

Shortest odd paths in undirected graphs with conservative weight functions

Alpár Jüttner^{a,b}, Csaba Király^{a,b}, Lydia Mirabel Mendoza-Cadena^{b,c,*},
Gyula Pap^{a,b}, Ildikó Schlotter^{d,e}, Yutaro Yamaguchi^f

^a*ELKH-ELTE Egerváry Research Group, Eötvös Loránd Research Network
(ELKH), Budapest, Hungary*

^b*Department of Operations Research, ELTE Eötvös Loránd
University, Budapest, Hungary*

^c*MTA-ELTE Matroid Optimization Research Group, Budapest, Hungary*

^d*HUN-REN Centre for Economic and Regional Studies, Budapest, Hungary*

^e*Budapest University of Technology and Economics, Budapest, Hungary*

^f*Department of Information and Physical Sciences, Graduate School of Information
Science and Technology, Osaka University, Osaka, Japan*

Abstract

We consider the SHORTEST ODD PATH problem, where given an undirected graph G , a weight function on its edges, and two vertices s and t in G , the aim is to find an (s, t) -path with odd length and, among all such paths, of minimum weight. For the case when the weight function is conservative, i.e., when every cycle has non-negative total weight, the complexity of the SHORTEST ODD PATH problem had been open for 20 years, and was recently shown to be NP-hard.

We give a polynomial-time algorithm for the special case when the weight function is conservative and the set E^- of negative-weight edges forms a single tree. Our algorithm exploits the strong connection between SHORTEST ODD PATH and the problem of finding two internally vertex-disjoint paths between two terminals in an undirected edge-weighted graph. It also relies on solving an intermediary problem variant called SHORTEST PARITY-CONSTRAINED ODD PATH where for certain edges we have parity constraints on their position along the path.

Also, we exhibit two FPT algorithms for solving SHORTEST ODD PATH. The first FPT algorithm is parameterized by $|E^-|$, the number of negative edges, or more generally, by the maximum size of a matching in the subgraph of G spanned by E^- , when the weight function is conservative. Our second FPT algorithm is parameterized by the treewidth of G , and the algorithm does not rely on conservativeness.

Keywords: shortest odd path, parity constrained odd path, fixed-parameter algorithms, treewidth, monadic second-order logic

1. Introduction

Lovász [24] asked for the complexity of the following question nearly 20 years ago: Given two vertices s and t in an undirected graph G and a weight function w on the edges of G such that each cycle in G has non-negative weight, find a minimum-weight odd path between s and t . Throughout the paper, a path is *odd* (or even) if it contains an odd (even, respectively) number of edges.

Recently, the problem was shown to be NP-hard by Schlotter and Sebő [23], even in the case when the weight function is conservative and takes values only from $\{-1, 1\}$. Note that replacing “odd” with “even” does not change the computational complexity of the problem: finding a minimum-weight odd (or even) (s, t) -path (i.e., a path between s and t) in G can easily be reduced to finding a minimum-weight even (or odd, respectively) path between s and a newly added vertex t' in the graph obtained from G by adding an edge tt' with weight 0. Note also that finding a shortest odd (even) path between two vertices s and t generalizes the shortest path problem: adding new vertices t_1 and t_2 along with 0-weight edges tt_2 , tt_1 , and t_1t_2 , a shortest odd (or even, resp.) (s, t_2) -path in the resulting graph yields a shortest (s, t) -path in the original graph. For undirected graphs with possibly negative but *conservative* edge-weights (meaning that there are no cycles with negative total weight), the shortest path problem is solvable in polynomial time [24, Chapter 29.2]; however, the usual shortest path algorithms for directed graphs are not applicable, since bidirecting an edge with negative weight creates a negative cycle.

If the weights are non-negative, a minimum-weight odd path between fixed vertices can be found via computing a maximum-weight matching problem in an auxiliary graph. Thomassen attributes this algorithm to Edmonds [25], while Grötschel and Pulleybank refer to it as “Waterloo folklore” [13]. A Dijkstra-like algorithm was given for the problem by Derigs [11].

*Corresponding author

Email addresses: `alpar@cs.elte.hu` (Alpár Jüttner), `csaba.kiraly@ttk.elte.hu` (Csaba Király), `lyd21@student.elte.hu` (Lydia Mirabel Mendoza-Cadena), `gyula.pap@ttk.elte.hu` (Gyula Pap), `schlotter.ildiko@krtk.hun-ren.hu` (Ildikó Schlotter), `yutaro.yamaguchi@ist.osaka-u.ac.jp` (Yutaro Yamaguchi)

For the unweighted case, Lapaugh and Papadimitriou [16] gave a linear-time algorithm that computes an odd (or even) path between two vertices, whenever such a path exists; Arkin et al. [1] achieved linear-time solvability even for finding a shortest odd (or even) path between two vertices.

The generalized problem of finding a shortest non-zero path in group-labeled graphs with non-negative edge weights has also been studied. A recent work of Iwata and Yamaguchi [14] shows a deterministic and strongly polynomial-time algorithm for this problem.

For directed graphs, the problem becomes much harder: already deciding whether a path of prescribed parity exists between two given vertices is NP-hard, as proved independently by Thomassen [25] and by Lapaugh and Papadimitriou [16]. Building on the connection between even-length cycles in a directed graph and Pfaffian orientations of bipartite graphs [26], Robertson, Seymour and Thomas [20], and independently, McCuaig [18] gave a polynomial-time algorithm for deciding whether a directed graph contains a cycle of even length. Using advanced algebraic techniques, Björklund, Husfeldt and Kaski [3] gave a randomized polynomial-time algorithm for finding a shortest even cycle in a directed, unweighted graph.

The closely related problem of finding parity-constrained induced cycles in undirected graphs is also of interest. Conforti, Cornuéjols, Kapoor, and Vušković [8] proved that detecting an even hole (where a *hole* is an induced cycle of length at least four) can be done in polynomial time; subsequently, Chudnovsky, Kawarabayashi, and Seymour [6] gave a simpler algorithm with a slightly better running time for the problem. Chudnovsky, Scott, Seymour, and Spirkl [7] recently showed that the problem of detecting if a graph contains an odd hole can also be solved in polynomial time. These algorithms inspired a wave of further research that aimed to improve running times and to solve more general problems in this area.

Our results. We consider the SHORTEST ODD PATH problem defined as follows:

SHORTEST ODD PATH:

Input: An undirected graph $G = (V, E)$, a weight function $w: E \rightarrow \mathbb{R}$, and two vertices s and t in G .

Goal: Find a minimum-weight odd (s, t) -path.

We solve the SHORTEST ODD PATH problem in polynomial time when the weight function is conservative and the set E^- of negative edges forms a tree in G . We also present two FPT algorithms for solving SHORTEST ODD

PATH on graphs with conservative weights. In the first FPT algorithm, our parameter is $|E^-|$, the number of negative edges; in fact, the approach we use can be extended to yield an FPT algorithm where the parameter is the size of a maximum matching in $G[E^-]$, the subgraph of G spanned by E^- . In the second FPT algorithm, the parameter is the treewidth of G ; in fact, this algorithm does not rely on conservativeness.

Our approach relies on solving a variant of the problem that we call SHORTEST PARITY-CONSTRAINED ODD PATH. In order to define this problem, we give the following definition. Let $F_{\text{even}} \subseteq E$ and $F_{\text{odd}} \subseteq E$ be two disjoint edge sets in G . An odd path Q is $(F_{\text{even}}, F_{\text{odd}})$ -constrained, if the sequence number of each edge in $Q \cap F_{\text{even}}$ and in $Q \cap F_{\text{odd}}$ is even and odd, respectively. Here, the *sequence number* of an edge e in Q is i , if e is the i -th edge on Q ; note that since Q is odd, e has an odd (even) sequence number independently whether we count the edges of Q starting from s or from t .

Using this terminology, we consider the following problem.

SHORTEST PARITY-CONSTRAINED ODD PATH:

Input: An undirected graph $G = (V, E)$, a weight function $w: E \rightarrow \mathbb{R}$, two vertices s and t in G , and disjoint edge sets $F_{\text{even}} \subseteq E$ and $F_{\text{odd}} \subseteq E$.

Goal: Find a minimum-weight $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path.

We show that this problem can be solved in polynomial time if all negative-weight edges are contained in $F_{\text{even}} \cup F_{\text{odd}}$; this result will be useful for us when developing our algorithms for SHORTEST ODD PATH.

Organization. Basic definitions and notation are introduced in Section 2. Section 3 presents some initial observations and key definitions that will be useful in our algorithms. In Section 4, we present a polynomial-time algorithm for the SHORTEST PARITY-CONSTRAINED ODD PATH problem. In Section 5, we show how to solve the SHORTEST ODD PATH problem when the weight function is conservative and the negative edges form a tree. Finally, in Section 6 we develop FPT algorithms for the SHORTEST ODD PATH problem.

2. Preliminaries

Basic notation. We denote the set of real numbers by \mathbb{R} , the rational numbers by \mathbb{Q} , and the non-negative real numbers by $\mathbb{R}_{\geq 0}$. Given subsets

$X, Y \subseteq E$, the *symmetric difference* of X and Y is denoted by $X \Delta Y := (X \setminus Y) \cup (Y \setminus X)$.

Graphs. Let a graph G be a pair (V, E) where V and E are the sets of vertices and edges, respectively. For two vertices u and v in V , an edge connecting u and v is denoted by uv or vu . For $X \subseteq V$ and $F \subseteq E$, the graph obtained by deleting X or F is denoted by $G - X$ or $G - F$, respectively. $G[F]$ denotes the graph *spanned* by F , that is, the subgraph of G containing only the edges of F and all the vertices incident to them. For an edge set $F \subseteq E$ and a vertex $v \in V$, let $d_F(v)$ denote the number of edges in F incident to v .

A *walk* W in G is a series e_1, e_2, \dots, e_ℓ of edges in G for which there exist vertices v_0, v_1, \dots, v_ℓ in G such that $e_i = v_{i-1}v_i$ for each $i \in \{1, 2, \dots, \ell\}$; note that both vertices and edges may appear repeatedly on a walk, and in this sense we think of the sets of vertices and edges of a walk as multisets in which an element might appear more than once. In contrast, we denote by $V(W)$ and $E(W)$ the sets of vertices and edges *contained by* W , respectively, where each item is taken with multiplicity one. We define the *length* of W as $|W| = \ell$, and we say that W is *odd* (or *even*) if ℓ is odd (or even, respectively). The *endpoints* of W are v_0 and v_ℓ , or in other words, it is a (v_0, v_ℓ) -walk, while all vertices on W that are not endpoints are *inner vertices*. If $v_0 = v_\ell$, then we say that W is a *closed walk*.

