

The diameter of the fractional matching polytope and its hardness implications

Laura Sanità

Department of Combinatorics & Optimization
University of Waterloo
Waterloo, ON, Canada
lsanita@uwaterloo.ca

Abstract—The (combinatorial) diameter of a polytope $P \subseteq \mathbb{R}^d$ is the maximum value of a shortest path between a pair of vertices on the 1-skeleton of P , that is the graph where the nodes are given by the 0-dimensional faces of P , and the edges are given the 1-dimensional faces of P . The diameter of a polytope has been studied from many different perspectives, including a computational complexity point of view. In particular, [Frieze and Teng, 1994] showed that computing the diameter of a polytope is (weakly) NP-hard.

In this paper, we show that the problem of computing the diameter is strongly NP-hard even for a polytope with a very simple structure: namely, the *fractional matching polytope*. We also show that computing a pair of vertices at maximum shortest path distance on the 1-skeleton of this polytope is an APX-hard problem. We prove these results by giving an *exact characterization* of the diameter of the fractional matching polytope, that is of independent interest.

Keywords—Polyhedral combinatorics, diameter of polytopes, matching theory.

I. INTRODUCTION

The (combinatorial) diameter of a polytope $P \subseteq \mathbb{R}^d$ is the maximum value of a shortest path between a pair of vertices on the 1-skeleton of P , which is the graph where the vertices correspond to the 0-dimensional faces of P , and the edges are given by the 1-dimensional faces of P . Giving bounds on the diameter of a polytope is a central question in discrete mathematics and computational geometry. Despite decades of studies, it is still not known whether the diameter of a d -dimensional polytope with n facets can be bounded by a polynomial function of n and d – this is currently referred to as the *polynomial Hirsch conjecture* [21]. Besides being a fundamental open question in polyhedral theory, the importance of the conjecture is also due to the fact that any polynomial pivot rule for the Simplex algorithm for Linear Programming requires the conjecture to hold. The existence of such rule would have significant consequences on the existence of a strongly-polynomial time algorithm for Linear Programming. The latter is a main open problem, mentioned in the list of the “mathematical problems for the next century” given by S. Smale [23].

The study of diameters of polytopes has a rich history, and we refer e.g. to [22] for a survey. The polynomial Hirsch conjecture is a generalization of a conjecture proposed by Hirsch in 1957, which states that the combinatorial diameter

of any d -dimensional polytope with n facets is at most $n-d$. This conjecture has been disproved first for unbounded polyhedra [15] and then for bounded polytopes [21], though it is known to hold for many classes of polytopes, such as for 0/1 polytopes [16]. For the currently best known upper bound on the diameter we refer to [20].

Besides providing general bounds, many researchers in the past 50 years have given bounds and/or characterizations of the diameter of polytopes that correspond to the set of feasible solutions of classical combinatorial optimization problems. Just to mention a few, such problems include matching [3, 8], TSP [18, 19], edge cover [13], fractional stable set [17], network flow and transportation problems [2, 5, 6, 7], stable marriage [12], and many more.

The diameter of a polytope has been studied from many different perspectives, including a computational complexity point of view. In particular, Frieze and Teng [9] showed that computing the diameter of a polytope is (weakly) NP-hard. Digging more into the complexity of computing the diameter of a polytope remains an interesting problem (see e.g. problem 10 on the list of 35 algorithmic problems in polytope theory, given in [14]).

In this paper, we show that computing the diameter of a polytope is a strongly NP-hard problem, and finding a pair of vertices at maximum (shortest path) distance on the 1-skeleton of a polytope is an APX-hard problem. In fact, what is probably more interesting, is that we can show hardness already for a polytope that has a very simple structure and it is quite well-understood: namely, the *fractional matching polytope*. We achieve these results by giving an exact characterization of the diameter of such polytope, which technically constitutes the main result of this paper, and is of independent interest. We are going to describe such characterization next.

One well-studied polytope for which a characterization of the diameter is known since the mid 70’s, is the *matching polytope*, that is the polytope given by the convex hull of the characteristic vectors of matchings of a graph. Formally, given a simple graph $G = (V, E)$ with n nodes and m edges, the matching polytope \mathcal{P}_M is as follows [11]:

$$\mathcal{P}_M := \{x \in \mathbb{R}^m : \begin{aligned} x(\delta(v)) &\leq 1 && \forall v \in V, \\ x(E(S)) &\leq \frac{|S|-1}{2} && \forall S \subseteq V : |S| \text{ odd}, \\ x &\geq 0 \end{aligned}\}. \quad (1)$$

Here $\delta(v)$ denotes the set of all edges of G incident into the node v , $E(S)$ denotes the set of edges with both endpoints in S , and for a set $F \subseteq E$, $x(F) = \sum_{e \in F} x_e$. The matching polytope is one of the most studied polytopes in combinatorial optimization. As shown in [3, 8] the diameter of the matching polytope is equal to the maximum cardinality of a matching in G , i.e. if we denote by $diam(\mathcal{P})$ and by $vert(\mathcal{P})$ the diameter and the set of vertices of a polytope \mathcal{P} , respectively, the result in [3, 8] states

$$diam(\mathcal{P}_M) = \max_{x \in vert(\mathcal{P}_M)} \{\mathbf{1}^T x\}.$$

The *fractional* matching polytope \mathcal{P}_{FM} is the polytope associated with the LP-relaxation of the standard IP formulation for the matching problem, and it is given by dropping the so-called odd-set inequalities from (1):

$$\mathcal{P}_{FM} := \{x \in \mathbb{R}^m : x(\delta(v)) \leq 1 \quad \forall v \in V, \quad x \geq 0\}. \quad (2)$$

As for the matching polytope, this polytope has been extensively studied in the optimization community. It is well known to be a half-integral polytope, and many structural results about its vertices, faces, and adjacency of the vertices are known. In particular, the support of a vertex x of \mathcal{P}_{FM} is the union of a matching, denoted by \mathcal{M}_x , given by the edges $\{e \in E : x_e = 1\}$, and a collection of node-disjoint odd cycles, denoted by \mathcal{C}_x , given by the edges $\{e \in E : x_e = \frac{1}{2}\}$, as shown in [1].

Our main result is a characterization of the diameter of this polytope, given in Theorem 1.

Theorem 1. *The diameter of the fractional matching polytope \mathcal{P}_{FM} is*

$$diam(\mathcal{P}_{FM}) = \max_{x \in vert(\mathcal{P}_{FM})} \left\{ \mathbf{1}^T x + \frac{|\mathcal{C}_x|}{2} \right\}.$$

As an observation, note that if the graph G is bipartite, then $\mathcal{C}_x = \emptyset$ for all vertices x of \mathcal{P}_{FM} , and our diameter bound for \mathcal{P}_{FM} reduces to the one for \mathcal{P}_M , as expected, since the odd-set inequalities are redundant for bipartite graphs.

As already mentioned, our result has an important algorithmic consequence regarding the hardness of computing the diameter of a polytope. As we will show later, combining Theorem 1 with an easy reduction from the NP-complete problem of finding a set of triangles partitioning the vertices of a graph, one can easily get that computing the diameter of a polytope is a strongly NP-hard problem.

Theorem 2. *Computing the diameter of a polytope is a strongly NP-hard problem.*

With some extra effort, the hardness result can be strengthened to get APX-hardness for the optimization problem of finding a pair of vertices at maximum distance on the 1-skeleton of a polytope. Due to space constraints, we omit the proof Theorem 3 from this extended abstract – it can be found in the full version of the paper.

Theorem 3. *Finding a pair of vertices at maximum (shortest path) distance on the 1-skeleton of a polytope is an APX-hard problem.*

We conclude this introduction with a brief description of the strategy we follow to give our characterization. We prove our result in two steps: first, we show that the value given in Theorem 1 is an *upper* bound on the diameter, and then we show that it is a *lower* bound on the distance between two specific vertices of \mathcal{P}_{FM} . Proving the first step in particular requires some effort.

The upper bound proof is based on a *token* argument. Specifically, given two distinct vertices z and y of \mathcal{P}_{FM} , we prove that the distance between z and y is bounded by $\mathbf{1}^T w + \frac{|\mathcal{C}_w|}{2}$ for some vertex w of \mathcal{P}_{FM} whose support graph is in the union of the support graphs of z and y . We will assign to each node $v \in V$ with $w(\delta(v)) = 1$ a token of value $\frac{1}{2}$, and to each cycle $C \in \mathcal{C}_w$ a token of value $\frac{1}{2}$ (note that the total sum of the token values equals $\mathbf{1}^T w + \frac{|\mathcal{C}_w|}{2}$). We will then define a path on the 1-skeleton of \mathcal{P}_{FM} from z to y , and show that each *move* along this path can be payed by using two tokens from some nodes or cycles, where a move refers to the process of traversing an edge on the 1-skeleton of \mathcal{P}_{FM} .

We would like to highlight that the selection of the moves to take is not straightforward. A trivial attempt to define a z - y path could be to (i) define a path from z to a 0/1-vertex \bar{z} obtained by “rounding” the half-valued coordinates of one of the cycles in \mathcal{C}_z at each step, (ii) define a path from y to a 0/1-vertex \bar{y} obtained by “rounding” the half-valued coordinates of one of the cycles in \mathcal{C}_y at each step, and (iii) define a path from \bar{z} to \bar{y} guided by the symmetric difference between the matchings corresponding to \bar{z} and \bar{y} (indeed, this corresponds to a \bar{z} - \bar{y} path on the 1-skeleton of the matching polytope \mathcal{P}_M). Unfortunately, it is not difficult to construct examples where *any* path of this form has a length strictly larger than the claimed bound. In order to reduce the number of moves, we will have to exploit better the adjacency properties of the vertices, and we will have to keep track of the tokens we use to pay for our moves in a very careful way.

II. PRELIMINARIES

In this section, we will introduce some notations and state some known structural results regarding the vertices of the

polytope \mathcal{P}_{FM} .

