Fixed-Parameter Tractability of Maximum Colored Path and Beyond*

Fedor V. Fomin† Petr A. Golovach† Tuukka Korhonen† Kirill Simonov‡
Giannos Stamoulis§

Abstract

We introduce a general method for obtaining fixed-parameter algorithms for problems about finding paths in undirected graphs, where the length of the path could be unbounded in the parameter. The first application of our method is a randomized algorithm, that given a colored n-vertex undirected graph, vertices s and t, and an integer k, finds an (s,t)-path containing at least k different colors in time $2^k n^{\mathcal{O}(1)}$. This is the first FPT algorithm for this problem, and it generalizes the algorithm of Björklund, Husfeldt, and Taslaman [SODA 2012] on finding a path through k specified vertices. It also implies the first $2^k n^{\mathcal{O}(1)}$ time algorithm for finding an (s,t)-path of length at least k.

Our method yields FPT algorithms for even more general problems. For example, we consider the problem where the input consists of an n-vertex undirected graph G, a matroid M whose elements correspond to the vertices of G and which is represented over a finite field of order q, a positive integer weight function on the vertices of G, two sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, and the task is to find p vertex-disjoint paths from S to T so that the union of the vertices of these paths contains an independent set of M of cardinality k and weight w, while minimizing the sum of the lengths of the paths. We give a $2^{p+\mathcal{O}(k^2 \log(q+k))} n^{\mathcal{O}(1)} w$ time randomized algorithm for this problem.

*The research leading to these results has received funding from the Research Council of Norway via the project BWCA (grant no. 314528). Kirill Simonov acknowledges support by the Austrian Science Fund (FWF, project Y1329). Giannos Stamoulis acknowledges support by the ANR project ESIGMA (ANR-17-CE23-0010) and the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027).

†Department of Informatics, University of Bergen, Norway.

‡Algorithms and Complexity Group, TU Wien, Austria.

§LIRMM, Université de Montpellier, CNRS, France.

1 Introduction

The study of long cycles and paths in graphs is a popular research direction in parameterized algorithms. Starting from the color-coding of Alon, Yuster, and Zwick [1], powerful algorithmic techniques have been developed [2, 3, 15, 16, 21, 29, 46, 49], see also [14, Chapter 10], for finding long cycles and paths in graphs. However, most of the known methods are applicable only in the scenario when the size of the solution is bounded by the parameter. Let us explain what we mean by that by the following example.

Consider two very related problems, k-Cycle and Longest Cycle. In both problems, we are given a graph[1] G and an integer parameter k. In k-Cycle we ask whether G has a cycle of length *exactly* k. In Longest Cycle, we ask whether G contains a cycle of length *at least* k. While in the first problem any solution should have exactly k vertices, in the second problem the solution could be even a Hamiltonian cycle on n vertices. The essential difference in applying color-coding (and other methods) to these problems is that for k-Cycle, a random coloring of the vertices of G in k colors will color the vertices of a solution cycle with different colors with probability e^{-k}. Such information about colorful solutions allows dynamic programming to solve k-Cycle (as well as the related k-Path problem, the problem of finding a path of length exactly k). However, since a solution cycle for Longest Cycle is not upper-bounded by a function of k, the coloring argument falls apart. As Fomin et al. write in [21] "This is why color-coding and other techniques applicable to k-Path do not seem to work here." Sometimes, like in the case of Longest Cycle, a simple "edge contraction" trick, see [14, Exercise 5.8], allows reducing the problem to k-Cycle. We are not aware of general methods for solving problems related to cycles and paths when the size of the solution is not upper-bounded by the parameter.

The main result of this paper is a theorem that allows deriving algorithms for various parameterized problems about paths, cycles, and beyond, in the scenario when the size of the solution is not upper-bounded by the parameter. We discuss numerous applications of the theorem in the next subsection.

Our theorem is about finding a *k-colored (S,T)-linkage* in a colored graph. Let G be a graph, S and T be sets of vertices of G, and p be a positive integer. An *(S,T)-linkage* of order p is a set $\mathcal{P}$ of $p=|\mathcal{P}|$ vertex-disjoint paths, each starting in S and ending in T. The set of vertices in the paths of $\mathcal{P}$ is denoted by $V(\mathcal{P})$. The *total length* (or often just the length) of an (S,T)-linkage is the total number of vertices in its paths, i.e., $|V(\mathcal{P})|$. For a coloring $c:V(G)\to[n]$ of G, an (S,T)-linkage $\mathcal{P}$ is called *k-colored* if $V(\mathcal{P})$ contains at least k different colors, i.e., $|c(V(\mathcal{P}))|\geq k$. Let us note that in the above definition the sets S and T are not necessarily disjoint and that the coloring c is not necessarily a proper coloring in the graph-coloring sense. We also note that for vertices $s,t\in V(G)$, an $(\{s\},\{t\})$-linkage of order 1 corresponds to an (s,t)-path.