A *path* is a walk on which no vertex appears more than once. We usually treat a path as a set $\{e_1, e_2, \dots, e_\ell\}$ of edges for which there exist distinct vertices v_0, v_1, \dots, v_ℓ in G such that $e_i = v_{i-1}v_i$ for each $i \in \{1, 2, \dots, \ell\}$. For any i and j with $0 \leq i \leq j \leq \ell$, we write $P[v_i, v_j]$ for the *subpath* of P between v_i and v_j , consisting of edges e_{i+1}, \dots, e_j . Note that since we associate no direction with P , we have $P[v_i, v_j] = P[v_j, v_i]$. For indices i_1, i_2, \dots, i_r , we say that the vertices $v_{i_1}, v_{i_2}, \dots, v_{i_r}$ *follow each other in this order* if $i_1 \leq i_2 \leq \dots \leq i_r$ or $i_1 \geq i_2 \geq \dots \geq i_r$.

The *sequence number* of edge $e_i = v_{i-1}v_i \in P$ with respect to v_0 , denoted by $\text{sqn}(e_i, P, v_0)$, is i ; observe that the sequence number of e_i w.r.t. v_ℓ is $\ell - i + 1$. Note that if P is an odd path, then the sequence number of an edge e on P w.r.t. v_0 has the same parity as the sequence number of e w.r.t. v_ℓ . Therefore, in such a case we may talk about the parity of the sequence number of some edge $e \in P$ without specifying from which endpoint of P we count the sequence number, and we may also write $\text{sqn}(e, P)$ to denote the sequence number of e with respect to one (arbitrarily fixed) endpoint of P .

Two paths P_1 and P_2 are *openly disjoint* if they share neither edges nor vertices except for endpoints of both P_1 and P_2 . A *cycle* is a set of edges that can be partitioned into two openly disjoint paths between the same pair

of endpoints.

As mentioned before, if we have a non-negative weight function, then we can find a solution for the SHORTEST ODD PATH problem efficiently. Since finding a minimum-weight even path can be reduced to finding a minimum-weight odd path (as explained in Section 1), we have the following.

Theorem 1 ([25, 13]). *Given an undirected graph $G = (V, E)$ with vertices s and t and a non-negative weight function $w: E \rightarrow \mathbb{R}_{\geq 0}$, we can find a minimum-weight odd (s, t) -path and a minimum-weight even (s, t) -path in strongly polynomial time, whenever such paths exist.*

A set $T \subseteq E$ of edges in G is *connected* if for every pair of edges e and e' in T , there is a path contained in T containing both e and e' . If T is connected and *acyclic*, i.e., contains no cycle, then T is a *tree* in G . Note that in accordance with this definition, we interpret primarily mostly as edge sets; nevertheless, sometimes we may also interpret a tree $T \subseteq E$ as the subgraph of G spanned by the edge set T , and refer to the vertices in this subgraph as vertices *on* T . Given two vertices a and b in a tree T , we denote by $T[a, b]$ the unique path contained in T whose endpoints are a and b . For an edge $uv \in T$ and a path P in T such that $uv \notin P$, either $T[u, p]$ contains v for every vertex p on P , or $T[v, p]$ contains u for every vertex p on P . In the former case, we say that v is *closer to* P in T than u .

We say that a weight function $w: E \rightarrow \mathbb{R}$ is *conservative*, if no cycle in G has negative weight. Given a weight function $w: E \rightarrow \mathbb{R}$ on the edges of G , the weight of an edge set $F \subseteq E$ is $w(F) = \sum_{e \in F} w(e)$; similarly, the weight of a walk $W = (e_1, \dots, e_\ell)$ is $w(W) = \sum_{i=1}^{\ell} w(e_i)$.

Derandomization and universal sets. Given a randomized algorithm, a standard tool for derandomization is the use of universal sets (see e.g., Cygan et al. [10]). An (n, k) -*universal set* on an n -element set V is a family \mathcal{U} of subsets of V such that for any subset $S \subseteq V$ of size k , the family $\{A \cap S : A \in \mathcal{U}\}$ contains all 2^k subsets of S .

Theorem 2 ([19]). *For any positive integers n and k we can construct an (n, k) -universal set \mathcal{U} of size $|\mathcal{U}| = 2^k \cdot k^{\mathcal{O}(\log k)} \cdot \log n$ in time $\mathcal{O}(n \cdot |\mathcal{U}|)$.*

Treewidth. Intuitively, treewidth measures how similar a graph is to a tree. For an undirected graph G , a *tree decomposition* is a pair $(\mathbb{T}, \mathcal{B})$ of a tree \mathbb{T} and a collection $\mathcal{B} = \{B_x : x \in V(\mathbb{T})\}$ of *bags*, where each bag B_x associated with some node x of \mathbb{T} is a subset of vertices in G , with the following properties. For each edge of G , there is a bag in \mathcal{B} that contains both of

its endpoints. Furthermore, for each vertex v of G , the nodes $\{x : v \in B_x\}$ in \mathbb{T} induce a subtree of \mathbb{T} . The *width* of $(\mathbb{T}, \mathcal{B})$ is defined as the largest bag size in \mathcal{B} minus one. The *treewidth* of G is the minimum width over all tree decompositions of G .

Korhonen [15] gives a 2-approximation for treewidth in the following result that has a better running time than the known exact algorithms computing treewidth (e.g., Bodlaender [4] with running time $2^{\mathcal{O}(k^3)}n$).

Theorem 3 ([15]). *Given a graph G on n vertices and an integer k , there is an algorithm that outputs a tree decomposition of G of width at most $2k + 1$ or determines that the treewidth of G is larger than k , in $2^{\mathcal{O}(k)}n$ time.*

3. Initial observations

Let $G = (V, E)$ be an undirected graph with a conservative weight function $w: E \rightarrow \mathbb{R}$. Let $E^- = \{e \in E : w(e) < 0\}$ denote the set of negative edges, and \mathcal{T} be the set of negative trees they form. More precisely, let \mathcal{T} be the set of connected components in the subgraph of G spanned by E^- , i.e., \mathcal{T} contains those maximal, non-empty subsets of E^- that are connected in G ; the acyclicity of each $T \in \mathcal{T}$ follows from the conservativeness of w .

Lemma 4. *If W is a closed walk that does not contain any edge of E^- more than once, then $w(W) \geq w(E(W)) \geq 0$.*

Proof. Let Z be the set of edges appearing an odd number of times on W . Since any edge used at least twice in W has non-negative weight, we have $w(Z) \leq w(E(W)) \leq w(W)$. Moreover, since W is a closed walk, each vertex of V has an even degree in the graph (V, Z) . Therefore, (V, Z) is the union of edge-disjoint cycles, and the conservativeness of w implies $w(Z) \geq 0$. \square

The following definition plays a crucial role in our algorithms.

Definition 5. For a tree $T \in \mathcal{T}$, a *T -leap* L with endpoints a and b is a path such that $V(L) \cap V(T) = \{a, b\}$ and $T \cap L = \emptyset$. A *leap* is simply a T -leap for some $T \in \mathcal{T}$. For a path Q , a *leap on Q* is a subpath of Q that is a leap. Furthermore, given a T -leap L with endpoints a and b , we refer to the cycle $C = L \cup T[a, b]$ as the *cycle induced by L* . We say that L is *parity-changing* if $|C|$ is odd. Finally, the *shadow* of a T -leap L on Q with endpoints a and b is the edge set $Q \cap T[a, b]$, see Figure 1 for an illustration.




Figure 1: Illustration of a leap. Tree T is represented by straight lines, while the (x, y) -path Q is shown by curvy lines. The thick orange part of Q forms a T -leap denoted by L . The shadow of L is highlighted with green.

Remark 6. [Computing parity-changing leaps.] Suppose that \mathcal{T} contains a unique tree $T \subseteq E$. Then, given two vertices a and b on T , it is not difficult to obtain a parity-changing leap with endpoints a and b having minimum-weight: we simply have to search for a minimum-weight (a, b) -path in $G - T - (V(T) \setminus \{a, b\})$ whose parity differs from the parity of $T[a, b]$. As w restricted to $G - T$ is non-negative, we can apply Theorem 1 to find such a path.

Let us introduce a technical result.

Lemma 7. *Let w be a conservative weight function and T be a tree in \mathcal{T} . For two edge sets $Q \subseteq E$ and $R \subseteq T$, if $d_Q(v) \equiv d_R(v) \pmod{2}$ for each $v \in V$, then $w(Q) \geq w(R)$.*

Proof. Let $Q' = Q \setminus R$ and $R' = R \setminus Q$. Since Q' and R' are disjoint sets with $d_{Q'}(v) \equiv d_{R'}(v) \pmod{2}$ for any $v \in V$, the edge set $Q' \cup R'$ can be partitioned into cycles. Therefore, $w(Q') + w(R') \geq 0$ as w is conservative. This together with $w(R') \leq 0$ shows that $w(Q') \geq 0 \geq w(R')$. Using that $w(Q) = w(Q') + w(Q \cap R)$ and $w(R) = w(R') + w(Q \cap R)$, we obtain $w(Q) \geq w(R)$. \square

The lemma gives the following corollary by observing that we can assume w.l.o.g. that a walk Q is an (x, y) -path: if Q contains cycles, then we can repeatedly delete any cycle from Q , possibly of length 2, so that in the end we obtain an (x, y) -path whose weight is at most the weight of Q , since w is conservative.

Corollary 8. *Let x, y, x', y' be four distinct vertices on a tree T in \mathcal{T} .*

- (1) If Q is an (x, y) -walk in G that does not use any edge with negative weight more than once, then $w(Q) \geq w(T[x, y])$.
- (2) If Q is an (x, y) -path, Q' is an (x', y') -path, and Q and Q' are vertex-disjoint, then $w(Q) + w(Q') \geq w(T[x, y] \Delta T[x', y'])$.

Lemma 9. *Let Q be a minimum-weight odd (s, t) -path in G that contains as few leaps as possible (among all minimum-weight odd (s, t) -paths). Then either Q contains no leaps, or it contains at least one parity-changing leap.*

Proof. We show that whenever an optimal solution contains a leap, then it contains at least one parity-changing leap. Let Q be a minimum-weight odd (s, t) -path with a minimum number of leaps, and suppose that Q contains at least one leap, but no parity-changing leap. Let T be a tree in \mathcal{T} for which Q contains a T -leap, and let x and y be the first and the last vertex, respectively, of Q on T when we traverse Q in an arbitrarily fixed direction. Clearly, $x \neq y$ as Q contains a T -leap and is a path.