To avoid confusion, we will always refer to the extreme points of a given polytope as *vertices*, and to the elements of V of a given graph $G = (V, E)$ as *nodes*. Furthermore, for a generic (sub)graph H , we will denote by $V(H)$ and $E(H)$ its set of nodes and edges, respectively. We say that a cycle C (resp. a path P) is odd if $|E(C)|$ (resp. $|E(P)|$) is odd. Given a matching M , an M -alternating path is a path that alternates edges in M and edges not in M . A node is M -exposed if it is not the endpoint of an edge in M . An M -augmenting path is an M -alternating path whose endpoints are M -exposed. We start by stating the well-known half-integrality property of the vertices of \mathcal{P}_{FM} .

Theorem 4 ([1]). *Every basic feasible solution to \mathcal{P}_{FM} has components equal to $0, 1$ or $\frac{1}{2}$, and the edges with half-integral components induce node-disjoint odd cycles. Furthermore, every half-integral vector $x \in \mathcal{P}_{FM}$ such that the half-integral components of x induce node-disjoint odd cycles, is a vertex of \mathcal{P}_{FM} .*

For a given vector $x \in \mathbb{R}^E$, we will refer to \mathcal{G}_x as the graph induced by the support of x , i.e. the graph induced by the set of edges $\{e \in E : x_e > 0\}$. As already mentioned in the introduction, if x is a vertex of \mathcal{P}_{FM} , then \mathcal{G}_x is the union of a matching (denoted by \mathcal{M}_x) induced by the edges $\{e \in E : x_e = 1\}$ and a collection of node-disjoint odd cycles (denoted by \mathcal{C}_x) induced by the edges $\{e \in E : x_e = \frac{1}{2}\}$. Furthermore, given two vectors $x, y \in \mathbb{R}^E$, we let $\mathcal{G}_x \Delta \mathcal{G}_y$ be the graph induced by the set of edges $\{e \in E : x_e \neq y_e\}$.

The following definition will be highly used later.

Definition 1. *We say that an odd cycle C of G is packed by a matching M (resp. by a vertex x), if $|M \cap E(C)| = \frac{|C|-1}{2}$ (resp. if $|\mathcal{M}_x \cap E(C)| = \frac{|C|-1}{2}$).*

The paper in [4] gives an adjacency characterization for the vertices of the fractional *perfect* b -matching polytope associated to a (possibly non-simple) graph. Using this result, it is easy to derive adjacency properties for the vertices of \mathcal{P}_{FM} . We here explicitly state a lemma that follows from Theorem 25 in [4]. This lemma gives some sufficient conditions for two vertices of \mathcal{P}_{FM} to be adjacent (we omit the proof from this extended abstract). These are all the conditions that we will use in Section III to prove our upper bound.

Lemma 1. *(follows from [4]) Let y and z be two vertices of \mathcal{P}_{FM} , and let $\mathcal{C}_z(y) \subseteq \mathcal{C}_z$ (resp. $\mathcal{C}_y(z) \subseteq \mathcal{C}_y$) be the set of cycles of \mathcal{G}_z (resp. \mathcal{G}_y) packed by \mathcal{M}_y (resp. packed by \mathcal{M}_z). Then y and z are adjacent if $\mathcal{G}_y \Delta \mathcal{G}_z$ contains exactly one component C with this component being*

- (a) *one even cycle in $\mathcal{M}_z \Delta \mathcal{M}_y$, or*
- (b) *one path in $\mathcal{M}_z \Delta \mathcal{M}_y$, or*
- (c) *one odd cycle in $\mathcal{C}_z(y) \cup \mathcal{C}_y(z)$, or*
- (d) *one odd cycle in $\mathcal{C}_z(y)$ and one odd cycle in $\mathcal{C}_y(z)$*

intersecting in exactly one node, or

- (e) *two node-disjoint odd cycles in $\mathcal{C}_z(y) \cup \mathcal{C}_y(z)$ and a path $P \subseteq \mathcal{M}_y \Delta \mathcal{M}_z$ intersecting the cycles at its endpoints, or*
- (f) *one odd cycle in $\mathcal{C}_z(y) \cup \mathcal{C}_y(z)$ and a path $P \subseteq \mathcal{M}_y \Delta \mathcal{M}_z$ intersecting the cycle at one of its endpoints.*

Recall that we refer to the process of traversing an edge of a polytope (i.e. moving from one vertex to an adjacent one) as a *move*.

III. A CHARACTERIZATION OF THE DIAMETER OF \mathcal{P}_{FM}

In this section we give a characterization of the diameter of the fractional matching polytope by providing a proof of Theorem 1. Recall that Theorem 1 states

$$\text{diam}(\mathcal{P}_{FM}) = \max_{x \in \text{vert}(\mathcal{P}_{FM})} \{ \mathbf{1}^T x + \frac{|\mathcal{C}_x|}{2} \}.$$

We will prove that the right-hand side of the above formula is an upper bound on the diameter of \mathcal{P}_{FM} in Sections III-A up to III-F, and then prove it is a lower bound in Section III-G.

A. Upper bound: general strategy

Let us recall the strategy sketched in the introduction, regarding how we are going to prove our upper bound. Given two distinct vertices z and y of \mathcal{P}_{FM} , we will prove that the distance between z and y is bounded by $\mathbf{1}^T w + \frac{|\mathcal{C}_w|}{2}$ for some vertex w of \mathcal{P}_{FM} that will be identified later.

We assign to each node in v in $V(\mathcal{G}_w)$ (i.e. with $w(\delta(v)) = 1$) a token of value $\frac{1}{2}$, and to each cycle $C \in \mathcal{C}_w$ a token of value $\frac{1}{2}$, so that the total sum of the token values equals $\mathbf{1}^T w + \frac{|\mathcal{C}_w|}{2}$. We define a path on the 1-skeleton of \mathcal{P}_{FM} from z to y , and show that each move along this path can be payed by using two tokens from some nodes or cycles. Specifically, our path will have the form

$$z \rightarrow w \rightarrow r \rightarrow y$$

where r and w are (non necessarily distinct) vertices of \mathcal{P}_{FM} with some nice structure, and the arrow “ \rightarrow ” indicates a path between the corresponding vertices on the 1-skeleton of \mathcal{P}_{FM} . While traversing our path, we will satisfy some invariants, which will be helpful for keeping track of the used tokens. The first set of invariants, valid for the path from z to w , are described in the next paragraph.

Invariants and definitions: For every pair of consecutive vertices $\ell, \bar{\ell}$ on the path $z \rightarrow w$ (with $\bar{\ell}$ coming after ℓ when traversing this path), the following invariant will hold:

$$V(\mathcal{G}_\ell) \subseteq V(\mathcal{G}_{\bar{\ell}}). \quad (3)$$

In other words, $\bar{\ell}$ is obtained by “augmenting” ℓ . This implies that each node u with $\ell(\delta(u)) = 1$ will satisfy $w(\delta(u)) = 1$, and therefore u has been assigned a token. It follows that we can use the tokens of the nodes of $V(\mathcal{G}_{\bar{\ell}})$ to pay for the move from ℓ to $\bar{\ell}$.

We also anticipate here that for every pair of consecutive vertices $\ell, \bar{\ell}$ on the path $z \rightarrow w$ (with $\bar{\ell}$ coming after ℓ when traversing this path), the following invariant will hold:

$$(\mathcal{C}_\ell \setminus \mathcal{C}_z) \subseteq \mathcal{C}_{\bar{\ell}}. \quad (4)$$

This implies that *if* an odd cycle C appears for the first time in the support of a vertex $\ell \neq z$, then C will be part of the support of w as well. This ensures that C has been assigned a token, and therefore we can use its token to pay for the move we performed to arrive to ℓ .

In the following, we will denote by $T(\ell)$ the subset of nodes of $V(\mathcal{G}_\ell)$ that still have an available token, i.e., a token that has *not* been used to pay for any of the moves performed to arrive from z to ℓ . Furthermore, for a given vertex ℓ , we let $\tilde{M}_\ell \subseteq \mathcal{M}_\ell$ be the subset of edges of \mathcal{M}_ℓ with both endpoints not in $T(\ell)$, i.e. satisfying: $\{u, v\} \in \tilde{M}_\ell \Leftrightarrow \{u, v\} \in \mathcal{M}_\ell$ and $u, v \notin T(\ell)$.

On our path from z to w , we will satisfy two additional invariants. The first one states that the edges of \mathcal{M}_ℓ whose endpoints have no tokens are a subset of the edges of \mathcal{G}_y :

$$\tilde{M}_\ell \subseteq E(\mathcal{G}_y). \quad (5)$$

The second one states that the nodes that used their tokens either can be “paired up” using edges of \tilde{M}_ℓ , or they belong to cycles in $\mathcal{C}_\ell \cap \mathcal{C}_y$:

$$\text{for each } v \in V(\mathcal{G}_\ell) \setminus V(\mathcal{C}_\ell \cap \mathcal{C}_y), \text{ we have } v \in T(\ell) \Leftrightarrow v \text{ is not an endpoint of an edge in } \tilde{M}_\ell. \quad (6)$$

B. Moving from z to w

The vertex w is obtained by “augmenting” z (if possible) using edges of \mathcal{G}_y . This is done in three main steps. To describe the first one, we need to introduce a definition (refer to Figure 1(a)).

Definition 2. Let ℓ be a vertex, and C be a cycle in $\mathcal{C}_y \setminus \mathcal{C}_\ell$, such that $|V(C) \cap T(\ell)| \geq 2$. We say that the cycle C is critical for ℓ if every odd path $P \subseteq C$ with nodes $\{v_1, \dots, v_k\}$ such that $V(P) \cap T(\ell) = \{v_1, v_k\}$, satisfies

- (i) $|E(P)| \geq 3$, and
- (ii) either $\{v_1, v_k\} \in \mathcal{M}_\ell$, or $v_1, v_k \in V(\bar{C})$ for some $\bar{C} \in \mathcal{C}_\ell$.

