions. For $\mathbf{u} \in \mathbb{R}^d$ and $\varepsilon > 0$, let $B_\varepsilon^d(\mathbf{u})$ denote the d -dimensional open ball of radius ε centered at \mathbf{u} . In order to construct a rock extension, they start with a polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq \mathbf{b}\}$ and a point o such that $B_\varepsilon^d(o) \subseteq P$. Then they construct a so-called *rock extension*, which is a simple polytope defined as

$$Q = \{(\mathbf{x}, z) \in \mathbb{R}^{d+1} : A\mathbf{x} + \mathbf{y}z \leq \mathbf{b} \text{ and } z \geq 0\}$$

where $\mathbf{y} \in \mathbb{R}_{>0}^d$ is a suitably chosen vector which forms part of their construction.

In particular, the projection of Q onto its first d coordinates is exactly P . Their construction is then built in such a way that $(o, 1)$ will be a vertex of Q and the unique maximizer for the linear program $\max_{(\mathbf{x}, z) \in Q} z$. The way Q is built is inductive by adding one inequality at a time. Initially, up to a reordering of the rows, it is the simplex:

$$P_{d+1} = \{\mathbf{x} \in \mathbb{R}^{d+1} : A_{[d+1]}\mathbf{x} + \mathbf{y}_{[d+1]}z \leq \mathbf{b}_{[d+1]}, z \geq 0\}.$$

This simplex has one vertex with positive z coordinate, which is exactly $(o, 1)$. More generally,

$$P_k = \{\mathbf{x} \in \mathbb{R}^{d+1} : A_{[k]}\mathbf{x} + \mathbf{y}_{[k]}z \leq \mathbf{b}_{[k]}, z \geq 0\}$$

for each $k \geq d + 1$. For each P_k , there is a subset V_k of the vertices of P_k consisting of all vertices with positive z coordinate. As they complete their construction, they note that there is a sequence of strictly increasing values $0 < \mu_{d+1} < \mu_{d+2} < \dots < \mu_m = \varepsilon$ such that $V_k \setminus V_{k-1} \subseteq B_{\mu_k}^{d+1}((o, 1)) \setminus B_{\mu_{k-1}}^{d+1}((o, 1))$. The way they ensure this is by choosing y_k such that the hyperplane $H_k = \{\mathbf{x} \in \mathbb{R}^d | A_k \mathbf{x} + y_k z = b_k\}$ is supporting for the ball $B_{\mu_{k-1}}^{d+1}((o, 1))$ and arguing any new vertex created must not be too much further away.

By virtue of their construction, each vertex in V_k is a vertex of the rock extension, and the vertices of the rock extension are $V_m \cup V_{m+1}$, where

$$V_{m+1} = \{(\mathbf{v}, 0) : \mathbf{v} \text{ is a vertex of } Q\}.$$

Every vertex in V_{m+1} is adjacent to a vertex in V_m . This gives rise to a simple algorithm to find a path of length at most $2(m - d)$ between any pair of vertices of Q . To do this, it suffices to find a path of a length at most $m - d$ from any vertex to $(o, 1)$ efficiently. For this, simply move to the neighbor that is closest to $(o, 1)$.

Namely, let \mathbf{v} be a vertex of Q . Let $\mathbf{v} \in V_k$. Then \mathbf{v} has a neighbor in V_{k-1} , and any such neighbor is in $B_{\mu_{k-1}}^{d+1}((o, 1))$, while any other neighbor is not. Hence, its closest neighbor to $(o, 1)$ is in V_{k-1} . Since $V_{d+1} = \{(o, 1)\}$ this path will reach $(o, 1)$ in at most $m - d$ steps. Computing the closest neighbor to $(o, 1)$ may be done in strongly polynomial time if $(o, 1)$ is known, since by simplicity, each vertex has only $d + 1$ neighbors that may be computed using a simplex tableau. If $(o, 1)$ is not known, it can be found in weakly polynomial time by the linear program maximizing z . Therefore, if $(o, 1)$ is known there is a strongly polynomial time algorithm to find a path of length at most $2(m - d)$ between any pair of vertices on Q . Otherwise, it can be done in weakly polynomial time. \square

5. CONCLUSION

Knowing that finding shortest paths on a simple polytope is hard does not exclude the possibility that one may find short paths efficiently on general simple polytopes beyond rock extensions. For example, approximation algorithms may be possible. Given the relevant results in the literature, we suspect this is an APX-hard problem and leave proving APX-hardness as an open question. Our argument does not yield any interesting APX-hardness results as one can always efficiently find a path of length one more than the shortest path that we use to model our decision problem.