Theorem 1. *There is a randomized algorithm, that given as an input an n-vertex graph G, a coloring $c:V(G)\to[n]$ of G, two sets of vertices $S,T\subseteq V(G)$, and integers p,k, in time $2^{k+p}n^{\mathcal{O}(1)}$ either returns a k-colored (S,T)-linkage of order p and of the minimum total length, or determines that G has no k-colored (S,T)-linkage of order p.*

Few remarks are in order. First, Theorem 1 cannot be extended to directed graphs. It is easy to show, see Proposition 1, that finding a 2-colored (s,t)-path in a 2-colored directed graph is already NP-hard. Second, by another simple reduction, see Proposition 2, it can also be observed that if the time complexity of Theorem 1 could be improved to $(2-\varepsilon)^{k+p}n^{\mathcal{O}(1)}$ for $\varepsilon>0$, even in the case when $p=1$, G is colored with k colors, and $S=T=V(G)$, then Set Cover

[1] In this paper, graphs are assumed to be undirected if it is not explicitly mentioned to be otherwise.

would admit a $(2-\varepsilon)^n(mn)^{\mathcal{O}(1)}$ time algorithm, contradicting the Set Cover Conjecture (SeCoCo) of Cygan et al. [13]. We also remark that actually we prove an even more general result than Theorem 1, our result in full generality will be stated as Theorem 4. It can be also observed that by a simple reduction that subdivides edges, the coloring could be on the edges of G instead of vertices (or on both vertices and edges).

The algorithm in Theorem 1 invokes DeMillo-Lipton-Schwartz-Zippel lemma for polynomial identity testing and thus is "heavily" randomized. We do not know whether Theorem 1 could be derandomized. The special case of Theorem 1 when the coloring is a bijection, the problem of finding an (S,T)-linkage of order p and of length at least k, can be reduced to the (rooted) topological minor containment. To see why, observe that if we enumerate all possible collections $\mathcal{P}$ of p paths of total length k, then we can check for each collection $\mathcal{P}$ if it is contained as a rooted topological minor in G. The topological minor containment admits a deterministic FPT algorithm parameterized by the size of the pattern graph [24]. However the running time of the algorithm of Grohe et al. [24] is bounded by a tower of exponents in k and p. Our next theorem gives a deterministic algorithm for computing an (S,T)-linkage of order p and of length at least k whose running time is single-exponential in the parameter k for any fixed value of p. The other advantage of the algorithm in Theorem 2 is that it works on directed graphs too. In the following statement, a *directed* (S,T)-linkage is defined analogously to an (S,T)-linkage, but is composed of directed paths from S to T.

Theorem 2. *There is a deterministic algorithm that, given an n-vertex digraph G, two sets of vertices $S,T\subseteq V(G)$, an integer p, and an integer k, in time $p^{\mathcal{O}(kp)}n^{\mathcal{O}(1)}$ either returns a directed (S,T)-linkage of order p and of total length at least k, or determines that G has no directed (S,T)-linkage of order p and total length at least k.*

1.1 Applications of Theorem 1

Theorem 1 implies FPT algorithms for several problems. It encompasses a number of fixed-parameter-tractability results and improves the running times for several fundamental well-studied problems.

Longest path/cycle. When the coloring $c:V(G)\to[n]$ is a bijection, and thus all vertices of G are colored in different colors, then an (S,T)-linkage is k-colored if and only if its length is at least k. In this case, Theorem 1 outputs an (S,T)-linkage of order p with at least k vertices in time $2^{k+p}n^{\mathcal{O}(1)}$. In particular, for $p=1$ it implies that Longest (s,t)-Path (i.e., for $s,t\in V(G)$ and $k\geq 0$, to decide whether there is an (s,t)-path of length at least k) is solvable in time $2^k n^{\mathcal{O}(1)}$. Since one can solve Longest Cycle (to decide whether G contains a cycle of length at least k) by solving for every edge $st\in E(G)$ the Longest (s,t)-Path problem, Theorem 1 also yields an algorithm solving Longest Cycle in time $2^k n^{\mathcal{O}(1)}$. To the best of our knowledge, the previous best known algorithm for Longest (s,t)-Path runs in time $4.884^k n^{\mathcal{O}(1)}$ [22] and the previous best known algorithm for Longest Cycle runs in time $1.66^{2k}n^{\mathcal{O}(1)}=2.76^k n^{\mathcal{O}(1)}$ [3, 49]. The latter algorithm follows by combining the result of Zehavi [49] stating that Longest Cycle is solvable in time $t(G,2k)n^{\mathcal{O}(1)}$, where $t(G,k)$ is the best known running time for solving k-Path, with the fastest algorithm for k-Path of Björklund et al. [3].