Algorithm 1 describes the moves we perform to arrive to the vertex w , starting from z , and proceeds as follows. In Step 2 we consider the current vertex ℓ (with $\ell := z$ at the beginning), and we look for a critical cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$. Given such a cycle, we perform a move which increases the number of nodes covered by \mathcal{M}_ℓ (see Figure 1(b)). In Step 3 we look for \tilde{M}_ℓ -augmenting paths whose edges are in $E(\mathcal{G}_y)$ and whose endpoints are not in $V(\mathcal{G}_\ell)$. If such a path is identified, we perform a move which again increases the number of nodes covered by \mathcal{M}_ℓ . In Step 4 we look for a cycle $C \in \mathcal{C}_y$, which is packed by ℓ , and has a node

$v \in V(C)$ which is not in $V(\mathcal{G}_\ell)$. If such a cycle is identified, we perform a move which increases the number of cycles in $\mathcal{C}_\ell \cap \mathcal{C}_y$.

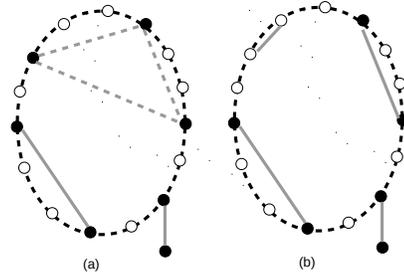


Figure 1. A critical cycle is shown in (a). Black edges represent the edges in the support of y , and gray edges represent the edges in the support of ℓ . Continuous lines represent edges of value 1, while dashed lines represent edges of value $\frac{1}{2}$. Nodes in $T(\ell)$ are colored black. The cycle contains three odd paths whose endpoints are black, and internal nodes are white. All these three paths satisfy conditions (i) and (ii) of Definition 2. The figure in (b) shows how the coordinates of ℓ change after performing one move $\langle \triangleleft \rangle$.

Algorithm 1 (from z to w):

1. Set $\ell := z$;
2. While there is a cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$ that is critical for ℓ :
 - 2.1 Let $P \subseteq C$ be an odd path with nodes $\{v_1, \dots, v_k\}$, such that $V(P) \cap T(\ell) = \{v_1, v_k\}$;
 - 2.2 If $\{v_1, v_k\} \in \mathcal{M}_\ell$, do the following move:
 - $\langle \triangleright \rangle$ Change the ℓ -coordinate of the edges $\{v_1, v_2\}, \{v_1, v_k\}, \{v_k, v_{k-1}\}$ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e = \{v_1, v_2\} \text{ or } e = \{v_k, v_{k-1}\} \\ 0 & \text{if } e = \{v_1, v_k\} \end{cases}$$

and let v_2 and v_1 use their tokens to pay for this move;

- 2.3 If $v_1, v_k \in V(\bar{C})$ for some $\bar{C} \in \mathcal{C}_\ell \setminus \mathcal{C}_y$ with $V(\bar{C}) = \{u_1 := v_1, u_2, \dots, u_{\bar{k}}\}$, do the move:
 - $\langle \triangleleft \rangle$ Change the coordinate ℓ_e of all the edges $e \in E(\bar{C}) \cup \{v_1, v_2\}$ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e = \{v_1, v_2\} \\ 1 & \text{if } e = \{u_{2i}, u_{2i+1}\} \text{ for } i \in \{1, \dots, \frac{\bar{k}-1}{2}\} \\ 0 & \text{if } e = \{u_{2i-1}, u_{2i}\} \text{ for } i \in \{1, \dots, \frac{\bar{k}-1}{2}\} \\ 0 & \text{if } e = \{v_{\bar{k}}, v_1\} \end{cases}$$

and let v_2 and v_1 use their tokens to pay for this move;

3. While there is an \tilde{M}_ℓ -augmenting path P in \mathcal{G}_y , with endpoints $u, v \notin V(\mathcal{G}_\ell)$, do the following move:
 - $\langle \circ \rangle$ Change the coordinate ℓ_e of all the edges of P by setting

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(P) \setminus \tilde{M}_\ell \\ 0 & \text{if } e \in E(P) \cap \tilde{M}_\ell \end{cases}$$

and let u and v use their tokens to pay for this move;

4. While there is a cycle $C \in \mathcal{C}_y$ that (i) is packed by ℓ , and (ii) contains a node $v \in V(C)$ with $v \notin \mathcal{V}(\mathcal{G}_\ell)$ do the following move:

(\diamond) Change the coordinates ℓ_e of the edges $e \in E(C)$ to $\frac{1}{2}$, and let the cycle C and the node v use their tokens to pay for this move;

5. Output $w := \ell$.

Lemma 2. *All steps of Algorithm 1 can be correctly performed, and invariants (3), (4), (5), and (6) are maintained.*

Proof: We first prove that Step 2 of Algorithm 1 can be correctly performed and maintains all the invariants, by induction on the number of moves executed.

Let C be a critical cycle for the vertex ℓ , and P be the corresponding path identified in Step 2.1. First, suppose that $\{v_1, v_k\} \in \mathcal{M}_\ell$ and let us focus on the move (\triangleright). We claim that the edges $\{v_1, v_2\}, \{v_1, v_k\}, \{v_k, v_{k-1}\}$ form an \mathcal{M}_ℓ -augmenting path, and therefore the move is a valid move, according to Lemma 1(b). To see this, note first that $v_{k-1} \neq v_2$, since by definition P has at least 3 edges, and so $k \geq 4$. Second, we claim that both v_2 and v_{k-1} are not in $V(\mathcal{G}_\ell)$. Assume for a contradiction that $v_2 \in V(\mathcal{G}_\ell)$ (the other case is similar). Since C is critical, we know that $v_2 \notin T(\ell)$, and therefore by invariants (5) and (6) (which hold by the inductive hypothesis), $\{v_2, v_3\} \in \bar{\mathcal{M}}_\ell$. Let ℓ' be the iteration where $\{v_2, v_3\}$ appears for the first time in the support of a vertex, and $\bar{\ell}$ be the vertex visited by the algorithm immediately before ℓ' . Then, necessarily C was critical for $\bar{\ell}$ and one of the two nodes v_2, v_3 was in $T(\bar{\ell})$. However, if $v_2 \in T(\bar{\ell})$, then the path with one single edge $\{v_1, v_2\}$ would contradict the fact that C was critical for $\bar{\ell}$. It follows that $v_3 \in T(\bar{\ell})$. However, the path \bar{P} identified at Step 2.1 of the algorithm has to contain the edge $\{v_2, v_3\}$, and its endpoints are v_3 and another node in $T(\bar{\ell})$. Necessarily, the other endpoint of \bar{P} is then v_1 . However, this contradicts the fact that \bar{P} is odd. It follows that $\{v_1, v_2\}, \{v_1, v_k\}, \{v_k, v_{k-1}\}$ form an \mathcal{M}_ℓ -augmenting path, and it is easy to see then that the move (\triangleright) maintains the claimed invariants.

Now, suppose that $\{v_1, v_k\} \in V(\bar{C})$ for some $\bar{C} \in \mathcal{C}_\ell \setminus \mathcal{C}_y$ and let us focus on the move (\triangleleft). Using an argument identical to the one used in the previous point, we can show that $v_2 \notin V(\mathcal{G}_\ell)$, and therefore the edge $\{v_1, v_2\}$ and the cycle \bar{C} have the structure described in Lemma 1(f), which allows us to perform a valid move. Once again, it is immediate to see that the move (\triangleright) maintains the claimed invariants.

We now argue that the other steps of Algorithm 1 can be correctly performed and maintain all the invariants. Each move in (\circ) and (\diamond) is indeed a valid move, according to Lemma 1(b) and Lemma 1(c), respectively. It is easy to

see that the all invariants are maintained in these steps. Furthermore, by invariants (3) and (4), each node v and each cycle C that used a token during the execution of Algorithm 1, satisfy $v \in V(\mathcal{G}_w)$ and $C \in \mathcal{C}_w$, and therefore had indeed a token to use. ■

We state a trivial observation which will be used later.

Observation 1. *Let C be a cycle in $\mathcal{C}_w \setminus \mathcal{C}_y$. Then C did not use its token during the execution of Algorithm 1.*

Before describing the subsequent moves on our path from w to y , we need to introduce an important notion: namely, the notion of *witnesses* of a cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_w$. This notion will be crucial to identify the nodes that will pay for the move in which the cycle C appears for the first time in the support of a vertex on our path from w to y .

C. Witnesses

Definition 3. *Let u be a node of a cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_w$. We say that u is a single witness of C if $u \in T(w)$, u is $(\mathcal{M}_w \cap E(C))$ -exposed, and C is packed by w .*

Definition 4. *Let (u, v) be a pair of nodes of a cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_w$. We say that (u, v) is a pair of witnesses of C if $u, v \in T(w)$ and there is an $\bar{\mathcal{M}}_w$ -augmenting path $Q(u, v)$ with endpoints u and v in $E(C)$.*

Note that for a given pair of witnesses (u, v) , the path $Q(u, v)$ is uniquely defined (since exactly one u - v path in $E(C)$ has odd length).

As already mentioned, we would like that the witnesses of a cycle C pay for the move in which the cycle C appears for the first time in the support of a vertex on our path from the vertex w to the vertex y . Unfortunately, it might be possible that if we do not select our witnesses carefully, we do not have enough tokens to perform our moves properly. We therefore impose some restrictions on our choice.