Consider the walk W that we obtain if we traverse $Q[x, y]$ in one direction with the modification that whenever we reach a T -leap L , we replace L by the path contained in T between the endpoints of L . Since no leap on Q is parity-changing, $|W|$ has the same parity as $|Q[x, y]|$. Furthermore, any walk in T between two fixed vertices x and y has the same parity as $|Q[x, y]|$, thus we can replace all edges of $Q[x, y]$ by the path $T[x, y]$ without changing the parity of Q . Therefore, $Q \setminus Q[x, y] \cup T[x, y]$ is an odd (s, t) -path, and its weight is at most $w(Q)$ by statement (1) of Corollary 8. Moreover, it contains fewer leaps than Q , contradicting the definition of Q . \square

We remark that if \mathcal{T} contains only a single tree, then even a stronger statement holds: a minimum-weight odd (s, t) -path either uses no leaps, or it uses *exactly one* parity-changing leap. However, even in this special case, a minimum-weight odd (s, t) -path may need to use an arbitrarily large number of leaps; see Figure 2 for an example. By contrast, a shortest (s, t) -path intersects each tree $T \in \mathcal{T}$ in a single (possibly empty) path, so it never uses leaps. The next result is quite trivial, but we present its proof for completeness.

Lemma 10. *Let P be a minimum-weight (s, t) -path in G , and let T be a tree in \mathcal{T} . Then for any two vertices u and v in $V(P) \cap V(T)$, it holds that $P[u, v] = T[u, v]$.*

Proof. If the statement of the lemma fails, then there exist vertices u and v in T such that the only vertices of P appearing on $T[u, v]$ are u and v , and




Figure 2: An example where there is a unique odd (s, t) -path P . The negative tree T^- is shown with purple double lines, and the unique parity-changing leap L is shown with an orange thick line. The path P must use all leaps in the figure.

w is not an edge of P . Then Lemma 4 implies $w(P[u, v]) + w(T[u, v]) \geq 0$. Thus, replacing the subpath $P[u, v]$ of P with $T[u, v]$ results in an (s, t) -path whose weight is strictly less than $w(P)$, a contradiction. \square

4. Parity Constrained Odd Paths

In this section, we consider the SHORTEST PARITY-CONSTRAINED ODD PATH problem, a variant of SHORTEST ODD PATH in which some edges have parity constraints on their position along the path. Let us recall the following definition stated in Section 1.


Definition 11. Let $F_{\text{even}} \subseteq E$ and $F_{\text{odd}} \subseteq E$ be two disjoint edge sets in G . An odd path Q is $(F_{\text{even}}, F_{\text{odd}})$ -constrained if the sequence number of each edge in $Q \cap F_{\text{even}}$ is even, and the sequence number of each edge in $Q \cap F_{\text{odd}}$ is odd. For an illustration, see Figure 3a.

SHORTEST PARITY-CONSTRAINED ODD PATH:

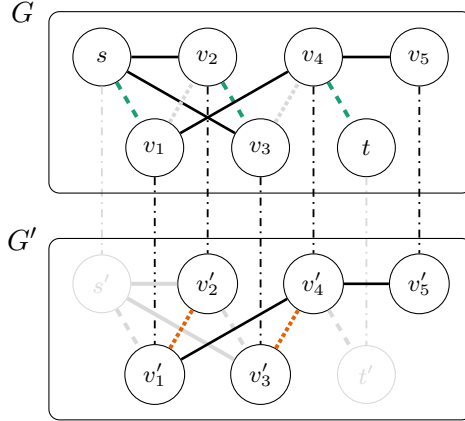
Input: An undirected graph $G = (V, E)$, a weight function $w: E \rightarrow \mathbb{R}$, two vertices s and t in G , and disjoint edge sets $F_{\text{even}} \subseteq E$ and $F_{\text{odd}} \subseteq E$.

Goal: Find a minimum-weight $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path.

This problem is a key step for our algorithms solving the SHORTEST ODD PATH problem. We show that SHORTEST PARITY-CONSTRAINED ODD PATH can be solved in polynomial time if all negative edges are contained in $F_{\text{even}} \cup F_{\text{odd}}$, assuming conservative weights.



(a) Illustration of Definition 11.



(b) Construction for Theorem 12.

Figure 3: (a) Edges in F_{even} are shown as dotted orange lines, F_{odd} as dashed green lines, and $E \setminus (F_{\text{even}} \cup F_{\text{odd}})$ as solid black. Note that there exists three $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -paths, namely $Q_1 = s, v_1, v_2, v_3, v_4, t$, $Q_2 = s, v_3, v_4, t$ and $Q_3 = s, v_1, v_4, t$. (b) Graph H obtained from (a). Removed arcs and vertices are shown in gray. Dashedotted lines have weight 0, while the rest have original weight.

Theorem 12. *If w is conservative and $F_{\text{even}} \cup F_{\text{odd}}$ contains all negative edges of the input graph $G = (V, E)$, then the SHORTEST PARITY-CONSTRAINED ODD PATH problem can be solved in $\mathcal{O}(mn + n^2 \log n)$ time, where $n = |V|$ and $m = |E|$.*

Proof. We present a generalization of the well-known algorithm for determining a shortest odd (or, equivalently, even) path between two vertices in an undirected graph with non-negative edge weights.

Let $G = (V, E)$, and let $(G, w, s, t, F_{\text{even}}, F_{\text{odd}})$ be our input for SHORTEST PARITY-CONSTRAINED ODD PATH such that $F_{\text{even}} \cup F_{\text{odd}}$ contains all edges with negative weight. An illustration of the following construction is shown in Figure 3b. We start by creating a copy G' of G where each vertex or edge x in G has a copy x' in G' . The copy of each edge $e \in E$ has the same weight as e , and additionally, we add a new edge between each vertex $v \in V$ and its copy v' of weight 0. Finally we delete F_{even} from G , and we delete the edge set $F'_{\text{odd}} := \{e' : e \in F_{\text{odd}}\}$ as well as vertices s' and t' from G' . Let H denote the graph thus obtained.

We claim that the minimum weight of an $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path in G equals the minimum weight of a perfect matching in H . First, let Q be an $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path. We define an

edge set M_Q as follows:

$$M_Q := \{e : e \in Q, \text{sqn}(e, Q) \text{ is odd}\} \cup \{e' : e \in Q, \text{sqn}(e, Q) \text{ is even}\} \\ \cup \{vv' : v \in V \setminus V(Q)\}.$$

Observe first that since Q is $(F_{\text{even}}, F_{\text{odd}})$ -constrained, each edge of M_Q is present in H . Therefore, it is easy to see that M_Q is a perfect matching in H , and its weight is exactly $w(Q)$.

Second, let M be a perfect matching in H . Let $M_1 = \{e \in E : e \in M\}$ and $M_2 = \{e \in E : e' \in M\}$. Consider the multiset $Q_M = M_1 \uplus M_2$ of edges. Note that for each vertex v in $V \setminus \{s, t\}$, either v is incident to one edge of M_1 and one edge of M_2 (counting multiplicities), or $vv' \in M$. Vertices s and t each are incident to one edge of M_1 and no edge of M_2 . Moreover, an edge $e \in E$ may be present both in M_1 and M_2 , but such an edge cannot have negative weight since $F_{\text{even}} \cup F_{\text{odd}}$ contains all edges with negative weight.

Therefore, each connected component of the subgraph $(V, M_1 \cup M_2)$ of G is either an odd (s, t) -path Q in G (alternating between M_1 and M_2), an even cycle in G , or a single edge in $E \setminus (F_{\text{even}} \cup F_{\text{odd}})$. Recall that each cycle in G has non-negative weight, as w is conservative. Moreover, each edge $e \in E \setminus (F_{\text{even}} \cup F_{\text{odd}})$ has non-negative weight by our assumption $E^- \subseteq F_{\text{even}} \cup F_{\text{odd}}$. Thus $w(Q) \leq w(M)$, and hence Q is an odd (s, t) -path of weight at most $w(M)$. Finally, observe that any edge of $F_{\text{even}} \cap Q$ belongs to M_2 (and not M_1), and thus has an even sequence number in Q , and similarly, any edge of $F_{\text{odd}} \cap Q$ belongs to M_1 (and not M_2), and thus has an odd sequence number in Q , by our construction of H . Therefore, Q is indeed an $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path, and has weight at most $w(M)$. This proves our claim.

As a consequence, in order to find a minimum-weight $(F_{\text{even}}, F_{\text{odd}})$ -constrained odd (s, t) -path, we simply have to construct H , find a minimum-weight perfect matching in H using, e.g., Gabow's algorithm [12], and build the corresponding odd (s, t) -path as explained above. The overall running time is $\mathcal{O}(mn + n^2 \log n)$. \square

5. Negative edges forming a tree

In this section, we consider the special case of SHORTEST ODD PATH when the weight function w is conservative, and the set of edges whose weight is negative forms a tree T^- in G . We give an algorithm for solving the problem in polynomial time, presented as Algorithm 1. An illustration can be seen in Figure 4.

The high-level idea of the algorithm is as follows. By Lemma 9, any shortest odd (s, t) -path either contains no leaps, or it contains a parity-changing leap, hence our algorithm checks both cases and it generates two types of paths. First, to find a minimum-weight odd (s, t) -path P among those that contain no leaps, we guess the first and the last vertices of P that appear on T^- by picking two vertices a and b . Then we partition the path $T^-[a, b]$ into two sets of alternating edges, and solve an instance of the SHORTEST PARITY-CONSTRAINED ODD PATH problem with this partition. This gives us the first type of path. Second, to find a minimum-weight odd (s, t) -path P among those that contain a parity-changing leap L , we guess the endpoints of L by picking two vertices a and b . Then we compute a minimum-weight parity-changing leap L' with endpoints a and b , and construct an odd (s, t) -path by using L' and two vertex-disjoint paths from $\{s, t\}$ to $\{a, b\}$ of minimum weight. This gives us the second type of path. The output of the algorithm is the best path among these two types of solutions.

Remark 13. The task on line 13 is to solve an instance of the SHORTEST TWO DISJOINT PATHS problem that, given an undirected graph with conservative edge weights and terminals s_1, s_2, t_1 , and t_2 , the goal is to find two vertex-disjoint paths leading from $\{s_1, s_2\}$ to $\{t_1, t_2\}$ with minimum total weight. This problem was shown to be NP-hard by Schlotter and Sebő [23] even with a conservative weight function of the form $w: E \rightarrow \{-1, 1\}$. Nevertheless, if the negative-weight edges form a tree, as in our case, then the problem can be solved in polynomial time [22].

The algorithm uses the following definition.

Definition 14. Given two distinct vertices a and b on a cycle C , we say that they *cut* C into paths C_1 and C_2 if both C_1 and C_2 have endpoints a and b , and $C_1 \cup C_2 = C$.