For $p=2$, the problem of finding an (S,T)-linkage of length at least k is equivalent to the problem of finding a cycle of length at least k passing through a given pair of vertices s,t. A randomized algorithm of running time $(2e)^k n^{\mathcal{O}(1)}$ for this problem, known as Longest (s,t)-Cycle, was given by Fomin et al. in [19, Theorem 4] (see also [20]).

As we already have mentioned the problem of finding an (S,T)-linkage of order p and of length at least k can be reduced to the (rooted) topological minor containment. For $p \geq 3$, Theorems 1 and 2 provide the *first* (randomized and deterministic) single-exponential in $k+p$ and single-exponential in k for constant p, respectively, algorithms for computing an (S,T)-linkage of order p and of length at least k. For directed graphs, Theorem 2 gives the first FPT algorithm for the problem parameterized by $k+p$.

T-cycle. In the T-Cycle problem, we are given a graph G and a set $T \subseteq V(G)$ of terminals. The task is to decide whether there is a cycle passing through all terminals [4, 28, 45]. By the celebrated result of Björklund, Husfeldt, and Taslaman [4], T-Cycle is solvable in time $2^{|T|}n^{\mathcal{O}(1)}$, and their algorithm in fact returns the shortest such cycle. To solve T-Cycle as an application of Theorem 1, we do the following. We pick a terminal vertex $t \in T$, create a twin vertex s of t (i.e., a vertex s with $N(s) = N(t)$), and color s and t with color 1. We then color all non-terminal vertices of G with color 1 too. The remaining terminal vertices $T \setminus \{t\}$ we color in $|T|-1$ colors from 2 to $|T|$, such that no color repeats twice. Then G has a T-cycle if and only if there is a $|T|$-colored $(\{s\},\{t\})$-linkage of order 1. Therefore, using the algorithm of Theorem 1, we can also find the shortest T-cycle in time $2^{|T|}n^{\mathcal{O}(1)}$. One could use Theorem 1 to generalize the algorithmic result of Björklund, Husfeldt, and Taslaman in different settings. For example, instead of a cycle passing through all terminal vertices, we can ask for a cycle containing at least k terminals from a set T of unbounded size, in time $2^k n^{\mathcal{O}(1)}$.

Another generalization of T-Cycle comes from covering terminal vertices by at most p disjoint cycles. For example, in the basic VRP (vehicle routing problem) one wants to route p vehicles, one route per vehicle, starting and finishing at the depot so that all the customers are supplied with their demands and the total travel cost is minimized [10]. In the simplified situation when the clients are viewed as terminal vertices T of a graph and routes in VRP are required to be disjoint, this problem turns into the problem of finding a "p-flower" of minimum total length containing all vertices of T. By p-flower we mean a family of p cycles that intersect only in one (depot) vertex s. To see this problem as a problem of finding a colored (S,T)-linkage, we replace the depot s by a set S of $2p$ vertices whose neighbors are identical to the neighbors of s. Then similar to T-Cycle, this variant of VRP reduces to computing a minimum length $(|T|+1)$-colored (S,S)-linkage of order p; thus it is solvable in time $2^{|T|+p}n^{\mathcal{O}(1)}$ by Theorem 1.

Colored paths and cycles. The problems of finding a path, cycle, or another specific subgraph in a colored graph with the maximum or the minimum number of different colors appear in different subfields of algorithms, graph theory, optimization, and operations research [6, 8, 11, 12, 25, 31, 32, 42, 47]. In particular, the seminal color-coding technique of Alon, Yuster, and Zwick [1], builds on an algorithm finding a colorful path in a k-colored graph, that is, a path of k vertices and k colors, in time $\mathcal{O}(2^k n)$.