Definition 5. *We say that $\mathcal{W} \subseteq V \times V$ is a good set of witnesses if it satisfies*

- (i) *for all $C \in \mathcal{C}_y \setminus \mathcal{C}_w$, the following holds: \mathcal{W} contains either exactly one pair of witnesses (u, v) of C , or exactly one single witness of C (indicated as $(u, u) \in \mathcal{W}$);*
- (ii) *for every $(u, v) \in \mathcal{W}$ with $u \neq v$, the following holds: if $|E(Q(u, v))| > 1$, then u and v belong to two distinct components of \mathcal{G}_w .*

In other words, condition (ii) of the above definition states that if $Q(u, v)$ has more than one edge, then $\{u, v\} \notin \mathcal{M}_w$, and there is no cycle $\bar{C} \in \mathcal{C}_w$ such that $V(\bar{C})$ contains both u and v . Next lemma shows that there exists a good set of witnesses.

Lemma 3. *There exists a good set of witnesses \mathcal{W} .*

Proof: It is enough to show that for every $C \in \mathcal{C}_y \setminus \mathcal{C}_w$, we can find either a single witness, or a pair of witnesses which satisfies the condition in (ii).

Let C be any cycle in $\mathcal{C}_y \setminus \mathcal{C}_w$. By definition, each edge $e \in \tilde{M}_w$ has its endpoints not in $T(w)$. Therefore, if there is an edge $\{u, v\} \in E(C)$ such that $u, v \in T(w)$ then the edge $\{u, v\}$ is clearly an \tilde{M}_w -augmenting path from u to v , and (u, v) are a pair of witnesses for C which satisfies the condition in (ii).

Let us assume there is no edge $\{u, v\} \in E(C)$ such that $u, v \in T(w)$. Consider the set $H := \{v \in V(C) : v \in T(w)\}$. First, we are going to show that $H \neq \emptyset$. Suppose otherwise. Then, by invariant (6), every node in $v \in V(C)$ satisfies either $w(\delta(v)) = 0$, or $w(\delta(v)) = 1$ and v is an endpoint of an edge in \tilde{M}_w . Note that, by invariant (5), if v is an endpoint of an edge $\{v, \bar{v}\}$ in \tilde{M}_w , then \bar{v} is also a node of C . Since $|V(C)|$ is odd but \tilde{M}_w is a matching, it follows that C contains either an \tilde{M}_w -augmenting path whose endpoints are not in $V(\mathcal{G}_w)$, or C is packed by \mathcal{M}_w and one node of C is not in $V(\mathcal{G}_w)$. In both cases, Algorithm 1 would not have stopped, a contradiction.

It follows that $H \neq \emptyset$. Consider the connected components C_1, \dots, C_k in the graph $C \setminus H$. Since the nodes in H are pairwise not adjacent in $E(C)$, and C is a cycle, we have that $|V(C)| = |H| + \sum_{i=1}^k |V(C_i)| = k + \sum_{i=1}^k |V(C_i)|$. Since $|V(C)|$ is odd, at least one component C_i has $|V(C_i)|$ equal to an even number. It follows that C_i is a path with an odd number of edges. By Step 3 of Algorithm 1, \tilde{M}_w has to contain a maximum matching of this component. It follows that C_i is an \tilde{M}_w -alternating path with the first and the last edge belonging to \tilde{M}_w . Let u (resp. v) be the node in H adjacent to the first (resp. last) edge of the path C_i in the graph C . If $u = v$ (i.e. $k = 1$), then C is packed by w , and u is a single witness for C . If instead $u \neq v$, then u, C_i, v yield an \tilde{M}_w -augmenting path from u to v in the subgraph induced by $\tilde{M}_w \cup E(C)$.

The above argument shows that for every cycle $C \in \mathcal{C}_w \setminus \mathcal{C}_y$ there exists either a single witness, or (at least) one pair of witnesses. It remains to show that, if C is not packed by w , then among all possible pairs of nodes that satisfy the definition of pair of witnesses, at least one satisfies the condition described in (ii). For the sake of a contradiction, assume that no pairs of witnesses satisfies the condition described in (ii). Then, this means that C is critical for the vertex w . Let $\bar{\ell}$ be the last vertex visited by Algorithm 1 during the execution of Step 2. Clearly, C cannot be critical for $\bar{\ell}$, otherwise the algorithm would have performed another iteration of Step 2. However, note that for all vertices ℓ' visited after $\bar{\ell}$ in Step 3 and Step 4, we have $T(\bar{\ell}) = T(\ell')$, and therefore C cannot be critical for any ℓ' . Since w is also visited after $\bar{\ell}$, C cannot be critical for w , a contradiction. ■

From now on, we fix \mathcal{W} to be a good set of witnesses (one such set exists because of last lemma), and based on that, we introduce the last two ingredients needed to describe our future moves: *target matchings*, and the *target graph*.

D. Target matchings and target graph

Definition 6. Let $C \in \mathcal{C}_y \setminus \mathcal{C}_w$. The target matching of C , denoted by M_C , is a matching that satisfies the following properties:

- (i) M_C is a maximum cardinality matching of C ;
- (ii) if $(u, v) \in \mathcal{W}$ is a pair of witnesses for C with $u \neq v$, then $M_C \cap E(Q(u, v))$ is a perfect matching of $Q(u, v)$;
- (iii) among all matchings satisfying (i) and (ii), M_C maximizes the quantity $|M_C \cap \mathcal{M}_w|$.

Definition 7. Let ℓ be a vertex. The target graph \mathcal{T}^ℓ is the graph induced by the edges

$$\mathcal{M}_\ell \Delta (\mathcal{M}_y \cup \{ \bigcup_{C \in \mathcal{C}_y \setminus \mathcal{C}_\ell} M_C \})$$

Note that the target graph is the symmetric difference of two matchings, namely \mathcal{M}_ℓ , and $\mathcal{M}_y \cup \{ \bigcup_{C \in \mathcal{C}_y \setminus \mathcal{C}_\ell} M_C \}$, and therefore it is the disjoint union of paths and even cycles. We call a component K of \mathcal{T}^ℓ a *path-component* if K is a path, and a *cycle-component* if K is a cycle.

Roughly speaking, our goal is to move from w to a vertex whose support graph does not contain any cycle in $\mathcal{C}_w \setminus \mathcal{C}_y$, by performing a sequence of moves, each possibly involving one component of the target graph. However, we would like not to use tokens belonging to witness nodes to pay for the moves, since as already mentioned, we would like to keep these tokens to pay for the moves where cycles in $\mathcal{C}_y \setminus \mathcal{C}_w$ show up. For this reason we introduce the following definition.

Definition 8. Let ℓ be a vertex. A component K of a target graph \mathcal{T}^ℓ is called *dangerous* if all the nodes of $V(K) \cap T(\ell)$ are witnesses of some cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$.

The next lemma gives a few properties that will be crucial for our analysis.

Lemma 4. Let K be a component of \mathcal{T}^w . Then:

- (i) $V(K) \cap T(w) \neq \emptyset$;
- (ii) If K is a dangerous cycle-component, then there exist at least two distinct cycles in $\mathcal{C}_y \setminus \mathcal{C}_w$ whose witnesses are in $V(K)$;
- (iii) If K is a dangerous path-component, then each endpoint v of K is either a single witness in \mathcal{W} , or $v \in V(\bar{C})$ for some cycle $\bar{C} \in \mathcal{C}_w \setminus \mathcal{C}_y$.

Proof: Let K be a component of the target graph \mathcal{T}^w . We first prove (i). Assume $V(K) \cap T(w) = \emptyset$. Then $E(K) \cap \mathcal{M}_w \subseteq \tilde{M}_w$, by invariant (6). Since $\tilde{M}_w \subseteq E(\mathcal{G}_y)$ by invariant (5), it follows that $K \subseteq \mathcal{G}_y$, and therefore K is a path. If K has even length, then $K \subset C$ for some $C \in \mathcal{C}_y$, and it is an M_C -alternating path. By switching along the edges of this path we could get another matching M'_C which would contradict our choice of the target matching for C , since M'_C would satisfy the conditions (i) and (ii) of Definition 6, but $|M'_C \cap \tilde{M}_w| > |M_C \cap \tilde{M}_w|$. It follows

that K has odd length, and since $\mathcal{M}_y \cup \{\cup_{C \in \mathcal{C}_y} M_C\}$ is a maximum matching in \mathcal{G}_y , K has to be \tilde{M}_w -augmenting. Then, if both endpoints c', c of K are not in $T(w)$, then $w(\delta(c')) = w(\delta(c)) = 0$ (since otherwise they would still have their token available, being non-endpoints of edges in \tilde{M}_w , by invariant (6)). However, this contradicts the termination of Step 3 of Algorithm 1. It follows that at least one of the endpoints has to be in $T(w)$.

Now we prove (ii). Let K be a dangerous cycle-component, and v be a node in $V(K) \cap T(w)$ (such node exists because of (i)). Since K is dangerous, v is a witness of a cycle $C \in \mathcal{C}_y$. Since K is a cycle, v cannot be a single witness (such nodes have degree at most 1 in the target graph). It follows that there is another node u such that (u, v) is a pair of witnesses for C . Let v' (resp. u') be the node such that $\{v, v'\} \in \mathcal{M}_w$ (resp. $\{u, u'\} \in \mathcal{M}_w$). Note that $v' \neq u$ (resp. $u' \neq v$), since otherwise \mathcal{W} would not be a good set of witnesses (the pair (u, v) would contradict the second condition of Definition 5). By invariant (5) both $u', v' \in T(w)$, and therefore, since K is dangerous, they are (in a pair of) witnesses for at least another cycle in \mathcal{C}_y .