We first prove the correctness of Algorithm 1.

Lemma 15. *Algorithm 1 correctly solves the SHORTEST ODD PATH problem when the weight function is conservative and negative edges form a tree.*

Proof. First we show that any path that the algorithm sets as S is an odd (s, t) -path. This is trivial for the paths on lines 7 and 9. Consider now the path S^* on line 20. Clearly, the paths P'_s and P'_t are vertex-disjoint, since P_s and P_t are vertex-disjoint by definition. Second, the vertex shared by P'_s and C is x , by the definition of x , so the paths P'_s, C_1 , and C_2 are pairwise openly disjoint. Similarly, by our choice of y , the paths P'_t, C_1 , and C_2 are

Algorithm 1 Solving SHORTEST ODD PATH with conservative weight function and the set of negative edges forms a tree.

Input: A graph $G = (V, E)$, a conservative weight function $w: E \rightarrow \mathbb{R}$, vertices $s, t \in V$, a set of negative edges $T^- = \{e \in E, w(e) < 0\}$ forming a tree.

Output: A minimum-weight odd (s, t) -path or \emptyset if no such path exists.

- 1: Set $S := \emptyset$, $w_S := \infty$.
 - 2: **for all** distinct vertices a and b on T^- **do**
 - \triangleright *Computations for first type of paths.*
 - 3: Set $F_{\text{even}}^a := \{e \in T^-[a, b] : \text{sqn}(e, T^-[a, b], a) \equiv 0 \pmod{2}\}$.
 - 4: Set $F_{\text{odd}}^a := \{e \in T^-[a, b] : \text{sqn}(e, T^-[a, b], a) \equiv 1 \pmod{2}\}$.
 - 5: Set $G' := G - (T^- \setminus T^-[a, b])$.
 - 6: Compute
 $Q := \text{SHORTEST PARITY-CONSTRAINED ODD PATH}(G', w, s, t, F_{\text{even}}^a, F_{\text{odd}}^a)$.
 - 7: **if** $w(Q) < w_S$ **then** set $S := Q$, $w_S := w(S)$.
 - 8: Compute
 $Q := \text{SHORTEST PARITY-CONSTRAINED ODD PATH}(G', w, s, t, F_{\text{odd}}^a, F_{\text{even}}^a)$.
 - 9: **if** $w(Q) < w_S$ **then** set $S := Q$, $w_S := w(S)$.
 - 10: **if** \exists parity-changing leap from a to b **then**
 - \triangleright *Computations for second type of paths.*
 - 11: Compute a minimum-weight parity-changing leap L with endpoints a and b in G . \triangleright *See Remark 6.*
 - 12: **if** \exists two vertex-disjoint paths from $\{s, t\}$ to $\{a, b\}$ **then**
 - 13: Compute two vertex-disjoint paths P_s and P_t from $\{s, t\}$ to $\{a, b\}$ in G that minimizes $w(P_s) + w(P_t)$, with $s \in V(P_s)$ and $t \in V(P_t)$. \triangleright *See Remark 13.*
 - 14: Set $C := L \cup T^-[a, b]$.
 - 15: Let x be the vertex of C closest to s on P_s , and set $P'_s := P_s[s, x]$.
 - 16: Let y be the vertex of C closest to t on P_t , and set $P'_t := P_t[y, t]$.
 - 17: Let C_1 and C_2 be the two paths into which x and y cut C .
 - 18: Set $S_1 := P'_s \cup C_1 \cup P'_t$ and $S_2 := P'_s \cup C_2 \cup P'_t$.
 - 19: **if** S_1 is odd **then** $S^* := S_1$ **else** $S^* := S_2$.
 - 20: **if** $w(S^*) < w_S$ **then** set $S := S^*$, $w_S := w(S^*)$.
 - 21: **if** $S \neq \emptyset$ **then return** S ;
 - 22: **else return** \emptyset .
-

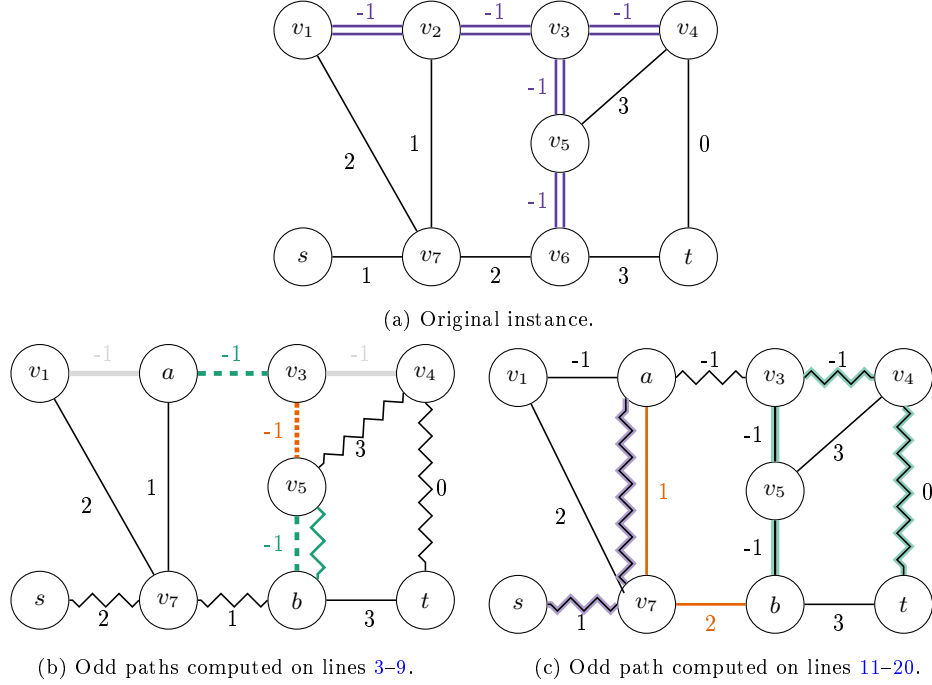


Figure 4: Example of SHORTEST ODD PATH instance with conservative weights when running Algorithm 1 for $a = v_2$ and $b = v_6$. With these specific a and b , the solution generated up to this point is given in Figure (c).

- (a) Original instance. Tree T^- is shown using purple double lines.
- (b) Paths obtained on lines 3–9. Deleted edges are shown in gray. Set F_{even}^a is shown by dotted orange lines, and F_{odd}^a by dashed green lines. Solving the SHORTEST PARITY-CONSTRAINED ODD PATH problem with instance $(G', w, s, t, F_{\text{odd}}^a, F_{\text{even}}^a)$ generates the path on vertices s, v_7, b, t which is a shortest odd (s, t) -path in $G - T^-$. Solving the instance $(G', w, s, t, F_{\text{even}}^a, F_{\text{odd}}^a)$ generates the zigzagged path.
- (c) Path obtained on lines 11–20 shown in zigzag. An (s, a) -path of minimum weight is highlighted in purple and a (t, b) -path of minimum weight is highlighted in green. The parity-changing leap $L = \{(a, v_7), (v_7, b)\}$ is shown in orange.

pairwise openly disjoint. Hence, $S_i = P'_s \cup C_i \cup P'_t$ is a path for both $i = 1, 2$. Moreover, since L is a parity-changing leap, we know that $|C_1| + |C_2|$ is odd. Therefore, exactly one of S_1 and S_2 is an odd path, and the algorithm considers it as S^* on line 20.

Thus, to prove the correctness of our algorithm, it suffices to show that if there exists an odd (s, t) -path of weight w^* , then our algorithm picks an odd (s, t) -path of weight at most w^* .

Let Q^* be a shortest odd (s, t) -path with weight $w(Q^*) = w^*$ containing as few leaps as possible. By Lemma 9, either Q^* contains no leaps, or it

contains a parity-changing leap. See an illustration of the algorithm on Figure 4.

First assume that Q^* contains no leaps. Then, by Lemma 10 either (i) $Q^* \cap T^- = \emptyset$, or (ii) $Q^* \cap T^-$ is a path containing at least one edge. In case (i), observe that Q^* is a path in $G - T^-$. Hence, it is an (F_1, F_2) -constrained path for every pair (F_1, F_2) of subsets of T^- . In fact, for every choice of a and b , Q^* is a solution for the two SHORTEST PARITY-CONSTRAINED ODD PATH instances constructed on lines 6 and 8. Consequently, the algorithm will find an odd (s, t) -path with weight at most w^* on both lines 6 and 8, and will select the one with minimum weight in either line 7 or line 9. In case (ii), let a and b be the endpoints of the path $Q^* \cap T^-$, and consider the iteration of Algorithm 1 when the algorithm picks vertices a and b in the for-loop on line 2. Let us define F_{even}^a and F_{odd}^a as on lines 3 and 4 of Algorithm 1. Clearly, Q^* is either $(F_{\text{even}}^a, F_{\text{odd}}^a)$ -constrained, or $(F_{\text{odd}}^a, F_{\text{even}}^a)$ -constrained for a pair of vertices a, b in T , moreover, it is a path in G' obtained by deleting all edges of T from G except those of $T^-[a, b]$. Hence, Algorithm 1 stores a path of weight at most $w(Q^*) = w^*$ either on line 7 or on line 9. This shows the correctness of the algorithm in the case when Q^* contains no leaps.

Assume now that Q^* contains a parity-changing leap L^* . Let a and b denote the endpoints of L^* . Let us consider the iteration of Algorithm 1 when the for-loop on line 2 chooses a and b ; we use the definitions on lines 11–20. Note that P_s and P_t exist, since $Q^* \setminus L^*$ decomposes into two vertex-disjoint paths from $\{s, t\}$ to $\{a, b\}$. Observe that $S_1 \triangle S_2 = C$, and since L is parity-changing we have that $|C|$ is odd. Hence, exactly one of S_1 and S_2 is an odd (s, t) -path. We claim that $w(S_1) \leq w^*$ and $w(S_2) \leq w^*$. Note that if our claim holds, then Algorithm 1 stores a path of weight at most w^* on line 20, and hence the correctness of the algorithm follows. Therefore, to show that Algorithm 1 is correct, it suffices to prove that

$$\max\{w(S_1), w(S_2)\} \leq w^*. \quad (1)$$

Let Q_s^* and Q_t^* be the two vertex disjoint paths formed by $Q^* \setminus L^*$. Since Q_s^* and Q_t^* lead from $\{s, t\}$ to $\{a, b\}$, the definition of P_s and P_t implies $w(P_s) + w(P_t) \leq w(Q_s^*) + w(Q_t^*)$. Moreover, since L is a minimum-weight parity-changing leap with endpoints a and b , we also know

$$w(P_s) + w(P_t) + w(L) \leq w(Q_s^*) + w(Q_t^*) + w(L^*) = w(Q^*) = w^*. \quad (2)$$

We distinguish between three cases to prove inequality (1).