In the Maximum Colored (s,t)-Path problem, we are given a graph G with a coloring $c : V(G) \to [n]$ and integer k. The task is to identify whether G contains a k-colored (s,t)-path, i.e., an (s,t)-path with at least k different colors. In the literature, this problem is also known as Maximum Labeled Path [12] and Maximum Tropical Path [11]. Theorem 1 yields the *first* FPT algorithm for Maximum Colored (s,t)-Path, as well as for Maximum Colored Cycle (decide whether G contains a k-colored cycle). It is also the first FPT algorithm for the even more restricted variant of deciding if a given k-colored graph contains any k-colored path. A recent paper of Cohen et al. [11] claims a $\mathcal{O}(2^k n^2)$ time deterministic algorithm for computing a shortest k-colored path in a given k-colored graph. Unfortunately, a closer inspection of the algorithm of Cohen et al. reveals that it computes a k-colored walk instead of a k-colored path.[2]

[2]The error in [11] occurs on p. 478. It is claimed that if P is a shortest (u,v)-path that uses the set C of colors

It is interesting to note that the minimization version of the colored (s,t)-path, i.e., to decide whether there is an (s,t)-path containing at most k different colors, is W[1]-hard even on very restricted classes of graphs [18].

Beyond graphs: frameworks. Frameworks provide a natural generalization of colored graphs. Following Lovász [35], we say that a pair (G,M), where G is a graph and $M=(V(G),\mathcal{I})$ is a matroid on the vertex set of G, is a *framework*. Then we seek for a path, cycle, or (S,T)-linkage in G maximizing the rank function of M. Note that frameworks (G,M) where M is a partition matroid generalize colored graphs. Indeed, the universe $V(G)$ of M is partitioned into color classes $L_1,\ldots,L_n$ and a set I is independent if $|I\cap L_i|\leq 1$ for every color $i\in[n]$. However, by plugging different types of matroids into the definition of the framework, we obtain problems that cannot be captured by colored graphs.

Frameworks, under the name *pregeometric graphs*, were used by Lovász in his influential work on representative families of linear matroids [34]. The problem of computing maximum matching in frameworks is strongly related to the matchoid, the matroid parity, and polymatroid matching problems. See the Matching Theory book of Lovász and Plummer [36] for an overview. In their book, Lovász and Plummer use the term *matroid graph* for frameworks. In his most recent monograph, [35], Lovász introduces the term frameworks, and this is the term we adopt in our work. More generally, the problems of computing specific subgraphs of large ranks in a framework, belong to the broad class of problems about submodular function optimization under combinatorial constraints [7, 9, 40].

Let (G,M) be a framework and let $r\colon 2^{V(G)}\to\mathbb{Z}_{\geq 0}$ be the rank function of the matroid M. The *rank of a subgraph* H of G is $r(V(H))$ and we denote it by $r(H)$. We say that an (S,T)-linkage $\mathcal{P}$ in a framework (G,M) is *k-ranked* if the rank of $\mathcal{P}$, that is the rank in M of the elements corresponding to the vertices of the paths of $\mathcal{P}$, is at least k. With additional work involving (lossy) randomized truncation of the matroid, it is possible to extend Theorem 1 from colored graphs to frameworks over a general class of representable matroids.

Theorem 3. *There is a randomized algorithm that, given a framework (G,M), where G is an n-vertex graph and M is represented as a matrix over a finite field of order q, sets of vertices $S,T\subseteq V(G)$, and an integer k, in time $2^{p+\mathcal{O}(k^2\log(q+k))}n^{\mathcal{O}(1)}$ either finds a k-ranked (S,T)-linkage of order p and of minimum total length, or determines that (G,M) has no k-ranked (S,T)-linkage of order p.*

With minor adjustments, Theorem 3 can be adapted for frameworks with matroids that are in general not representable over a finite field of small order. For example, uniform matroids, and more generally transversal matroids, are representable over a finite field, but the field of representation must be large enough. Despite this, we can apply Theorem 3 to transversal matroids. Similarly, it is possible to apply Theorem 3 in the situation when M is represented by an integer matrix over rationals with entries bounded by $n^{\mathcal{O}(k)}$.

Weighted extensions. Theorem 1 can be extended into a weighted version in two different settings. The first setting is to have weights on edges that affect the length of the (S,T)-linkage. It is easy to see that by subdividing edges, coloring the subdivision vertices with a new "dummy color", and increasing k by one, all our algorithms work in the setting when the edges have polynomially-bounded positive integer weights.