Finally we prove (iii). Let K be a dangerous path-component, with $V(K) := \{v_1, \dots, v_k\}$. Let i be the smallest index such that $v_i \in T(w)$ (such index exists because of (i)). Since K is dangerous, v_i is a witness (either a single one, or a paired one) of a cycle $C \in \mathcal{C}_y$. If v_i is a single witness, then it is necessarily the endpoint of K , i.e. $i = 1$, since single witnesses have degree at most one in the target graph. Assume now that v_i is in a pair of witnesses. In this case, v_i cannot be an endpoint of an edge in \mathcal{M}_w : using invariant (6), we can see that if $\{v_i, v_{i+1}\} \in \mathcal{M}_w$, then necessarily v_{i+1} is the other witness of C , i.e. $Q(v_i, v_{i+1}) = \{v_i, v_{i+1}\}$ contradicting that this edge is in the target graph, and if $\{v_i, v_{i-1}\} \in \mathcal{M}_w$, our choice of i is contradicted. It follows that $v_i \in V(\bar{C})$ for some $\bar{C} \in \mathcal{C}_w \setminus \mathcal{C}_y$. This implies that v_i have degree one in the target graph, and therefore it is the endpoint of K (i.e. $i = 1$). We can apply the same argument to the biggest index such that $v_i \in T(w)$, and get the same conclusion for the other endpoint of K . ■

E. Moving from w to r

We move from w to a vertex r with the property that $\mathcal{C}_r \subseteq \mathcal{C}_y$, by eliminating up to two cycles in $\mathcal{C}_w \setminus \mathcal{C}_y$ at each move.

The algorithm maintains six invariants for every vertex ℓ visited during its execution. The first three invariants guarantee that the conditions of Lemma 4 hold for every vertex visited by the algorithm:

For every component K of \mathcal{T}^ℓ , $V(K) \cap T(\ell) \neq \emptyset$. (7)

If K is a dangerous cycle-component of \mathcal{T}^ℓ , then $V(K)$ contains the witnesses of at least two cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$. (8)

If K is a dangerous path-component of \mathcal{T}^ℓ , then each endpoint v of K is either a single witness in \mathcal{W} , or $v \in V(\bar{C})$ for some cycle $\bar{C} \in \mathcal{C}_\ell \setminus \mathcal{C}_y$. (9)

The fourth and fifth invariants will be useful to keep track of which nodes in \mathcal{T}^ℓ still have their tokens. The fourth invariant states that for every edge of \mathcal{M}_ℓ that is in the target graph, either both the endpoints have the token, or both the endpoints used their tokens.

$\forall \{u, v\} \in \mathcal{M}_\ell \cap E(\mathcal{T}^\ell)$, we have $u \in T(\ell) \Leftrightarrow v \in T(\ell)$. (10)

Note that (10) holds for w , because of invariant (6) maintained by Algorithm 1. The fifth invariant states that every edge of \tilde{M}_ℓ that is in the target graph, is an edge of $E(\mathcal{C}_y)$.

$\forall \{u, v\} \in \tilde{M}_\ell \cap E(\mathcal{T}^\ell)$, $\{u, v\} \in E(\mathcal{C}_y)$. (11)

Invariant (11) holds for w , because of invariant (5) maintained by Algorithm 1, together with the fact that if $e \in \tilde{M}_\ell \cap \mathcal{M}_y$, then $e \notin E(\mathcal{T}^\ell)$. Finally, the last invariant establishes that the witnesses of cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$ have their token available.

$\forall C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$, if $v \in V(C)$ is a witness of C , then $v \in T(\ell)$. (12)

Clearly invariant (12) holds for w , because of the definition of witnesses.

The algorithm selects one cycle $C \in \mathcal{C}_\ell \setminus \mathcal{C}_y$ at the time, and performs a move which involves at most one path-component of the target graph. Note that C intersects each component of \mathcal{T}^ℓ in at most two nodes. To see this, recall that each component K of \mathcal{T}^ℓ is either an \mathcal{M}_ℓ -alternating path, or an \mathcal{M}_ℓ -alternating cycle. Therefore, C can only intersect K at the endpoints of an \mathcal{M}_ℓ -alternating path. Since C has an odd number of nodes, it follows that there exists either one node v in $V(C)$ which is not a node of \mathcal{T}^ℓ , or one component K of \mathcal{T}^ℓ such that $|V(C) \cap V(K)| = 1$. Based on this observation, we have the following definition.

Definition 9. Let $C \in \mathcal{C}_\ell \setminus \mathcal{C}_y$. The least-intersecting component for C is a component H of the graph $(V(\mathcal{T}^\ell \cup C), E(\mathcal{T}^\ell))$ that minimizes $|V(H) \cap V(C)|$.

Note that the graph $(V(\mathcal{T}^\ell \cup C), E(\mathcal{T}^\ell))$ is simply the graph \mathcal{T}^ℓ with (possibly) additional singleton nodes that belong to $V(C)$. By the above reasoning, a least-intersecting component H for C will always have $|V(H) \cap V(C)| = 1$. In particular, either H will be a path-component of \mathcal{T}^ℓ with exactly one endpoint in $V(C)$, or H will be a singleton node v for some $v \in V(C)$ (this can happen if there exists a node $v \in V(C)$ with $v \notin V(\mathcal{T}^\ell)$). Algorithm 2 formally describes the moves we perform to go from w to r .

Algorithm 2 (from w to r): _____

1. Set $\ell := w$;
2. While there exists a cycle $C \in \mathcal{C}_\ell \setminus \mathcal{C}_y$ do:

2.1 Let H be the least-intersecting component for C , and $V(H) := \{v_1, \dots, v_k\}$ with $v_1 \in V(C)$;

2.2 If v_k is a single witness for some $\tilde{C} \in \mathcal{C}_y \setminus \mathcal{C}_\ell$, let N be the perfect matching of C exposing v_1 and do the following move:

(Δ) Change the coordinate ℓ_e of all the edges of $H \cup C \cup \tilde{C}$ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(H) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(H) \cap \mathcal{M}_\ell \\ 1 & \text{if } e \in N \\ 0 & \text{if } e \in E(C) \setminus N \\ \frac{1}{2} & \text{if } e \in E(\tilde{C}) \end{cases}$$

and let the node v_k and the cycle C use their tokens to pay for this move;

2.3 Else if $v_k \in V(\tilde{C})$ for some $\tilde{C} \in \mathcal{C}_\ell \setminus \mathcal{C}_y$, $\tilde{C} \neq C$, let N be the perfect matching of C exposing v_1 , let \tilde{N} be the perfect matching of \tilde{C} exposing v_k , and do the following move:

(∇) Change the coordinate ℓ_e of all the edges of $H \cup C \cup \tilde{C}$ to

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(H) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(H) \cap \mathcal{M}_\ell \\ 1 & \text{if } e \in N \\ 0 & \text{if } e \in E(C) \setminus N \\ 1 & \text{if } e \in \tilde{N} \\ 0 & \text{if } e \in E(\tilde{C}) \setminus \tilde{N} \end{cases}$$

and let the cycles C and \tilde{C} use their tokens to pay for this move;

2.4 Else, let j be any index such that $v_j \in V(H) \cap T(\ell)$ and v_j is not a witness in \mathcal{W} , let N be the perfect matching of C exposing v_1 , and do the move:

(\square) Change the coordinate ℓ_e of the edges of $H \cup C$ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(H) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(H) \cap \mathcal{M}_\ell \\ 1 & \text{if } e \in N \\ 0 & \text{if } e \in E(C) \setminus N \end{cases}$$

and let the node v_j and the cycle C use their tokens to pay for this move;

3. Output $r := \ell$.

Lemma 5. *All steps of Algorithm 2 can be correctly performed, and invariants (7), (8), (9), (10), (11) and (12) are maintained.*

Proof: First we argue that the moves are indeed valid moves. The move (Δ) is a valid move according to either Lemma 1(d) (if $v_k = v_1$) or Lemma 1(e) (if $v_k \neq v_1$).

The move (∇) is a valid move according to Lemma 1(e). Furthermore, the move (\square) is a valid move according to either Lemma 1(c) (if $v_k = v_1$), or Lemma 1(f) (if $v_k \neq v_1$).

Second, we argue that the invariants hold, by induction on the number of iterations of Step 2 performed by the algorithm. Let C be a cycle in $\mathcal{C}_\ell \setminus \mathcal{C}_y$ considered at a Step 2 of the algorithm, and H be the component in Step 2.1, with nodes $\{v_1, \dots, v_k\}$ (possibly, $v_k = v_1$). Let $\bar{\ell}$ be the vertex visited by the algorithm right after ℓ . Invariants (7), (10) and (11) follow by two things: (i) the nodes paying for the moves are in $V(H)$, but $V(H)$ and $E(H)$ will not be present in $\mathcal{T}^{\bar{\ell}}$, (ii) if K is a component of $\mathcal{T}^{\bar{\ell}}$ but not a component of \mathcal{T}^ℓ , then K contains at least one edge $e \in \mathcal{M}_{\bar{\ell}} \setminus \mathcal{M}_\ell$. Necessarily, the endpoints of e are in $V(\mathcal{C}_\ell) \setminus \{v_1, v_k\}$, i.e. $e \notin \tilde{\mathcal{M}}_\ell$, and therefore still have their token available.

Let K be a dangerous cycle-component of $\mathcal{T}^{\bar{\ell}}$. If K is also a dangerous cycle-component of \mathcal{T}^ℓ , then the condition of invariant (8) clearly still hold. Otherwise, K contains at least one edge $e = \{u, \bar{u}\}$ in $\mathcal{M}_{\bar{\ell}} \setminus \mathcal{M}_\ell$. By the same reasoning as above, u and \bar{u} are in $T(\bar{\ell})$, and since K is dangerous, they must be witnesses in \mathcal{W} . However, u and \bar{u} cannot be a pair of witnesses for the same cycle in \mathcal{C}_y , since otherwise this would contradict condition (ii) of Definition 5. It follows that invariant (8) holds. The argument for invariant (9) is identical to the one used in Lemma 4 (we omit it from this extended abstract). Invariant (12) holds trivially, since if we use a token of a witness node v_k to pay for moving from ℓ to $\bar{\ell}$, then v_k was a single witness and the cycle witnessed by v_k is in $\mathcal{C}_{\bar{\ell}} \cap \mathcal{C}_y$.