Case A: x and y are both on $T^-[a, b]$. W.l.o.g., the vertices a, x, y, b follow each other in this order on $T^-[a, b]$, as otherwise we switch the names for a and b . Note that $P_s \setminus P'_s$ and $P_t \setminus P'_t$ are two vertex-disjoint paths leading from $\{x, y\}$ to $\{a, b\}$, so by statement (2) of Corollary 8, we have

$$w(P_s \setminus P'_s) + w(P_t \setminus P'_t) \geq w(T^-[x, a] \cup T^-[b, y]). \quad (3)$$

Observe also that among the paths S_1 and S_2 , the one containing leap L has the larger weight since $w(T^-[x, y]) < 0$. Hence, $w(C\Delta T^-[x, y]) > 0$. Therefore, inequality (3) together with inequality (2) implies

$$\begin{aligned} \max\{w(S_1), w(S_2)\} &= w(P'_s) + w(T^-[x, a]) + w(L) + w(T^-[b, y]) + w(P'_t) \\ &\leq w(P'_s) + w(P_s \setminus P'_s) + w(L) + w(P_t \setminus P'_t) + w(P'_t) \\ &= w(P_s) + w(L) + w(P_t) \\ &\leq w^*. \end{aligned}$$

Case B: x is on $T^-[a, b]$, but y is on L . We may assume that $x \in V(P_s)$ and $y \in V(P_t)$, as otherwise we can switch the names of s and t . Next, we may also assume that P_s ends at b , and P_t ends at a , as otherwise we can switch the names of a and b . Furthermore, we may even assume that $a \in V(S_1)$ and $b \in V(S_2)$, as otherwise we can switch the names of S_1 and S_2 . This means that $P_s \setminus P'_s$ is an (x, b) -path, and $P_t \setminus P'_t$ is an (a, y) -path; recall that they are vertex-disjoint as well, by the definition of P_s and P_t . Hence, $(P_s \setminus P'_s) \cup L[y, b] \cup (P_t \setminus P'_t)$ is an (x, a) -walk, and moreover, it does not contain any negative-weight edge more than once (recall that $L \cap T^- = \emptyset$, since L is a leap). By statement (1) of Corollary 8, this implies

$$w(P_s \setminus P'_s) + w(L[y, b]) + w(P_t \setminus P'_t) \geq w(T^-[x, a]).$$

From this, using also equation (2) we obtain that

$$\begin{aligned} w(S_1) &= w(P'_s) + w(T^-[x, a]) + w(L[y, a]) + w(P'_t) \\ &\leq w(P'_s) + w(P_s \setminus P'_s) + w(L[y, b]) + w(P_t \setminus P'_t) + w(L[y, a]) + w(P'_t) \\ &= w(P_s) + w(L) + w(P_t) \\ &\leq w^*. \end{aligned}$$

To prove $w(S_2) \leq w^*$ as well, we need to apply statement (1) of Corollary 8 for the (b, x) -path $P_s \setminus P'_s$ to get

$$w(P_s \setminus P'_s) \geq w(T^-[b, x]). \quad (4)$$

Furthermore, note that $(P_t \setminus P'_t) \cup L[y, a]$ is a closed walk that does not contain any edge of T^- more than once, since $P_t \setminus P'_t$ is a path and $L \cap T^- = \emptyset$. Thus, Lemma 4 yields

$$w(P_t \setminus P'_t) + w(L[y, a]) \geq 0. \quad (5)$$

Putting inequalities (2), (4) and (5) together we obtain

$$\begin{aligned} w(S_2) &= w(P'_s) + w(T^-[b, x]) + w(L[y, b]) + w(P'_t) \\ &\leq w(P'_s) + w(P_s \setminus P'_s) + w(L) - w(L[y, a]) + w(P'_t) \\ &\leq w(P_s) + w(L) + w(P_t \setminus P'_t) + w(P'_t) \\ &= w(P_s) + w(L) + w(P_t) \\ &\leq w^*. \end{aligned}$$

This shows inequality (1) for Case B.

Case C: Both x and y are on L . W.l.o.g., assume that the vertices a, x, y, b follow each other in this order on the leap L , as otherwise we can switch the names of x and y . We may also assume that $S_1 = P'_s \cup L[x, y] \cup P'_t$ while $S_2 = S_1 \Delta C$, as otherwise we can switch the names of S_1 and S_2 . Note that $P_s \setminus P'_s$ and $P_t \setminus P'_t$ are two paths from $\{x, y\}$ to $\{a, b\}$ and they are vertex-disjoint. Thus, their union together with $L[x, y]$ yields a walk from a to b that contains no negative-weight edge more than once. Statement (1) of Corollary 8 then implies

$$w(P_s \setminus P'_s) + w(L[x, y]) + w(P_t \setminus P'_t) \geq w(T^-[a, b]).$$

Hence, using also inequality (2), we obtain

$$\begin{aligned} w(S_2) &= w(P'_s) + w(L[x, a]) + w(T^-[a, b]) + w(L[b, y]) + w(P'_t) \\ &\leq w(P'_s) + w(L[x, a]) + w(P_s \setminus P'_s) + w(L[x, y]) + w(P_t \setminus P'_t) \\ &\quad + w(L[b, y]) + w(P'_t) \\ &= w(P_s) + w(L) + w(P_t) \\ &\leq w^*. \end{aligned}$$

To show $w(S_1) \leq w^*$, observe that the paths $P_s \setminus P'_s$ and $P_t \setminus P'_t$ together with $L[a, x]$ and $L[b, y]$ form either one or two closed walks, depending on whether P_s ends at a or at b . Moreover, no edge of T^- appears more than once on these closed walks. Therefore, Lemma 4 yields

$$w(P_s \setminus P'_s) + w(P_t \setminus P'_t) + w(L[a, x]) + w(L[b, y]) \geq 0.$$

Thus, we obtain

$$\begin{aligned}
w(S_1) &= w(P'_s) + w(L[x, y]) + w(P'_t) \\
&\leq w(P'_s) + w(L[x, y]) + w(P'_t) + w(P_s \setminus P'_s) + w(P_t \setminus P'_t) \\
&\quad + w(L[a, x]) + w(L[b, y]) \\
&= w(P_s) + w(L) + w(P_t) \\
&\leq w^*. \quad \square
\end{aligned}$$

Theorem 16. SHORTEST ODD PATH can be solved in strongly polynomial time, if the weight function is conservative and the set of negative edges forms a tree.

Proof. By Lemma 15 we know that Algorithm 1 is correct, so it remains to bound its running time. Let $n = |V|$ and $m = |E|$ denote the number of vertices and edges in the input graph G .

Clearly, the **for**-loop on line 2 is iterated $\binom{n}{2}$ times. Performing lines 3–9 can be done in $\mathcal{O}(mn + n^2 \log n)$ by Theorem 12. For line 11, by Remark 6 we can obtain a parity-changing leap of minimum weight in strongly polynomial time, in fact, in time $\mathcal{O}(mn + n^2 \log n)$ [24].

To perform line 13, we rely on Remark 13. Lines 15–20 can be performed in $\mathcal{O}(n + m)$ time. Hence, if solving an instance of the SHORTEST TWO DISJOINT PATHS problem on G on line 13 takes $F(n, m)$ time, then we obtain an overall running time of $\mathcal{O}(n^2(mn + n^2 \log n + F(n, m)))$. Note that currently the only known upper bound on $F(n, m)$ is $F(n, m) = \mathcal{O}(n^9)$ [22], but probably this can be improved. \square

We remark that there is a strong connection between the SHORTEST ODD PATH and the SHORTEST TWO DISJOINT PATHS problems. On the one hand, line 13 of our algorithm for SHORTEST ODD PATH relies on finding two vertex-disjoint paths from two source terminals to two sink terminals. The need to use such an algorithm as a subroutine seems inevitable: if the input graph admits a unique parity-changing leap L , then in order to find a shortest odd path between s and t , we may need to find two vertex-disjoint paths that lead from s and t to two vertices of the odd cycle induced by L .

On the other hand, there is a very simple reduction from SHORTEST TWO DISJOINT PATHS to SHORTEST ODD PATH. To see this, consider the equivalent formulation of SHORTEST TWO DISJOINT PATHS where instead of two source and two sink terminals in G , we are only given two terminals, s and t , and the task is to find two openly disjoint (s, t) -paths in G while minimizing their total weight. To solve this problem using an algorithm

for SHORTEST ODD PATH, we first make a copy s' of s and a copy t' of t , and connect the vertices adjacent to s to s' , and the vertices adjacent to t to t' . Then, subdivide each edge of this graph (halving the weights as well), and add a new edge tt' of weight 0. Let us denote by G' this resulting graph. We obtain our two (s, t) -paths by computing an odd (s, s') -path of minimum weight in graph G' . Indeed, any pair of openly disjoint (s, t) -paths in G corresponds to an odd (s, s') -path of the same weight in G' , and vice versa. This reduction shows that the connection between these two problems is bidirectional.

6. FPT algorithms

In this section, we present two FPT algorithms for solving SHORTEST ODD PATH when the weight function is conservative. In Section 6.1 we present two algorithms related to E^- the set of edges of negative weight: the first one is regarding the cardinality of E^- and the second one uses the size of the maximum matching in $G[E^-]$. Later, in Section 6.2 we propose an FPT algorithm where the parameter is the treewidth of the input graph; this algorithm does not rely on conservativeness but works for arbitrary weights.

6.1. Parametrization by negative edges

This section is divided in two. First, we show a simple FPT algorithm when the parameter is the number of negative edges. Then, we proceed to a more interesting parameter, namely the size of a maximum matching in the graph spanned by all negative-weight edges.

The key idea of the algorithms is to guess for each negative edge e whether e has an even or odd sequence number on an optimal path. With this guess, we run an instance of the SHORTEST PARITY-CONSTRAINED ODD PATH problem and obtain our solution.

6.1.1. The number of negative edges as parameter

We show that we can solve SHORTEST ODD PATH by an FPT algorithm with $|E^-|$ as the parameter.

Theorem 17. *There is an FPT algorithm for SHORTEST ODD PATH with conservative weights when parameterized by the number $|E^-|$ of negative edges, running in $\mathcal{O}\left(2^{|E^-|}(mn + n^2 \log n)\right)$ time.*

Proof. Let (G, w, s, t) be our input for SHORTEST ODD PATH with graph $G = (V, E)$, such that $E^- \subseteq E$ denotes the set of all edges with negative

weight. We guess a partition $(E_{\text{even}}^-, E_{\text{odd}}^-)$ of E^- so that E_{even}^- and E_{odd}^- contains all negative edges with even and odd sequence number in an optimal solution for (G, w, s, t) , respectively. Then we solve the SHORTEST PARITY-CONSTRAINED ODD PATH problem with instance $(G, w, s, t, E_{\text{even}}^-, E_{\text{odd}}^-)$ using Theorem 12, and return the path Q obtained.