The second weighted extension is more interesting. It is to have weights on vertices, and asking for an (S,T)-linkage containing a combination of weights and colors in a specific sense. In this

and P' is a (w,t)-sub-path of P using colors $C'\subseteq C$, then P' must be a shortest (w,t)-path among all (w,t)-paths using colors C'. This claim is correct for walks but not for paths.

setting, we have in addition to the coloring $c : V(G) \to [n]$ a weight function $\texttt{we} : V(G) \to \mathbb{Z}_{\geq 1}$. For integers k, w, we say that an (S,T)-linkage $\mathcal{P}$ is (k,w)-colored if its vertices $V(\mathcal{P})$ contain a set $X \subseteq V(\mathcal{P})$ so that $|X| = k$, all vertices of X have different colors, and the total weight of X is exactly $\texttt{we}(X) = \sum_{v \in X} \texttt{we}(v) = w$. This weighted version does not follow by direct reductions, but instead by a modification of Theorem 1 (in our main proof, we will directly prove Theorem 4 instead of Theorem 1).

Theorem 4. *There is a randomized algorithm that, given as an input an n-vertex graph G, a coloring $c : V(G) \to [n]$ of G, a weight function $\texttt{we} : V(G) \to \mathbb{Z}_{\geq 1}$, two sets of vertices S, T, and three integers p, k, w, in time $2^{k+p} n^{\mathcal{O}(1)} w$ either returns a (k,w)-colored (S,T)-linkage of order p and of minimum total length, or determines that no (k,w)-colored (S,T)-linkages of order p exist.*

Note that Theorem 4 implies Theorem 1 by setting all vertex weights to 1 and $w = k$. Theorem 4 allows to derive some applications of our technique that do not directly follow from Theorem 1, which we proceed to describe.

Longest T-cycle. Recall that in the T-Cycle problem the task is to find a cycle passing through a given set T of terminal vertices. Both the algorithm of Björklund, Husfeldt, and Taslaman [4], and the application of the algorithm of Theorem 1 find in fact the shortest T-cycle. A natural generalization of the T-Cycle problem is the Longest T-Cycle problem, where in addition to the set T we are given an integer k and the task is to find a cycle of length at least k passing through the terminals T. Theorem 4 can be used to solve Longest T-Cycle in time $2^{\max(|T|,k)} n^{\mathcal{O}(1)}$ as follows. First, if $|T| \geq k$, any T-cycle has length at least k and we just use the algorithm for T-Cycle. Otherwise, like in the reduction for T-Cycle, we first pick a terminal $t \in T$ and create a twin s of it. Then, we color s and t with color 1, and all the other vertices with different colors from 2 to n. We also assign weight 3 to the terminal vertices T, weight 1 to the vertex s, and weight 2 to all other vertices. We invoke Theorem 4 to find an $(\{s\}, \{t\})$-linkage of order 1 that contains a set X of vertices with distinct colors, size $|X| = k$, and weight $\texttt{we}(X) = 2k + |T|$. Any such set X must be a superset of T and not contain s, and therefore the found path must correspond to a cycle of length at least k passing through the terminals T.

Vehicle routing with profits. With Theorem 4, we can give an algorithm for the vehicle routing problem in a bit more general setting. In particular, we consider the situation where the depot has k parcels, p vehicles, and for each vertex v we know that we obtain a profit $\texttt{we}(v)$ for delivering a parcel to that vertex. We can use Theorem 4 with the same reduction as used for VRP earlier, but instead letting the coloring of the vertices to be a bijection, to obtain a $2^{k+p} n^{\mathcal{O}(1)} w$ time algorithm for determining the shortest routing by cycles intersecting only at the depot that yields a total profit of w.

Longest k-colored (S,T)-linkage. Theorem 4 can be also used to derive a longest path version of Theorem 1, in particular an algorithm that given a graph G, a coloring $c : V(G) \to [n]$, two sets of vertices $S, T \subseteq V(G)$, three integers k, p, ℓ, in time $2^{p+\ell+k} n^{\mathcal{O}(1)}$ outputs a k-colored (S,T)-linkage of order p and length at least ℓ. The reduction is as follows. First, if $p \geq \ell$, then any (S,T)-linkage of order p has length at least ℓ, so we can use Theorem 1. Otherwise, we are looking for a k-colored (S,T)-linkage that contains at least $\ell - p$ edges. We subdivide every edge, and for each created subdivision vertex we assign a new color and weight $2k$. For the original vertices we keep their colors and assign weight 1. Now, any k-colored (S,T)-linkage of order p and length at least ℓ corresponds to an (S,T)-linkage of order p that contains a set X of vertices with distinct colors, size $|X| = k + \ell - p$, and weight exactly $\texttt{we}(X) = (\ell - p) \cdot 2k + k$ (note that here we use the property that we are looking for an exact weight instead of maximum weight).