Finally, we argue that there are enough tokens to pay for the moves. In Step 2.2, the cycle C has a token by Observation 1, and v_k has a token by invariant (11). In Step 2.3, the cycles C and \tilde{C} have their tokens by Observation 1. In Step 2.4, the cycle C has a token by Observation 1, and v_j exists: invariant (7) guarantees that $V(H) \cap T(\ell)$ is not empty, and H cannot be dangerous, otherwise by invariant (9) we would perform the operation in Step 2.3 or 2.2. ■ We conclude this section with a lemma that lists some useful properties satisfied by vertex r (for its proof, we refer to the full version of the paper).

Lemma 6. *The vertex r satisfies the following properties:*

- (a) *Each component of the target graph \mathcal{T}^r contains at least two nodes that are in $T(r)$.*
- (b) *Let $C \in \mathcal{C}_y \setminus \mathcal{C}_r$ be a cycle with a single witness $(u, u) \in \mathcal{W}$. Then C is packed by r , and u is the endpoint of some path-component of \mathcal{T}^r .*
- (c) *Let $C \in \mathcal{C}_y \setminus \mathcal{C}_r$ be a cycle with a pair of witnesses $(u, v) \in \mathcal{W}$, $u \neq v$. If $u \in V(K)$ for some component K of \mathcal{T}^r , then $Q(u, v) \subseteq K$. Furthermore, u (resp. v) is incident into at least two edges of K .*
- (d) *Let $C \in \mathcal{C}_y \setminus \mathcal{C}_r$, and K be a dangerous cycle-component of \mathcal{T}^r such that $|V(C) \cap V(K)| \neq \emptyset$. Then K contains the witnesses of C .*

F. Moving from r to y

Given r from the previous procedure, we move to y using Algorithm 3. In Step 2, we perform moves involving path-components of the target graph, and in Step 3, we perform moves involving non-dangerous cycle-components of the target graph. In both cases, each component has (at least) two nodes with an available token to pay for the move. In particular, for a path-component, if any of these nodes is a single witness of some cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$, then the cycle will appear in the support of next vertex.

Dangerous cycle-components are considered last, in Step 4, and for each such component K , we perform two moves. Roughly speaking, a move which switches the coordinate values along the edges of $E(K)$ would require two tokens which are witnesses of some cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$, but differently from dangerous path-components, here such move will not make C appearing in the support of next vertex. Therefore, we first make a move which guarantees that two distinct cycles C and \bar{C} becomes packed, and then perform a second move which guarantees that both C and \bar{C} appear in the support of our vertex. We use the tokens of the two witnesses of C to pay for the first move, and the tokens of the two witnesses of \bar{C} to pay for the second move.

To describe our moves formally, we introduce one definition (see Figure 2 for an illustration).

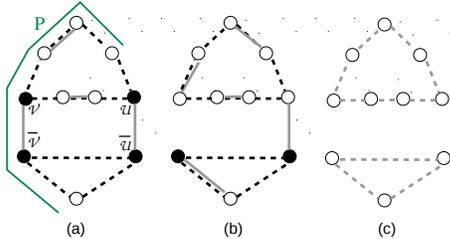


Figure 2. Two cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$ intersecting a dangerous component are shown in (a). Black edges represent the edges in the support of y , and gray edges represent the edges in the support of ℓ . Continuous lines represent edges of value 1, while dashed lines represent edges of value $\frac{1}{2}$. Black nodes represent nodes in $T(\ell)$. The dangerous component is the even cycle given by the edges $\{v, \bar{v}\}, \{u, \bar{u}\}$, the odd u - v path ($Q(u, v)$) contained in the top dashed black cycle, and the odd \bar{u} - \bar{v} path ($Q(\bar{u}, \bar{v})$) contained in the bottom dashed black cycle. The dark green line indicates the path P as in Definition 10. The figure in (b) shows how the coordinates of ℓ change after performing the move (\otimes) , while the figure in (c) shows how the coordinates of ℓ change after performing the move (\odot) .

Definition 10. Let K be a dangerous cycle-component of \mathcal{T}^ℓ , and let $C, \bar{C} \in \mathcal{C}_y \setminus \mathcal{C}_\ell$ be two distinct cycles with $V(C) \cap V(K) \neq \emptyset$ and $V(\bar{C}) \cap V(K) \neq \emptyset$. Let (u, v) (resp. (\bar{u}, \bar{v})) be the pair of witnesses of C (resp. \bar{C}). We call a path P a K -linking path between C and \bar{C} if:

- (i) P is an \mathcal{M}_ℓ -augmenting path;
- (ii) $E(P) \cap E(K) = \{v, \bar{v}\}$;
- (iii) for all $e' \in E(P)$ with $e' \neq \{v, \bar{v}\}$, we have $e' \in E(C \setminus Q(u, v)) \cup E(\bar{C} \setminus Q(\bar{u}, \bar{v}))$.

In Step 5, we perform our last set of moves, which make all remaining cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$ appear, one by one, in the support of next vertex. For any vertex ℓ visited by Algorithm 3, we will maintain invariant (12) and the following invariant:

If K is a component of \mathcal{T}^ℓ , then K is a component of \mathcal{T}^r and $V(K) \cap T(\ell) = V(K) \cap T(r)$ (13)

Algorithm 3 (from r to y): _____

1. Set $\ell := r$;
2. While there exists a path-component K of \mathcal{T}^ℓ :
 - 2.1 Let $V(K) := \{v_1, \dots, v_k\}$, let $E(K) := \{\{v_1, v_2\}, \dots, \{v_{k-1}, v_k\}\}$ and let v_i (resp. v_j) be the node in $V(K) \cap T(\ell)$ with the smallest (resp. biggest) index;
 - 2.2 Do the following move:
 - (\oplus) Change the coordinates of ℓ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(K) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(K) \cap \mathcal{M}_\ell \\ \frac{1}{2} & \text{if } v_i \text{ is a single witness for some} \\ & C \in \mathcal{C}_y \setminus \mathcal{C}_\ell \text{ and } e \in E(C) \\ \frac{1}{2} & \text{if } v_j \text{ is a single witness for some} \\ & \bar{C} \in \mathcal{C}_y \setminus \mathcal{C}_\ell \text{ and } e \in E(\bar{C}) \end{cases}$$

and let v_i and v_j use their tokens to pay for this move;

3. While there exists a non-dangerous cycle-component K of \mathcal{T}^ℓ :

- 3.1 Let v and u be two nodes in $V(K) \cap T(\ell)$, such that v, u are not witnesses in \mathcal{W} , and do the move:

(\ominus) Change the coordinate ℓ_e of the edges of K to

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(K) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(K) \cap \mathcal{M}_\ell \end{cases}$$

and let v and u use their tokens to pay for this move;

4. While there exists a dangerous cycle-component K of \mathcal{T}^ℓ :

- 4.1 Let P be a K -linking path between two distinct cycles C and \bar{C} in $\mathcal{C}_y \setminus \mathcal{C}_\ell$. Let (u, v) be the pair of witnesses for C and let (\bar{u}, \bar{v}) be the pair of witnesses for \bar{C} ;

- 4.2 Do the following moves:

(\otimes) Change the coordinate ℓ_e of the edges of P to

$$\ell_e = \begin{cases} 1 & \text{if } e \in E(P) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in E(P) \cap \mathcal{M}_\ell \end{cases}$$

and let u and v use their tokens to pay for this move;

(\odot) Change the coordinate ℓ_e of all the edges in $E(C) \cup E(\bar{C}) \cup (E(K) \setminus \{v, \bar{v}\})$ by setting

$$\ell_e = \begin{cases} 1 & \text{if } e \in (E(K) \setminus (E(P \cup C \cup \bar{C}))) \setminus \mathcal{M}_\ell \\ 0 & \text{if } e \in (E(K) \setminus (E(C \cup \bar{C}))) \cap \mathcal{M}_\ell \\ \frac{1}{2} & \text{if } e \in E(C) \cup E(\bar{C}) \end{cases}$$

and let \bar{u} and \bar{v} use their tokens to pay for this move;

5. While there exists a cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$, do the move:

(\odot) Change the coordinate ℓ_e of the edges in $E(C)$ to $\frac{1}{2}$

and let the witnesses of C use their tokens to pay for this move;

6. Output ℓ .

Lemma 7. *All steps of Algorithm 3 can be correctly performed, invariants (12) and (13) are maintained, and the vertex output in Step 6 satisfies $\ell = y$.*

Proof: We argue that each step of Algorithm 3 can be correctly performed and that the claimed invariants are maintained, by induction on the number of moves performed by the algorithm.

In Step 2, we consider a path-component K of \mathcal{T}^ℓ . Invariant (13), which holds by the inductive hypothesis, guarantees that K is a component of \mathcal{T}^r , and that if any of the endpoints is a single witness of a cycle, then this cycle is packed by ℓ . Therefore, the move (\oplus) is a valid move, according to either Lemma 1(b), or Lemma 1(e), or Lemma 1(f). Furthermore, since $V(K) \cap T(\ell) = V(K) \cap T(r)$, property (a) of Lemma 6 guarantees that there are at least two nodes in $V(K)$ with an available token which can pay for the move.

We now show that the invariants are maintained in this step. Let v_i be the node chosen in Step 2.1. We claim that if v_i is a witness, then it can only be a single witness. Assume for a contradiction that v_i is in a pair of witnesses $(v_i, v_{i'}) \in \mathcal{W}$, with $i \neq i'$ for some cycle $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$. Note that, $Q(v_i, v_{i'}) \subseteq K$ by property (c) of Lemma 6. Since $v_{i'} \in T(\ell)$ by invariant (12), it follows $i' > i$. Then, by definition of target matching, $\{v_i, v_{i+1}\} \in M_C$, and since v_i has degree 2 in K by property (c) of Lemma 6, it follows that $\{v_{i-1}, v_i\} \in \mathcal{M}_\ell$. However, this also implies that $v_{i-1} \in T(\ell)$ (using invariant (13) together with (10) which holds for r). We get a contradiction with our choice of i . We can apply a similar argument to v_j . This yields that invariant (12) is maintained. Finally, after the move is performed, the component K is no longer a component of the target graph, while all other components of the target graph remain the same. This shows that invariant (13) is maintained.