To check the correctness of this algorithm, let Q^* be an optimal solution for SHORTEST ODD PATH. As Q is an odd (s, t) -path, $w(Q) \geq w(Q^*)$. Conversely, if our guesses are correct, then Q^* is an $(E_{\text{even}}^-, E_{\text{odd}}^-)$ -constrained odd (s, t) -path, and thus $w(Q) \leq w(Q^*)$.

For the running time, it is clear that there are $2^{|E^-|}$ possible guesses, and getting a solution for the SHORTEST PARITY-CONSTRAINED ODD PATH problem can be done in $\mathcal{O}(mn + n^2 \log n)$ time by Theorem 12. Hence, the total running time is as stated. \square

6.1.2. The size of a maximum matching on the negative edges as parameter

Let us now generalize the algorithm of Theorem 17. Let $\mu(E^-)$ be the size of a maximum matching in the graph $G[E^-]$. Note that $\mu(E^-) \leq |E^-|$ is trivial, and it is also clear that $|E^-|$ can be arbitrarily large for any fixed value of $\mu(E^-)$; to see this, consider the case when E^- is a collection of $\mu(E^-)$ stars of arbitrary sizes. Thus, the parameter $\mu(E^-)$ is substantially weaker than the parameter $|E^-|$. In the following, $\mu(E^-)$ is simply denoted by μ .

Theorem 18. *There is a randomized FPT algorithm for SHORTEST ODD PATH with conservative weights that returns an optimal solution to the problem with probability at least $1 - e^{-1} > 0.632$ in $\mathcal{O}(2^{2\mu}(mn + n^2 \log n))$ time.*

Proof. Let (G, w, s, t) be an instance of the SHORTEST ODD PATH problem with conservative weights. Let $E^- \subseteq E$ denote the set of all edges with negative weight. We would like to get an instance of the SHORTEST PARITY-CONSTRAINED ODD PATH problem to solve our problem. Again, as in Theorem 17, we guess a random partition $(E_{\text{even}}^-, E_{\text{odd}}^-)$ of E^- by placing each edge $e \in E^-$ into E_{even}^- with probability $\frac{1}{2}$, and additionally setting $E_{\text{odd}}^- = E^- \setminus E_{\text{even}}^-$. Then we obtain an optimal solution Q for the SHORTEST PARITY-CONSTRAINED ODD PATH problem with instance $(G, w, s, t, E_{\text{even}}^-, E_{\text{odd}}^-)$ using Theorem 12. We repeat this procedure τ times and pick the path Q with the minimum weight, where the value of τ is chosen later.

The key observation is that if Q^* is an optimal solution of the SHORTEST ODD PATH problem, then the set $Q_{\text{even}}^* \subseteq E^-$ of edges with an even sequence number in Q^* form a matching in G , and similarly, the set $Q_{\text{odd}}^* \subseteq E^-$ of edges with an odd sequence number in Q^* form a matching as well. This

implies $|Q_{\text{even}}^*| \leq \mu$ and $|Q_{\text{odd}}^*| \leq \mu$, so Q^* contains at most 2μ negative-weight edges. Hence, we need to guess at most 2μ edges of E^- correctly.

The probability of getting a guess where $Q_{\text{even}}^* \subseteq E_{\text{even}}^-$ is $(\frac{1}{2})^\mu$. Similarly, the probability of getting a guess where $Q_{\text{odd}}^* \subseteq E_{\text{odd}}^-$ is $(\frac{1}{2})^\mu$. Since there might be several optimal paths, the probability of getting a good set E_{even}^- is at least $(\frac{1}{2})^{2\mu}$. Thus, we return a wrong path Q with probability at most

$$\left(1 - \left(\frac{1}{2}\right)^{2\mu}\right)^\tau \leq \left(e^{-\left(\frac{1}{2}\right)^{2\mu}}\right)^\tau,$$

where we use the inequality $1 + x \leq e^x$. This implies that we succeed in returning a minimum-weight path with probability $1 - \left(e^{-\left(\frac{1}{2}\right)^{2\mu}}\right)^\tau$. Setting $\tau = 2^{2\mu}$ as the number of repetitions, our algorithm succeeds with probability at least $1 - e^{-1} > 0.632$.

The running time depends on the time we need to compute each path Q . By Theorem 12 and the value of τ , we obtain $\mathcal{O}(2^{2\mu}(mn + n^2 \log n))$ as overall running time. \square

It is straightforward to derandomize the algorithm using a universal set.

Theorem 19. *There is a deterministic FPT algorithm for SHORTEST ODD PATH with conservative weights when parameterized by μ , running in $\mathcal{O}(f(\mu) \cdot (mn \log n + n^2(\log n)^2))$ time, where $f(\mu) = 2^{2\mu} \cdot \mu^{\mathcal{O}(\log \mu)}$.*

Proof. Let (G, w, s, t) be our instance of SHORTEST ODD PATH; we use the notation in the proof of Theorem 18. By Theorem 2, one can construct an $(m^-, 2\mu)$ -universal set family \mathcal{U} on E^- of size $|\mathcal{U}| = 2^{2\mu} \cdot \mu^{\mathcal{O}(\log \mu)} \cdot \log m^-$ in time $\mathcal{O}(m^- \cdot |\mathcal{U}|)$, where $m^- := |E^-| = \mathcal{O}(m) = \mathcal{O}(n^2)$. Hence, instead of guessing E_{even}^- randomly as in Theorem 18, we can try all sets $U \in \mathcal{U}$, and define $E_{\text{even}}^- = U \cap E^-$. By the definition of an $(n, 2\mu)$ -universal set family we obtain that \mathcal{U} contains at least one set U^* such that $Q_{\text{even}}^* \subseteq U^* \cap E^-$ and $Q_{\text{odd}}^* \cap U^* = \emptyset$, which suffices for our algorithm to produce a correct output. The total running time is $\mathcal{O}(|\mathcal{U}| \cdot (mn + n^2 \log n))$. \square

6.2. Parametrization by treewidth

In this section we show that SHORTEST ODD PATH can be solved in linear time on any graph whose treewidth is bounded by a constant, even when the weight function is not conservative. The main idea is to show that this problem can be described by monadic second-order formulas.

First, we provide a brief introduction into monadic second-order logic in Section 6.2.1, and then present our FPT algorithm for SHORTEST ODD PATH parameterized by treewidth in Section 6.2.2.

6.2.1. Monadic second-order logic on graphs

Monadic second-order logic can be used to describe graph properties. A *first-order formula on graphs* is composed of atomic formulas “ $x = y$ ” and “ $\text{Adj}(x, y)$ ” over variables using logical connectives \wedge (**and**), \vee (**or**) and \neg (**negation**) and logical quantifiers \forall and \exists . Every formula defines a graph property by interpreting variables as vertices, and the relation $\text{Adj}(x, y)$ as the vertices corresponding to x and y being adjacent. To facilitate the reading, we also make use of the usual symbols \Rightarrow , \Leftrightarrow , and \nexists .

A wider class of graph properties can be described by monadic second-order logic. A *monadic second-order formula* has four types of variables: variables for vertices and edges (denoted by lower-case letters) and variables for sets of vertices and sets of edges (denoted by upper-case letters). The atomic formulas here also contain the $\text{Inc}(e, v)$ relation, interpreted as edge e being incident to vertex v , and the relation of inclusion, denoted as “ $x \in X$ ”. For example, the graph property of being bipartite is expressed by the formula **Bipartite** defined as

$$\begin{aligned} \mathbf{Bipartite} \equiv & \exists V_1, V_2 : (\forall v : (v \in V_1 \vee v \in V_2) \wedge \neg(v \in V_1 \wedge v \in V_2)) \\ & \wedge (\forall v_1, v_2 : \text{Adj}(v_1, v_2) \\ & \Rightarrow ((v_1 \in V_1 \wedge v_2 \in V_2) \vee (v_1 \in V_2 \wedge v_2 \in V_1))), \end{aligned}$$

where v , v_1 , and v_2 are vertex variables, while V_1 and V_2 are vertex set variables. We usually denote a monadic second-order formula by φ . We follow the notation by Satzinger [21], which allows us to use abbreviations such as $\exists v_1, v_2 : \varphi(v_1, v_2)$ instead of $\exists v_1 \exists v_2 \varphi(v_1, v_2)$, and $\nexists v \in V : \varphi(v)$ instead of $\nexists v (v \in V \wedge \varphi(v))$. Although each variable in φ has an associated type, we will omit to provide these explicitly whenever they are clear from the context.

Note that whenever we need to decide whether a given graph G satisfies formula φ , we also need to provide an interpretation of every free variable in φ ; this way we can use certain vertices or edges in G as constants. For instance, the following formula expresses whether a fixed vertex s is contained in a triangle of G :

$$\mathbf{Triangle}_s \equiv \exists v_1, v_2 : \text{Adj}(s, v_1) \wedge \text{Adj}(s, v_2) \wedge \text{Adj}(v_1, v_2)$$

For further definitions and examples, see Lovász [17] and Cygan et al. [10].

Courcelle proved in [9] that every graph property definable in monadic second-order logic can be decided in linear time on graphs that have bounded

treewidth. This metatheorem is not enough to solve the SHORTEST ODD PATH problem on graphs of bounded treewidth, but the extension of monadic second-order logic developed by Arnborg, Lagergren and Seese [2] enables us to deal with edge weights. Arnborg et al. defined *linear EMS extremum problems*, which include problems where the input graph $G = (V, E)$ is equipped with an evaluation function f , and the task is to find the maximum (or the minimum) of the evaluation term $f(P) := \sum_{e \in P} f(e)$ over all sets $P \subseteq E$ such that G satisfies a given monadic-second order formula $\varphi(P)$ with P as a free edge set variable in φ .¹

Theorem 20 ([2]). *For each integer k , every linear EMS extremum problem can be solved in linear time, assuming that a tree-decomposition of width at most k is provided for the input graph.*

6.2.2. FPT algorithm with parameter treewidth

We are now ready to prove the following result.

Theorem 21. *For each fixed integer k , the SHORTEST ODD PATH problem can be solved in linear time on graphs with treewidth at most k .*

Theorem 21 can be proved by using Theorem 20. For this, we need the following key lemma.