Weighted frameworks. We consider a generalization of frameworks into weighted frameworks. In particular, we say that a triple $(G, M, \mathtt{we})$, where G is a graph, $M = (V(G), \mathcal{I})$ is a matroid, and $\mathtt{we} : V(G) \to \mathbb{Z}_{\geq 1}$ is a weight function, is a *weighted framework*. Now we can say that an (S, T)-linkage $\mathcal{P}$ in a weighted framework $(G, M, \mathtt{we})$ is (k, w)-*ranked* if $V(\mathcal{P})$ contains a set X of vertices with $X \in \mathcal{I}$, size $|X| = k$, and weight $\mathtt{we}(X) = w$. By using the same reduction as from Theorem 1 to Theorem 3, we obtain the following theorem.

Theorem 5. *There is a randomized algorithm that given a weighted framework $(G, M, \mathtt{we})$, where G is an n-vertex graph and M is represented as a matrix over a finite field of order q, sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, in time $2^{p+\mathcal{O}(k^2 \log(q+k))} n^{\mathcal{O}(1)} w$ either finds a (k, w)-ranked (S, T)-linkage of order p and of minimum total length, or determines that $(G, M, \mathtt{we})$ has no (k, w)-ranked (S, T)-linkages of order p.*

Note that Theorem 5 implies Theorem 3 by setting all vertex weights to 1 and $w = k$.

Finally, we remark that even though the correctness argument of our algorithm is technical, the algorithm itself is simple and practical, consisting of only simple dynamic programming over walks in the graph. In particular, the observed practicality of the algorithm of Björklund, Husfeldt, and Taslaman [4] for T-Cycle on graphs with thousands of vertices holds also for our algorithm.

Organization of the paper. The rest of the paper is organized as follows. In Section 2 we overview our techniques and outline our algorithms. In Section 3 we recall definitions and preliminary results. In Section 4 we prove the main result, i.e., Theorem 4 (recall that Theorem 4 implies Theorem 1). In Section 5 we give the extensions of our results from colored graphs to frameworks, i.e., Theorem 5. In Section 6 we prove Theorem 2. Finally, we conclude in Section 7.

2 Techniques and outline

The techniques behind Theorem 1 build on the idea of exploiting cancellation of monomials, a fundamental tool in the area [2, 3, 4, 5, 29, 30, 33, 42, 46]. In particular, we build on the cycle-reversal-based cancellation for T-Cycle introduced by Björklund, Husfeldt, and Taslaman [4], and on the bijective labeling-based cancellation introduced by Björklund [2] (see also [3]). The algorithm of Theorem 2 builds on color-coding [1], generalizing ideas that appeared in [19] for finding an (s, t)-cycle of length at least k.

In Section 2.1 we explain the new ideas of the techniques behind Theorem 1 in comparison to the earlier works, in Section 2.2 we give a more detailed outline of the proof of Theorem 1, and in Section 2.3 we give an outline of the proof of Theorem 2.

2.1 New techniques for Theorem 1

Let us first focus on the single path case of Theorem 1, i.e., $p = |S| = |T| = 1$, corresponding to the question of finding a k-colored (s, t)-path. Our algorithm is analogous to the algorithm of Björklund, Husfeldt, and Taslaman [4] for T-Cycle, but instead of having the "interesting set" of vertices T fixed in advance, our algorithm can choose any interesting set $X \subseteq V(G)$ of vertices of size $|X| = k$ included in the path "on the fly" in the dynamic programming over the walks. In particular, our dynamic programming over walks can choose whether it gives a label to a vertex or not. This is the crucial difference to the earlier works where there would be a set of vertices $Y \subseteq V(G)$ fixed in advance so that a vertex of Y would always be given a label if encountered in the walk and the vertices $V(G) \setminus Y$ would never be given labels [2, 3, 4, 42]. This would impose a limitation that because these algorithms work in time exponential in the number of labels used

(i.e. 2^k, where k is the number of labels), the intersection of the found path with the set Y would have to be bounded in the parameter. This explains why the previous techniques could not yield an FPT-algorithm for Maximum Colored Path, as no such suitable set Y can be fixed in advance.