In Step 3, we consider a non dangerous cycle-component K of \mathcal{T}^ℓ . The move (\ominus) is a valid move, according to Lemma 1(a). Since K is not dangerous, there is at least one node in $V(K) \cap T(\ell) = V(K) \cap T(r)$ that is not the witness of any cycle. However, using invariant (10), which holds for r , we know that $|V(K) \cap T(\ell)|$ is even. Furthermore, since K is a cycle, K also contains an even number of nodes that are witnesses in \mathcal{W} . It follows that K contains at least two nodes in $V(K) \cap T(\ell)$ that are not witnesses of any cycle, and therefore can pay for the move. This shows that invariant

(12) is maintained. Invariant (13) is also maintained, for the exact same argument used in the previous step.

In Step 4, we consider a dangerous cycle-component K of \mathcal{T}^ℓ . Let us argue that Step 4.1 can be performed, i.e. that there exists a K -linking path P . Let $e = \{v, \bar{v}\}$ be any edge in \mathcal{M}_ℓ with both endpoints in $T(\ell)$. Such an edge exists, because invariant (13), which holds by the inductive hypothesis, guarantees that K is a component of \mathcal{T}^r and $V(K) \cap T(\ell) = V(K) \cap T(r)$, therefore we can rely on property (a) of Lemma 6, and on invariant (10) which holds for r . Since the component is dangerous, v and \bar{v} are both witness nodes. Let $C \in \mathcal{C}_y \setminus \mathcal{C}_\ell$ be the cycle which has v as one of its witnesses. Note that \bar{v} cannot be a witness for the same cycle C : if this was the case, then either (i) $\{v, \bar{v}\} = Q(v, \bar{v})$, implying that $\{v, \bar{v}\} \subseteq M_C$, contradicting that the edge is in the target graph, or (ii) $K = Q(v, \bar{v}) \cup \{v, \bar{v}\}$, contradicting that K contains the witnesses of at least two distinct cycles. Let \bar{C} be the cycle which has \bar{v} as one of its witnesses. Neither v or \bar{v} can be a single witness, since such nodes appear only in path-components, by property (b) of Lemma 6. Let u and \bar{u} be such that $(u, v) \in \mathcal{W}$, and $(\bar{u}, \bar{v}) \in \mathcal{W}$. By invariant (13), we know that all the components of \mathcal{T}^ℓ are dangerous cycle-components. By property (d) of Lemma 6, $V(C) \cap V(K') = \emptyset$ for all components $K' \neq K$ in \mathcal{T}^ℓ . It follows that $(M_C \Delta \mathcal{M}_\ell) \cap E(C) = Q(u, v)$, where M_C is the target matching of C . This implies that there exists a \mathcal{M}_ℓ -alternating path P_1 from v to the unique M_C -exposed node in $V(C)$, which does not use edges in $Q(u, v)$ (refer again to Figure 2). The same argument shows that there exists a \mathcal{M}_ℓ -alternating path P_2 from \bar{v} to the unique $M_{\bar{C}}$ -exposed node in $V(\bar{C})$, which does not use edges in $Q(\bar{u}, \bar{v})$. Combining P_1 , $\{v, \bar{v}\}$, and P_2 , yields a K -linking path P . The move (\otimes) is then a valid move, according to Lemma 1(b), and u and v have their token available by invariant (12).

After the move (\otimes) is performed, C becomes packed, with u being its unique $(\mathcal{M}_\ell \cap E(C))$ -exposed node, and \bar{C} becomes packed, with \bar{u} being its unique $(\mathcal{M}_\ell \cap E(\bar{C}))$ -exposed node. The move (\odot) is then a valid move, according to Lemma 1(f). The nodes \bar{u} and \bar{v} have their token available by invariant (12). Since both C and \bar{C} are now in the support of the current vertex, invariant (12) is maintained. Invariant (13) is also maintained, for the exact same argument used in the previous steps.

Let ℓ be the vertex after Step 4 terminates. Note that all cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$ are packed by \mathcal{M}_ℓ , and $\mathcal{M}_\ell \cap \mathcal{M}_y = \mathcal{M}_y$. All cycles in $\mathcal{C}_y \setminus \mathcal{C}_w$ which had a unique witness in \mathcal{W} , have appeared in the support of the current vertex during some iteration of Step 2, because of invariant (13) together with property (b) of Lemma 6. Therefore, all cycles in $\mathcal{C}_y \setminus \mathcal{C}_\ell$ have two distinct witnesses, with their token available at the beginning of Step 5 because of invariant (12). It follows that every move (\odot), which is a valid move by Lemma 1(c), can be paid. It is easy to see then that the output vertex is y . ■

G. Lower bound

We here argue that the quantity $\max_{x \in \text{vert}(\mathcal{P}_{FM})} \{ \mathbf{1}^T x + \frac{|C_x|}{2} \}$ is a lower bound on the value of the diameter of \mathcal{P}_{FM} . Let w be the vertex at which the above maximum value is achieved. We will show that the distance between the vertex w and the $\mathbf{0}$ -vertex (i.e. the vertex corresponding to an empty matching), is at least $\mathbf{1}^T w + \frac{|C_w|}{2}$.

Suppose to introduce a non-negative slack variable for each inequality of the form $x(\delta(v)) \leq 1$ of \mathcal{P}_{FM} . We get a polytope that naturally corresponds to the set of feasible solutions of a fractional *perfect* matching problem on a modified graph \bar{G} , defined as follows. Let $\bar{G} = (V, \bar{E})$ be the graph obtained from G by adding a loop edge on each node $v \in V$. We let $\bar{E} := E \cup L$, with L being the set of loop edges introduced. We can interpret the slack variable associated to a node v as the variable associated to its loop edge e_v , and define $\delta_{\bar{G}}(v) := \delta(v) \cup e_v$ (note that the loop edge is counted once in this set). We define

$$\bar{\mathcal{P}}_{FM} := \{ \bar{x} \in \mathbb{R}^{\bar{E}} : \sum_{e \in \delta_{\bar{G}}(v)} \bar{x}_e = 1, \forall v \in V, \bar{x} \geq 0 \}$$

For a vertex x of \mathcal{P}_{FM} we let \bar{x} denote the correspondent vertex of $\bar{\mathcal{P}}_{FM}$, and vice versa. Note that two vertices x and y are adjacent vertices of \mathcal{P}_{FM} if and only if \bar{x} and \bar{y} are adjacent vertices of $\bar{\mathcal{P}}_{FM}$.

Let \bar{x} be the vertex of $\bar{\mathcal{P}}_{FM}$ corresponding to $x = \mathbf{0}$, i.e. the empty matching in G , and \bar{w} be the vertex of $\bar{\mathcal{P}}_{FM}$ corresponding to w . The support graph of \bar{x} , denoted by $\bar{\mathcal{G}}_{\bar{x}}$ contains $|V|$ odd cycles, all of unit length, given by the $|V|$ loop edges. The support graph of \bar{w} , denoted by $\bar{\mathcal{G}}_{\bar{w}}$ contains $|C_w| + |V \setminus V(\mathcal{G}_w)|$ odd cycles: there are $|C_w|$ odd cycles of length at least 3, and $|V \setminus V(\mathcal{G}_w)|$ odd cycles of unit length, given by loop edges associated to nodes that are not in the support graph \mathcal{G}_w of w .

The following claim will be used to give a lower bound on the distance of two vertices of $\bar{\mathcal{P}}_{FM}$, which depends on the number of odd cycles that are not in common in the support graphs of the vertices.

Claim. Let \bar{y} and \bar{z} be two adjacent vertices of $\bar{\mathcal{P}}_{FM}$. Let $\bar{C}_{\bar{y}}$ be the set of odd cycles in the support graph of \bar{y} , and $\bar{C}_{\bar{z}}$ be the set of odd cycles in the support graph of \bar{z} . Then $|\bar{C}_{\bar{y}} \Delta \bar{C}_{\bar{z}}| \leq 2$.

Proof of claim. A proof of the claim can be derived from the results given in [4]. We report the details for completeness.

First, we claim that there is at most one component of $\bar{\mathcal{G}}_{\bar{y}} \cup \bar{\mathcal{G}}_{\bar{z}}$ that contains an edge f with $\bar{y}_f \neq \bar{z}_f$. Suppose for a contradiction that there exist two such components, namely K_1 and K_2 . Let \tilde{z} be the point defined as $\tilde{z}_e = \bar{z}_e$ for all $e \in \bar{E} \setminus E(K_1)$, and $\tilde{z}_e = \bar{y}_e$ for all $e \in E(K_1)$. Similarly, let \tilde{y} be the point defined as $\tilde{y}_e = \bar{z}_e$ for all $e \in \bar{E} \setminus E(K_2)$, and $\tilde{y}_e = \bar{y}_e$ for all $e \in E(K_2)$. It is not difficult to see that \tilde{z} and \tilde{y} are both vertices of $\bar{\mathcal{P}}_{FM}$. However, $\frac{1}{2}\tilde{z} + \frac{1}{2}\tilde{y} = \frac{1}{2}\bar{z} + \frac{1}{2}\bar{y}$. This is a contradiction, since if two vertices of a polytope

are adjacent, there is a unique way to express their midpoint as a convex combination of vertices.

Let K be the component of $\bar{\mathcal{G}}_{\bar{y}} \cup \bar{\mathcal{G}}_{\bar{z}}$ that contains an edge f with $\bar{y}_f \neq \bar{z}_f$, and let k be the number of nodes of this component. If K has at most $k+1$ edges, then K can be seen to be a tree spanning its k nodes, plus two additional (possibly loop) edges: it is easy to realize then that K can have at most 2 odd cycles. We are left to show that K contains at most $k+1$ edges.