Finally, the core motivation for understanding these hardness questions is to approach the problem of whether there is a polynomial time version of the simplex method. In particular, one could show the answer is no conditional on $P \neq NP$ by showing that computing a polynomial length path in the graph of a simple polytope is NP-hard at least with a more standard Phase 1 procedure than that of constructing a rock extension. All hardness results thus far have relied on showing that determining the existence of a short path is hard. However, if the polynomial Hirsch conjecture holds, then a polynomial length path always exists and so this approach could not resolve the question of existence of a polynomial time simplex method. Our final open question is whether one can encode a hard search problem and prove TFNP-hardness for finding a short path on a simple polytope for which we know that short paths exist. This would, in particular, contrast with our observation for rock extensions.

ACKNOWLEDGMENTS

We would like to warmly thank Christian Nöbel and Laura Sanità for interesting and helpful discussions on this subject as well as Kirill Kukhareenko for help regarding rock extensions. The idea for the proof of the reduction to Partition arose in part from conversations with Chat GPT. However, all the work here was completely written and verified carefully by the (human) authors.

REFERENCES

1. I. Adler, C. Papadimitriou, and A. Rubinstein, *On simplex pivoting rules and complexity theory*, International Conference on Integer Programming and Combinatorial Optimization, Springer, 2014, pp. 13–24. [1](#)
2. O. Aichholzer, J. Cardinal, T. Huynh, K. Knauer, T. Mütze, R. Steiner, and B. Vogtenhuber, *Flip distances between graph orientations*, *Algorithmica* **83** (2021), no. 1, 116–143. [2](#)
3. N. Amenta and G. Ziegler, *Deformed products and maximal shadows*, *Contemporary Math.* **223** (1998), 57–90. [1](#)
4. D. Avis and V. Chvátal, *Notes on Bland’s pivoting rule*, *Polyhedral Combinatorics*, Springer, 1978, pp. 24–34. [1](#)
5. A. Black, *Exponential lower bounds for many pivot rules for the simplex method*, *Mathematical Programming* (2026). [1](#), [4](#)
6. A. Black, C. Nöbel, and R. Steiner, *Short circuit walks in fixed dimension*, Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2026, pp. 563–573. [2](#)
7. K. Borgwardt, *The simplex method: A probabilistic analysis*, vol. 1, Springer-Verlag, Berlin, 1987. [1](#)
8. S. Borgwardt, W. Grewe, S. Kafer, J. Lee, and L. Sanità, *On the hardness of short and sign-compatible circuit walks*, *Discrete Applied Mathematics* **367** (2025), 129–149. [2](#)
9. J. Cardinal and R. Steiner, *Inapproximability of shortest paths on perfect matching polytopes*, *Mathematical Programming* **210** (2025), no. 1, 147–163. [2](#)
10. ———, *Shortest paths on polymatroids and hypergraphic polytopes*, *Combinatorial Theory* **5(3)** (2025). [2](#), [4](#)
11. J. De Loera, S. Kafer, and L. Sanità, *Pivot rules for circuit-augmentation algorithms in linear optimization*, *SIAM Journal on Optimization* **32** (2022), no. 3, 2156–2179. [2](#)
12. Y. Disser, O. Friedmann, and A. Hopp, *An exponential lower bound for Zadeh’s pivot rule*, *Mathematical Programming* (2022). [1](#)
13. Y. Disser, G. Loho, M. Maat, and N. Mosis, *Lower bounds for ranking-based pivot rules*, arXiv preprint arXiv:2512.16684 (2025). [1](#)
14. Y. Disser and N. Mosis, *A Unified Worst Case for Classical Simplex and Policy Iteration Pivot Rules*, 34th International Symposium on Algorithms and Computation (ISAAC 2023) (Dagstuhl, Germany) (Satoru Iwata and Naonori Kakimura, eds.), Leibniz International Proceedings in Informatics (LIPIcs), vol. 283, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, pp. 27:1–27:17. [1](#)
15. Y. Disser and M. Skutella, *The simplex algorithm is np-mighty*, *ACM Transactions on Algorithms (TALG)* **15** (2018), no. 1, 1–19. [1](#)
16. J. Dorfer, *Flip distance of triangulations of convex polygons / rotation distance of binary trees is np-complete*, 2026. [2](#), [3](#)
17. J. Fearnley and R. Savani, *The complexity of the simplex method*, Proceedings of the forty-seventh annual ACM symposium on Theory of computing, 2015, pp. 201–208. [1](#)
18. O. Friedmann, *A subexponential lower bound for Zadeh’s pivoting rule for solving linear programs and games*, Proceedings of 15th International Conference on Integer Programming and Combinatorial Optimization (Oktay Günlük and Gerhard J. Woeginger, eds.), Springer, 2011, pp. 192–206. [1](#)
19. O. Friedmann, T. Hansen, and U. Zwick, *Subexponential lower bounds for randomized pivoting rules for the simplex algorithm*, Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, 2011, pp. 283–292. [1](#)
20. A. Frieze and S. Teng, *On the complexity of computing the diameter of polytope*, *Comput. Complex.* **4** (1994), no. 3, 207–219. [2](#), [4](#), [16](#)
21. D. Goldfarb, *Worst case complexity of the shadow vertex simplex algorithm*, preprint, Columbia University (1983). [1](#)
22. D. Goldfarb and W. Sit, *Worst case behavior of the steepest edge simplex method*, *Discrete Applied Mathematics* **1** (1979), no. 4, 277–285. [1](#)
23. T. Hansen and U. Zwick, *An improved version of the random-facet pivoting rule for the simplex algorithm*, Proceedings of the forty-seventh annual ACM symposium on Theory of computing, 2015, pp. 209–218. [1](#)
24. F. Holt and V. Klee, *Many polytopes meeting the conjectured hirsch bound*, *Discrete & Computational Geometry* **20** (1998), no. 1, 1–17. [4](#)
25. T. Ito, N. Kakimura, N. Kamiyama, Y. Kobayashi, and Y. Okamoto, *Shortest reconfiguration of perfect matchings via alternating cycles*, *SIAM Journal on Discrete Mathematics* **36** (2022), no. 2, 1102–1123. [2](#)