Our on the fly labeling of vertices allows our algorithm to find paths that visit the same color multiple times, while still making sure that at least k different colors are visited. In particular, the interesting set $X \subseteq V(G)$ of vertices in the path that we want to label is any set of size k that contains k different colors. While our dynamic programming is still a straightforward dynamic programming over walks, the main difficulty over previous works is the argument that if no solution exists, then the polynomial that we compute is zero, i.e., all unwanted walks cancel out.

First, the argument of cancellation in the case when two vertices of the same color are given a label is a now-standard application of the bijective labeling based cancellation of Björklund [2]. Therefore, our main focus is on a cancellation argument for walks where k vertices of different colors have been labeled. Here, our starting point is the cycle reversal based cancellation argument for T-Cycle [4], but in our case significantly more arguments are needed. In particular, the main difference to earlier works caused by the introduction of the on the fly labeling is that a vertex can occur in a walk as both labeled and unlabeled. Very much oversimplified, this case is handled by a new label-swap cancellation argument, where a label is moved from a labeled occurrence of a vertex into an unlabeled occurrence of the vertex. While in isolation this argument is simple, it causes significant complications when combining with the cycle reversal based cancellation, in particular because of the "no labeled digons" property we have to impose to the labeled walk. However, we manage to combine these two arguments into a one very technical cancellation argument.

Then, let us move from one (s,t)-path to an (S,T)-linkage. This generalization of using cancellation of monomials to find multiple paths is foreshadowed by an algorithm for minimum cost flow by Lingas and Persson [33]. However, their arguments are considerably simpler due to not having labels on the walks.

To find (S,T)-linkages, we use a similar dynamic programming to the one path case, extending the set of walks from S to T one walk at a time. Here, we must introduce a new cancellation argument for the case when two different walks intersect. This argument is again simple in isolation: take the intersection point of the two intersecting walks and swap the suffixes of them starting from this point. First, to make sure that this operation does anything we need to make sure that the suffixes are not equal. We do this by enforcing that the ending vertices of the walks are different already in the dynamic programming, which adds the extra 2^p factor to the time complexity. The second complication is that again, this suffix swap operation does not play well together with the other cancellation arguments, and we need to again significantly increase the complexity of the combination of the three cancellation arguments. In the end, we have to consider 18 different cases in our cancellation argument, see Definition 3.

The extension from Theorem 1 to the weighted setting of Theorem 4 is a simple modification of the dynamic programming so that also the weight of the labeled vertices X is stored. Interestingly, this argument could be extended to look for paths containing a set of vertices X with any property of X that could be efficiently evaluated in dynamic programming.

2.2 Outline of Theorem 1

We first give the outline of the algorithm for the single path case of finding a k-colored (s,t)-path, and then discuss the generalization to (S,T)-linkage.

Superficially, our approach follows the one of Björklund, Husfeldt, and Taslaman [4] developed for the T-cycle problem. Similar to Björklund et al., for every length $\ell \geq 1$, we define a certain family of walks $\mathcal{C}_\ell$ and a polynomial $f(\mathcal{C}_\ell)$ so that over a finite field of characteristic 2, the polynomial

$f(\mathcal{C}_\ell)$ is non-zero if the graph contains a k-colored (s,t)-path of length ℓ, and the polynomial $f(\mathcal{C}_\ell)$ is the identically zero polynomial if the graph does not contain any k-colored (s,t)-path of length $\leq \ell$. Then by making use of the DeMillo-Lipton-Schwartz-Zippel lemma [44, 50], finding a k-colored (s,t)-path of minimum length boils down to evaluating the polynomial $f(\mathcal{C}_\ell)$ at a random point for increasing values of ℓ.

With k-colored path, the role similar to the role of terminal vertices in T-CYCLE is played by a subset X of k vertices of the path with k different colors. However, a priori we do not know this set X, and there could be n^k candidates so we cannot enumerate them. Because of that, we define the polynomial f on families of *labeled* (s,t)-walks in the graph G. A labeled (s,t)-walk of length ℓ is a pair of sequences $W = ((v_1,\ldots,v_\ell),(r_1,\ldots,r_\ell))$, where $v_1,\ldots,v_\ell$ is an (s,t)-walk of length ℓ, and $r_1,\ldots,r_\ell$ is a sequence of numbers from $[0,k]$ indicating a labeling. The interpretation of the labeling is that $r_i = 0$ indicates that the index i of the walk is not labeled, and $r_i \geq 1$ indicates that the index i is labeled with the label r_i, with the interpretation that the vertex v_i at this index is selected to the set X.