Let $\bar{G}[V(K)]$ be the subgraph of \bar{G} induced by the nodes in $V(K)$. Consider the fractional perfect matching polytope associated to the graph $\bar{G}[V(K)]$. Now one can see that $\bar{z}|_{E(\bar{G}[V(K)])} \in \mathbb{R}^{E(\bar{G}[V(K)])}$, obtained from the vector \bar{z} by taking only the coordinates in $E(\bar{G}[V(K)])$, is a vertex of this polytope, and the same holds for $\bar{y}|_{E(\bar{G}[V(K)])}$ (defined similarly). Furthermore, these vertices must be adjacent if \bar{z} and \bar{y} are adjacent.

Let A be the incidence matrix of the graph $\bar{G}[V(K)]$, where we have a row for every node of $V(K)$, and a column for every edge of $\bar{G}[V(K)]$. Note that the column associated to a loop edge has only one non-zero entry. Clearly, the rank of A is $|V(K)| = k$. Furthermore, the constraint matrix of the polytope $\bar{\mathcal{P}}_{FM}$ for the graph $\bar{G}[V(K)]$ is $\begin{pmatrix} A \\ I \end{pmatrix}$, with I being the identity matrix of order $E(\bar{G}[V(K)])$. Let M be the submatrix of $\begin{pmatrix} A \\ I \end{pmatrix}$ corresponding to the constraints

that are tight for both $\bar{y}|_{E(\bar{G}[V(K)])}$ and $\bar{z}|_{E(\bar{G}[V(K)])}$. Since $\bar{y}|_{E(\bar{G}[V(K)])}$ and $\bar{z}|_{E(\bar{G}[V(K)])}$ are adjacent vertices, the rank of M is equal to $|E(\bar{G}[V(K)])| - 1$. Note that M contains A as a submatrix, and since the rank of A is k , it follows that at least $|E(\bar{G}[V(K)])| - 1 - k$ variables are zero in both $\bar{z}|_{E(\bar{G}[V(K)])}$ and $\bar{y}|_{E(\bar{G}[V(K)])}$, and therefore the union of their support graphs contains at most $k+1$ edges. \square

Now let us discuss how the claim implies our desired lower bound. By the above claim, if \bar{y} and \bar{z} are two (non necessarily adjacent) vertices of $\bar{\mathcal{P}}_{FM}$, then the quantity $\frac{|C_{\bar{y}} \Delta C_{\bar{z}}|}{2}$ is a lower bound on the distance between \bar{y} and \bar{z} on the 1-skeleton of $\bar{\mathcal{P}}_{FM}$, since the size of the symmetric difference of the sets of odd cycles can be reduced by at most 2 at each move.

We can use this to bound the number of moves needed on a path from x to w . The distance between x and w on the 1-skeleton of \mathcal{P}_{FM} is equal to the distance between \bar{x} and \bar{w} on the 1-skeleton of $\bar{\mathcal{P}}_{FM}$. The distance between \bar{x} and \bar{w} is at least the cardinality of the symmetric difference of the odd cycles in their support graphs divided by 2, i.e., $(|C_w| + |V(\mathcal{G}_w)|)/2$. Note that $\frac{|V(\mathcal{G}_w)|}{2} = \sum_{v \in V(\mathcal{G}_w)} \frac{1}{2} = \mathbf{1}^T w$. It follows that the distance between x and w is at least $\mathbf{1}^T w + \frac{|C_w|}{2}$, as desired.

IV. HARDNESS OF COMPUTING THE DIAMETER OF \mathcal{P}_{FM}

With Theorem 1 at hands, one can easily prove that computing the diameter of a polytope is strongly NP-hard.

Proof: (Proof of Theorem 2) We give an easy reduction from the *Partition Into Triangles* (PIT) problem. In an instance the PIT problem, we are given a simple graph $G = (V, E)$ with $|V| = 3q$ for some integer $q > 0$. We want to decide whether there exists a partition of V into q sets V_1, \dots, V_q , each containing 3 nodes, with the property that the subgraph of G induced by every set V_i is a triangle. This problem is strongly NP-hard [10].

Let $G = (V, E)$ be an instance of PIT, and \mathcal{P}_{FM} the fractional matching polytope associated to G . We claim that the $\text{diam}(\mathcal{P}_{FM}) = \frac{2}{3}|V|$ if and only if there is a yes-answer to the PIT instance described by G , i.e. the edges of G allow for a partition of V into triangles.

Suppose that the diameter of \mathcal{P}_{FM} is $\frac{2}{3}|V|$. From Theorem 1, it follows that there exists a vertex x of \mathcal{P}_{FM} such that $\text{diam}(\mathcal{P}_{FM}) = \mathbf{1}^T x + \frac{|C_x|}{2}$. Note that for any $y \in \text{vert}(\mathcal{P}_{FM})$, we have $\mathbf{1}^T y \leq \frac{|V|}{2}$ and $|C_y| \leq \frac{|V|}{3}$. Therefore, $\mathbf{1}^T x + \frac{|C_x|}{2} = \frac{2}{3}|V|$ implies that $\mathbf{1}^T x = |V|/2$ and $|C_x| = |V|/3$. Since C_x is a set of node-disjoint odd cycles of a simple graph, it follows that C_x contains a set of $|V|/3 = q$ triangles.

Suppose G is a yes-instance, and let $T \subseteq E$ be the set of edges of the triangles induced by the sets V_i , for $i = 1, \dots, q$. We construct a fractional matching x by letting $x_e = \frac{1}{2}$ for all $e \in T$, and $x_e = 0$ otherwise. By construction, $\mathbf{1}^T x = \frac{|V|}{2}$ and $|C_x| = \frac{|V|}{3}$, and x is indeed a vertex of \mathcal{P}_{FM} according to Theorem 4. As we mentioned before, for any $y \in \text{vert}(\mathcal{P}_{FM})$, we have $\mathbf{1}^T y \leq \frac{|V|}{2}$ and $|C_y| \leq \frac{|V|}{3}$. Therefore, by Theorem 1, it follows that $\text{diam}(\mathcal{P}_{FM}) = \mathbf{1}^T x + \frac{|C_x|}{2} = \frac{2}{3}|V|$. ■

REFERENCES

- [1] M. L. Balinski, *On maximum matching, minimum covering and their connections*, Proceedings of the Princeton symposium on mathematical programming (1970).
- [2] M. L. Balinski, *The Hirsch conjecture for dual transportation polyhedra*, Mathematics of Operations Research, pages 629-633, (1984).
- [3] M. L. Balinski, A. Russakoff, *On the assignment polytope*, SIAM Review, 16(4):516-525, (1974).
- [4] R. E. Behrend, *Fractional perfect b-matching polytopes. I: General theory*, Linear Algebra and its Applications, 439(12):3822-3858, (2013).
- [5] S. Borgwardt, E. Finhold, R. Hemmecke, *Quadratic diameter bounds for dual network flow polyhedra*, Mathematical Programming, Volume 159, Issue 1-2, pages 237-251, (2016).
- [6] S. Borgwardt, J. A. De Loera, E. Finhold, *The diameters of network-flow polytopes satisfy the Hirsch conjecture*, Mathematical Programming, (2017).
- [7] G. Brightwell, J. Heuvel, L. Stougie, *A linear bound on the diameter of the transportation polytope*, Combinatorica, pages 133-139, (2006).
- [8] V. Chvátal, *On certain polytopes associated with graphs*, Journal of Combinatorial Theory, Series B, 18(2):138-154, (1975).
- [9] A. M. Frieze, S. H. Teng, *On the complexity of computing the diameter of a polytope*, Computational Complexity, 4: 207-219, (1994).
- [10] M. R. Garey, D. S. Johnson, *Computers and Intractability, A Guide to the Theory of NP-Completeness*, W.H. Freeman & Co., New York, (1979).
- [11] J. Edmonds, *Maximum matching and a polyhedron with 0,1-vertices*, Journal of Research of the National Bureau of Standards: Section B Mathematics and Mathematical Physics, 69B(1-2):125-130, (1965).
- [12] P. Eirinakis, D. Magos, I. Mourtos, *From One Stable Marriage to the Next: How Long Is the Way?*, SIAM Journal on Discrete Mathematics, 28(4):1971-1979, (2014).
- [13] J. Hurkens, *On the diameter of the edge cover polytope*, Journal of Combinatorial Theory, Series B Volume 51, Issue 2, pages 271-276, (1991).
- [14] V. Kaibel, M. E. Pfetsch, *Some Algorithmic Problems in Polytope Theory*, Algebra, Geometry and Software Systems, pages 23-47, (2003).
- [15] V. Klee, D. W. Walkup, *The d-step conjecture for polyhedra of dimension $d > 6$* , Acta Mathematica, (133):53-78, (1967).
- [16] D. Naddef, *The Hirsch conjecture is true for (0,1)-polytopes*, Mathematical Programming, 45:109-110, (1989).
- [17] C. Michini, A. Sassano, *The Hirsch conjecture for the fractional stable set polytope*, Mathematical Programming, 147(1):309-330, (2014).
- [18] M. Padberg, M. R. Rao, *The travelling salesman problem and a class of polyhedra of diameter two*, Mathematical Programming, pages 32-45, (1974).
- [19] F. J. Rispoli, S. Cosares, *A bound of 4 for the diameter of the symmetric traveling salesman polytope*, SIAM Journal on Discrete Mathematics, pages 373-380, (1998).
- [20] N. Sukegawa, *Improving bounds on the diameter of a polyhedron in high dimensions*, Discrete Mathematics, Volume 340, Issue 9, pages 2134-2142, (2017).
- [21] F. Santos, *A counterexample to the Hirsch conjecture*, Annals of Mathematics, 176(1):383-412, (2012).
- [22] F. Santos, *Recent progress on the combinatorial diameter of polytopes and simplicial complexes*, Top 21, pages 426-460, (2013).
- [23] S. Smale, *Mathematical problems for the next century*, Mathematics: frontiers and perspectives, American Mathematics Society, Providence, RI, pages 271-294, (2000).