26. R. Jeroslow, *The simplex algorithm with the pivot rule of maximizing criterion improvement*, Discrete Mathematics **4** (1973), no. 4, 367–377. [1](#)
27. V. Kaibel and K. Kukharensko, *Rock extensions with linear diameters*, SIAM Journal on Discrete Mathematics **38** (2024), no. 4, 2982–3003. [5](#), [27](#)
28. V. Kaibel and M. Pfetsch, *Some algorithmic problems in polytope theory*, Algebra, geometry and software systems, Springer, 2003, pp. 23–47. [3](#)
29. G. Kalai, *A subexponential randomized simplex algorithm*, Proceedings of the twenty-fourth annual ACM symposium on Theory of computing, 1992, pp. 475–482. [1](#)
30. R. Karp, *Reducibility among combinatorial problems*, 50 Years of Integer Programming 1958-2008: from the Early Years to the State-of-the-Art, Springer, 2009, pp. 219–241. [6](#)
31. V. Klee and G. Minty, *How good is the simplex algorithm*, Inequalities : III : proceedings of the 3rd Symposium on inequalities (1972), 159–175. [1](#)
32. K. Kukharensko, *Short paths for the simplex algorithm*, Ph.D. thesis, Dissertation, Magdeburg, Otto-von-Guericke-Universität Magdeburg, 2025, 2025. [5](#)
33. J. Matoušek, M. Sharir, and E. Welzl, *A subexponential bound for linear programming*, Algorithmica **16** (1996), no. 4-5, 498–516. [1](#)
34. K. Murty, *Computational complexity of parametric linear programming*, Mathematical programming **19** (1980), no. 1, 213–219. [1](#)
35. B. Natura, *Circuit diameter of polyhedra is strongly polynomial*, arXiv preprint arXiv:2602.06958 (2026). [5](#)
36. C. Nöbel and R. Steiner, *Complexity of polytope diameters via perfect matchings*, Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2025, pp. 2234–2251. [2](#)
37. J. Orlin, *A polynomial time primal network simplex algorithm for minimum cost flows*, Mathematical Programming **78** (1997), no. 2, 109–129. [1](#)
38. L. Sanità, *The diameter of the fractional matching polytope and its hardness implications*, 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS), 2018, pp. 910–921. [2](#), [3](#)
39. F. Santos, *A counterexample to the Hirsch conjecture*, Annals of Mathematics (2012), 383–412. [1](#)
40. D. Sleator, R. Tarjan, and W. Thurston, *Rotation distance, triangulations, and hyperbolic geometry*, Proceedings of the eighteenth annual ACM symposium on Theory of computing, 1986, pp. 122–135. [3](#)
41. D. Spielman and S. Teng, *Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time*, Journal of the ACM (JACM) **51** (2004), no. 3, 385–463. [1](#)
42. Lasse Wulf, *Computing the polytope diameter is even harder than np-hard (already for perfect matchings)*, to appear in FOCS 2025 (2025). [2](#), [3](#)