Lemma 22. *There exists a monadic second-order formula $\text{OddPath}_{s,t}(P)$ defined on a graph G containing two fixed vertices s and t which holds for a set P of edges in G if and only if P is an odd (s, t) -path in G .*

Proof. To define $\text{OddPath}_{s,t}(P)$, we first define a series of simple formulas of monadic second-order logic on our graph $G = (V, E)$. Recall that there is an atomic symbol $\text{Inc}(e, v)$ which holds exactly if $e \in E$, $v \in V$, and e is incident to v in G .

We start with formulas $\text{Deg0}(F, v)$ and $\text{Deg1}(F, v)$ which hold for an edge set $F \subseteq E$ and a vertex $v \in V$ if and only if v has degree 0 and 1, respectively, in $G[F]$:

$$\text{Deg0}(F, v) \equiv \nexists e : e \in F \wedge \text{Inc}(e, v),$$

$$\text{Deg1}(F, v) \equiv \exists e : e \in F \wedge \text{Inc}(e, v) \wedge (\forall f : (f \in F \wedge \text{Inc}(f, v)) \Rightarrow (e = f)).$$

¹We remark that such problems constitute only a small subset of linear EMS extremum problems; for the precise definition, see the paper by Arnborg et al. [2] or the introductory description by Satzinger [21, Section 3.3].

Using these, we can easily construct a formula $\text{Matching}(M)$ which holds if and only if the edge set M is a matching in G :

$$\text{Matching}(M) \equiv \forall v : \text{Deg0}(M, v) \vee \text{Deg1}(M, v).$$

We next define formula $\text{PartitionE}(F, F_1, F_2)$ which holds for edge sets F , F_1 , and F_2 if and only if (F_1, F_2) is a partition of F :

$$\begin{aligned} \text{PartitionE}(F, F_1, F_2) \equiv & (\nexists e \in E : (e \in F_1 \wedge e \in F_2)) \\ & \wedge (\forall e \in E : (e \in F_1 \vee e \in F_2) \Leftrightarrow (e \in F)). \end{aligned}$$

We proceed by defining the formula $\text{ConnectedE}(S)$ which holds if and only if the subgraph $G[F]$ is connected for a given set F of edges:

$$\begin{aligned} \text{ConnectedE}(F) \\ \equiv \forall F_1, F_2 : & (\text{PartitionE}(F, F_1, F_2) \wedge (\exists e_1 \in F_1) \wedge (\exists e_2 \in F_2)) \\ \Rightarrow & (\exists v, e_1, e_2 : e_1 \in F_1 \wedge e_2 \in F_2 \wedge \text{Inc}(e_1, v) \wedge \text{Inc}(e_2, v)). \end{aligned}$$

We are now ready to define $\text{OddPath}_{s,t}(P)$ as follows:

$$\begin{aligned} \text{OddPath}_{s,t}(P) \equiv & \text{ConnectedE}(P) \\ & \wedge (\exists P_1, P_2 : \text{PartitionE}(P, P_1, P_2) \wedge \text{Matching}(P_1) \\ & \wedge \text{Matching}(P_2) \wedge \text{Deg1}(P_1, s) \wedge \text{Deg0}(P_2, s) \\ & \wedge \text{Deg1}(P_1, t) \wedge \text{Deg0}(P_2, t)). \end{aligned}$$

To see the correctness of the formula, first assume that $S \subseteq E$ is an odd (s, t) -path in G . Let S_1 and S_2 be the set of edges in S with an odd and even sequence number, respectively. Then (i) (S_1, S_2) is a partition of S , (ii) S_1 and S_2 are both matchings, (iii) s and t both have degree 1 in S_1 and degree 0 in S_2 , and (iv) $G[S]$ is connected. Hence, G satisfies $\text{OddPath}_{s,t}(S)$.

Conversely, suppose that there exists an edge set P such that G satisfies $\text{OddPath}_{s,t}(P)$. By definition, this means that P is connected and can be partitioned into two matchings, P_1 and P_2 , satisfying also that s and t both have degree 1 in P_1 but have degree 0 in P_2 . Observe that a connected graph can only be partitioned into two matchings if and only if it is a path or a cycle. The degree constraints on s and t then ensure that P is a path with s and t as its endpoints, and moreover, it is an odd path (since P_1 is incident to both s and t). This proves the correctness of $\text{OddPath}_{s,t}(P)$. \square

Proof of Theorem 21. By Lemma 22, the SHORTEST ODD PATH problem can be expressed as a *linear EMS extremum problem* [2], using the monadic

second-order formula $\text{OddPath}_{s,t}(P)$ with the evaluation function being the edge weight function w , so that the evaluation term for P becomes $w(P)$. Therefore, Theorem 20 implies that for any class \mathcal{K} of graphs of bounded treewidth, the value

$$\min \{w(P) : P \subseteq E, \text{OddPath}_{s,t}(P)\}$$

can be computed in linear time for any graph $G = (V, E) \in \mathcal{K}$, edge weight function $w: E \rightarrow \mathbb{R}$, and vertices $s, t \in V$, assuming that a tree-decomposition of G is given as part of the input. Since it is possible to compute a tree-decomposition of width at most $2k + 1$ of any n -vertex graph with treewidth at most k in $2^{\mathcal{O}(k)}n$ time by Theorem 3, the theorem follows. \square

We remark that it is also possible to give a direct FPT-algorithm for SHORTEST ODD PATH parameterized by the treewidth of the input graph using dynamic programming techniques on a tree-decomposition. Although the naive method for solving SHORTEST ODD PATH on a graph with n vertices and a tree-decomposition of width k would yield an algorithm running in $k^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ time, it is possible to obtain an $c^k k^{\mathcal{O}(1)}n$ algorithm for some constant c using the rank-based approach of Bodlaender et al. [5] in almost the same way as they do for the TRAVELING SALESMAN problem; see Appendix A for the details.

7. Conclusion

We have identified islands of tractability for the computationally hard SHORTEST ODD PATH problem. We gave a polynomial-time algorithm for the case when the weight function is conservative and the set of negative-weight edges forms a tree. We developed an FPT algorithm for the problem when parameterized by the number of negative-weight edges, or more generally, by the size of a maximum matching in the graph spanned by all negative-weight edges, assuming again conservative weights. We also proved that SHORTEST ODD PATH is linear-time solvable on graphs of bounded treewidth, irrespective of the weight function.

A natural direction for further research is to identify additional cases when the problem becomes computationally tractable. A starting point could be to generalize our results to the case when the negative-weight edges form $c \geq 2$ trees. Is SHORTEST ODD PATH solvable in polynomial time when c is a constant? If so, is it fixed-parameter tractable with c as the parameter?

Acknowledgement. We would like to thank Kristóf Bérczi who provided a technical review of the manuscript and was part of the organizers of the 12th Emléktábla Workshop, Hungary, July 2022, where the collaboration of the authors was hosted and supported. We also would like to thank Jesper Nederlof for personal communication about their results in [5]. We are grateful to the reviewers for their helpful comments; in particular, the proof of Corollary 8 has been substantially simplified.

Funding. Csaba Király was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences; by the ÚNKP-21-5 New National Excellence Program of the Ministry for Innovation and Technology; and by the Hungarian Scientific Research Fund (OTKA grants FK128673 and PD138102). Lydia Mirabel Mendoza-Cadena was supported by the Lendület Programme of the Hungarian Academy of Sciences – grant number LP2021-1/2021; and by the Ministry of Innovation and Technology of Hungary – grant number ELTE TKP 2021-NKTA-62. Ildikó Schlotter is supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2021-2) and its János Bolyai Research Scholarship; and by the Hungarian Scientific Research Fund (OTKA grants K128611 and K124171). Yutaro Yamaguchi was supported by JSPS KAKENHI Grant Numbers 20K19743 and 20H00605; and by the Osaka University Research Abroad Program.

Declarations of interest. None.

References

- [1] E. M. Arkin, C. H. Papadimitriou, and M. Yannakakis. Modularity of cycles and paths in graphs. *Journal of the ACM*, 38(2):255–274, 1991.
- [2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991.
- [3] A. Björklund, T. Husfeldt, and P. Kaski. The shortest even cycle problem is tractable. In *STOC 2022: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, page 117–130, 2022.
- [4] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.
- [5] H. L. Bodlaender, M. Cygan, S. Kratsch, and J. Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation*, 243:86–111, 2015.

- [6] M. Chudnovsky, K. Kawarabayashi, and P. Seymour. Detecting even holes. *Journal of Graph Theory*, 48(2):85–111, 2005.
- [7] M. Chudnovsky, A. Scott, P. Seymour, and S. Spirkl. Detecting an odd hole. *Journal of the ACM*, 67(1):5:1–5:12, 2020.
- [8] M. Conforti, G. Cornuéjols, A. Kapoor, and K. Vušković. Even-hole-free graphs part ii: Recognition algorithm. *Journal of Graph Theory*, 40(4):238–266, 2002.
- [9] B. Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.
- [10] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized algorithms*. Springer, Cham, 2015.
- [11] U. Derigs. An efficient Dijkstra-like labeling method for computing shortest odd/even paths. *Information Processing Letters*, 21(5):253–258, 1985.
- [12] H. N. Gabow. Data structures for weighted matching and nearest common ancestors with linking. In *SODA '90: Proceedings of the 1st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 434–443, 1990.
- [13] M. Grötschel and W. R. Pulleyblank. Weakly bipartite graphs and the max-cut problem. *Operations Research Letters*, 1(1):23–27, 1981.
- [14] Y. Iwata and Y. Yamaguchi. Finding a shortest non-zero path in group-labeled graphs. *Combinatorica*, 42(Suppl 2), 2022.
- [15] T. Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *FOCS 2021: Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science*, pages 184–192, 2021.
- [16] A. S. Lapauwh and C. H. Papadimitriou. The even-path problem for graphs and digraphs. *Networks*, 14(4):507–513, 1984.
- [17] L. Lovász. *Large networks and graph limits*, volume 60 of *Colloquium publications*. American Mathematical Society, 2012.
- [18] W. McCuaig. Pólya’s permanent problem. *The Electronic Journal of Combinatorics*, 11(1):Research paper R79, 2004.

- [19] M. Naor, L. J. Schulman, and A. Srinivasan. Splitters and near-optimal derandomization. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 182–191, 1995.
- [20] N. Robertson, P. D. Seymour, and R. Thomas. Permanents, Pfaffian orientations, and even directed circuits. *Annals of Mathematics*, 150(3):929–975, 1999.
- [21] C. Satzinger. On tree-decompositions and their algorithmic implications for bounded-treewidth graphs. diplomathesis, Vienna University of Technology, 2014.
- [22] I. Schlotter. Shortest two disjoint paths in conservative graphs. In *STACS 2024: Proceedings of the 41st International Symposium on Theoretical Aspects of Computer Science*, volume 289 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 57:1–57:17, 2023.
- [23] I. Schlotter and A. Sebő. Odd paths, cycles and T -joins: Connections and algorithms, 2022. arXiv:2211.12862[cs.CC].
- [24] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, Berlin, 2003.
- [25] C. Thomassen. Even cycles in directed graphs. *European Journal of Combinatorics*, 6(1):85–89, 1985.
- [26] V. V. Vazirani and M. Yannakakis. Pfaffian orientations, 0–1 permanents, and even cycles in directed graphs. *Discrete Applied Mathematics*, 25(1):179–190, 1989.

Appendix A. Dynamic programming for bounded treewidth

We assume that the reader is familiar with the technique and terminology developed by Bodlaender et al. in their paper [5]. We show the necessary tools for solving the SHORTEST ODD PATH problem when the graph has treewidth k , in a similar fashion as the dynamic programming approach for the TRAVELING SALESMEN problem in [5]. Using the results in [5], one can show that the running time is $c^k k^{\mathcal{O}(1)} n$, for some constant c .

Let us consider a nice tree decomposition $(\mathbb{T}, \mathcal{B})$ as defined in [5]. Let r be the root vertex of \mathbb{T} ; we may assume that the root bag contains only vertices $\{s, t\}$, and we denote it by B_r . Given a node x of \mathbb{T} , we denote its bag by B_x , and by V_x the set of all vertices contained in the bag of some

descendant of x . The subgraph G_x is then defined as (V_x, E_x) where E_x is the set of edges introduced in some descendant of x .

Observe that given a solution to our instance (G, w, s, t) , that is, an odd (s, t) -path P , the edges of $P \cap E_x$ span a set of vertex-disjoint paths in the subgraph G_x . For the dynamic programming algorithm, we use $\mathbf{d}: B_x \rightarrow \{0, 1, 2\}$ to encode the degree of each vertex in B_x in such a partial solution $P \cap E_x$. Note also that $P \cap E_x$ naturally defines a perfect matching on those vertices of B_x that have degree 1 in $P \cap E_x$. We encode this information in a partition $M \in \Pi_2(\mathbf{d}^{-1}(1))$ where $\Pi_2(S)$ denotes the set of all partitions over a set S in which each block has size 2. Finally, we encode the parity of $|P \cap E_x|$ in an integer p .

For each node $x \in V(\mathbb{T})$, we can now define the set $\mathcal{E}_x(\cdot)$ of partial solutions we aim to capture and the table $A_x(\cdot)$ that we are going to compute at x . More precisely, given a function $\mathbf{d} \in \{0, 1, 2\}^{B_x}$, a partition $M \in \Pi_2(\mathbf{d}^{-1}(1))$ and integer $p \in \{0, 1\}$ we define

$$\begin{aligned} \mathcal{E}_x(M, \mathbf{d}, p) = & \left\{ F \subseteq E_x : \forall v : v \in B_x \Rightarrow \deg_F(v) = \mathbf{d}(v) \right. \\ & \wedge \forall v : v \in V_x \setminus B_x \Rightarrow \deg_F(v) \in \{0, 2\} \\ & \wedge G_x[F] \text{ contains no cycle} \\ & \wedge \forall u, v \in B_x : \{u, v\} \in M \\ & \quad \Rightarrow u \text{ and } v \text{ are connected in } G_x[F] \\ & \left. \wedge |F| \equiv p \pmod{2} \right\}; \\ A_x(\mathbf{d}, p) = & \left\{ \left(M, \min_{F \in \mathcal{E}_x(M, \mathbf{d}, p)} w(F) \right) : M \in \Pi_2(\mathbf{d}^{-1}(1)) \right. \\ & \left. \wedge \mathcal{E}_x(M, \mathbf{d}, p) \neq \emptyset \right\}. \end{aligned}$$

Observe that the weight of a minimum-weight odd (s, t) -path can be found by taking the weight stored in the table $A_r(\mathbf{d}, 1)$ for the partition $\{\{s, t\}\}$, where $\mathbf{d}(s) = \mathbf{d}(t) = 1$ and $\mathbf{d}(v) \in \{0, 2\}$ for each $v \in B_r \setminus \{s, t\}$. We consider the following recurrence for $A_x(\mathbf{d}, p)$, depending on the type of x . Note that we give the recurrence formulas using the operators over weighted partitions defined in [5]. Although Bodlaender et al. [5] use non-negative integer weights, their machinery also works for integer weights.

Leaf bag B_x . We have $B_x = \emptyset$. Then $A_x(\cdot)$ is empty for all possible values, except for

$$A_x(\emptyset, 0) = \{(\emptyset, 0)\}.$$

Introduce vertex v with bag B_x . Let y be the child of x , that is, $B_x = B_y \cup \{v\}$. For all $\mathbf{d} \in \{0, 1, 2\}^{B_x}$ and $p \in \{0, 1\}$ we set

$$A_x(\mathbf{d}, p) = \begin{cases} A_y(\mathbf{d}|_{B_y}, p) & \text{if } \mathbf{d}(v) = 0, \\ \emptyset & \text{otherwise,} \end{cases}$$

as v has just been introduced, and so it has no neighbours in G_x ; therefore, there is no change in the connectivity or in the parity of partial solutions.

Forget vertex v with bag B_x . Let y be the child of x , that is, $B_y = B_x \cup \{v\}$. For all $\mathbf{d} \in \{0, 1, 2\}^{B_x}$ and $p \in \{0, 1\}$ we set

$$A_x(\mathbf{d}, p) = A_y(\mathbf{d}[v \rightarrow 0], p) \uplus A_y(\mathbf{d}[v \rightarrow 2], p).$$

Either v is contained in an optimal solution P or not. If v is contained in P , then since v cannot have neighbours outside V_x , the two neighbours of v in P must be in V_x . Hence, the degree of v in a partial solution must be 0 or 2. Again, there is no change in parity or in connectivity; however, we need to discard dominated partitions.

Introduce edge $e = uv$ with bag B_x . Let y be the child of x . We have $B_x = B_y$. For all $\mathbf{d} \in \{0, 1, 2\}^{B_x}$ and $p \in \{0, 1\}$ we set

$$A_x(\mathbf{d}, p) = A_y(\mathbf{d}, p) \uplus \begin{cases} \emptyset & \mathbf{d}(u) = 0 \vee \mathbf{d}(v) = 0, \\ \text{glue}_w(uv, A_y(\mathbf{d}[u, v \rightarrow 0], 1 - p)) & \mathbf{d}(u) = 1 \wedge \mathbf{d}(v) = 1, \\ \text{proj}(\{v\}, \text{glue}_w(uv, A_y(\mathbf{d}[u \rightarrow 0, v \rightarrow 1], 1 - p))) & \mathbf{d}(u) = 1 \wedge \mathbf{d}(v) = 2, \\ \text{proj}(\{u\}, \text{glue}_w(uv, A_y(\mathbf{d}[u \rightarrow 1, v \rightarrow 0], 1 - p))) & \mathbf{d}(u) = 2 \wedge \mathbf{d}(v) = 1, \\ \text{proj}(\{u, v\}, \text{glue}_w(uv, A_y(\mathbf{d}[u, v \rightarrow 1], 1 - p))) & \mathbf{d}(u) = 2 \wedge \mathbf{d}(v) = 2. \end{cases}$$

Clearly, a partial solution may not use the edge e , and we keep track of such partial solutions by keeping all undominated entries from the table $A_y(\mathbf{d}, p)$ computed for the child y . If e is contained in a partial solution F for x compatible with \mathbf{d} and p , then the degree of u in $F \cap E_y = F \setminus \{e\}$ is smaller by 1 than its degrees in F , and the same holds for v . In particular, $\mathbf{d}(u) = 0$ or $\mathbf{d}(v) = 0$ is not possible. Furthermore, F and $F \cap E_y$ have different parities. The simple case is when both $\mathbf{d}(u) = 1$ and $\mathbf{d}(v) = 1$: then we can safely add the edge e and use the parity that differs from p ; note that this means inserting the block $\{u, v\}$ into the partition, which we can achieve by using the glue_w operation. If $\mathbf{d}(u) = 2$ or $\mathbf{d}(v) = 2$, then we first need to glue the partitions containing u and v , and then we need

to use the **proj** operation to remove those vertices from the partition that have degree 2 in F but degree 1 in $F \cap E_y$ (such vertices can be u , v , or both). Note that the **proj**($\{u, v\}, \cdot$) operation filters out those partitions where projecting out u and v decreases the number of blocks; hence, entries of $A_y(\mathbf{d}, p)$ with partitions containing $\{u, v\}$ will not be used in the case $\mathbf{d}(u) = \mathbf{d}(v) = 2$. This way we avoid creating cycles.

Join with bag B_x . Let y and z be the children of x , that is, $B_x = B_y = B_z$; let us now denote the “left” degree vector for child y by \mathbf{l} and the “right” degree vector for child z by \mathbf{r} . For all $\mathbf{d} \in \{0, 1, 2\}^{B_x}$ and $p \in \{0, 1\}$ we set

$$A_x(\mathbf{d}, p) = \bigsqcup_{\substack{\mathbf{l} + \mathbf{r} = \mathbf{d}, \\ p_1 + p_2 \equiv p \pmod{2}}} \text{proj}(\mathbf{d}^{-1}(2) \setminus (\mathbf{l}^{-1}(2) \cup \mathbf{r}^{-1}(2)), \text{join}(A_y(\mathbf{l}, p_1), A_z(\mathbf{r}, p_2))).$$

Observe that $\mathbf{l} + \mathbf{r} = \mathbf{d}$ is a vector summation as $\mathbf{l}, \mathbf{r}, \mathbf{d} \in \{0, 1, 2\}^{B_x}$. Combining $(M_1, w_1) \in A_y(\mathbf{l}, p_1)$ and $(M_2, w_2) \in A_z(\mathbf{r}, p_2)$ to get an entry for $A_x(\mathbf{d}, p)$ can be done if and only if $M_1 \cup M_2$ is acyclic and we have the correct parity. Maintaining acyclicity is equivalent to asking that the vertices in $\mathbf{d}^{-1}(2) \setminus (\mathbf{l}^{-1}(2) \cup \mathbf{r}^{-1}(2))$, which are the vertices that have degree 1 in the partial solutions for y and for z but degree 2 in the partial solution for x , are connected to vertices in $\mathbf{d}^{-1}(1)$ in the resulting partition $M_1 \sqcup M_2$. We use the projection operation for ensuring this.

Using the above formulas and applying the framework of [5], it is straightforward to obtain an algorithm for SHORTEST ODD PATH that runs in $c^k k^{\mathcal{O}(1)} n$ time for some constant c on graphs with treewidth at most k , where n denotes the number of vertices.