Next we present the definition of the polynomial f. The polynomial f is over $\mathrm{GF}(2^{3+\lceil \log_2 n \rceil})$, which is a field of characteristic 2 and order $\geq 8n$. With every edge $uv \in E(G)$ we associate a variable $f_e(uv)$, with every vertex $v \in V(G)$ we associate a variable $f_v(v)$, and with every color-label pair $(x,y) \in [n] \times [k]$ we associate a variable $f_c(x,y)$. For a labeled walk $W = ((v_1,\ldots,v_\ell),(r_1,\ldots,r_\ell))$ we associate the monomial

$$f(W) = \prod_{i=1}^{\ell-1} f_e(v_i v_{i+1}) \cdot \prod_{i \in [\ell] | r_i \neq 0} f_v(v_i) \cdot f_c(c(v_i), r_i).$$

For the family of walks $\mathcal{C}_\ell$, which we will define immediately, we are interested in the polynomial

$$f(\mathcal{C}_\ell) = \sum_{W \in \mathcal{C}_\ell} f(W).$$

For vertices s,t and integers k,ℓ, the family $\mathcal{C}_\ell$ is the family of all labeled (s,t)-walks

$$W = ((v_1 = s, v_2, \ldots, v_\ell = t),(r_1,\ldots,r_\ell))$$

of length ℓ that satisfy the following two properties. The first property is that the labeling $(r_1,\ldots,r_\ell)$ is bijective, meaning that every label from $[k]$ is used exactly once. Note that this implies that every monomial of $f(\mathcal{C}_\ell)$ has degree $\ell - 1 + 2k$, being a product of $\ell - 1$ edge variables, k vertex variables, and k color-label pair variables. The second property is that the labeled walk W has no *labeled digons*. By that we mean that W cannot have a subwalk $v_{i-1}v_iv_{i+1}$ with v_i being a labeled vertex (with $r_i \neq 0$) and $v_{i-1} = v_{i+1}$. It is not immediately clear that having no labeled digons is useful, but this will turn out to be crucial similarly to the property of having no T-digons in the algorithm for T-cycle [4].

It is not difficult to prove that when a graph has a k-colored (s,t)-path of length ℓ, then $f(\mathcal{C}_\ell)$ is a non-zero polynomial. Indeed, a path has no repeated vertices and thus has no labeled digons, so if we take a k-colored (s,t)-path $v_1,\ldots,v_\ell$ and let the labels $r_1,\ldots,r_\ell$ take the values from $[k]$ on k vertices with k different colors, then the labeled walk $W = ((v_1,\ldots,v_\ell),(r_1,\ldots,r_\ell))$ appears in $\mathcal{C}_\ell$, and thus a corresponding monomial $f(W)$ appears in $f(\mathcal{C}_\ell)$. Because $v_1,\ldots,v_\ell$ is a path and the labeled vertices have different colors, we can recover the labeled walk W uniquely from the monomial $f(W)$, and therefore the monomial $f(W)$ must occur exactly once in the polynomial $f(\mathcal{C}_\ell)$ (i.e. with coefficient 1), and therefore $f(\mathcal{C}_\ell)$ is non-zero.

The proof of the opposite statement—absence of a k-colored (s,t)-path of length $\leq \ell$ implies that $f(\mathcal{C}_\ell)$ is zero—is more complicated. We have to show that in this case each monomial $f(W)$ for

labeled walks $W \in \mathcal{C}_\ell$ occurs an even number of times in the polynomial $f(\mathcal{C}_\ell)$, in particular that there is an even number of labeled walks $W \in \mathcal{C}_\ell$ for every monomial $f(W)$. The proof is based on constructing an *f-invariant fixed-point-free involution* ϕ on $\mathcal{C}_\ell$, that is a function $\phi : \mathcal{C}_\ell \to \mathcal{C}_\ell$ such that for every $W \in \mathcal{C}_\ell$ it holds that $f(W) = f(\phi(W))$, $\phi(W) \neq W$, and $\phi(\phi(W)) = W$.

Let us start with the easy part of the proof, that is, constructing such ϕ for labeled walks where two vertices of the same color are labeled (which could be two different occurrences of the same vertex). In this case, let $1 \leq i < j \leq \ell$ be the lexicographically smallest pair of indices so that $c(v_i) = c(v_j)$, $r_i \neq 0$, and $r_j \neq 0$. The function ϕ works by swapping r_i with r_j. Because each label from $[k]$ occurs in $r_1, \ldots, r_\ell$ exactly once, this results in a different labeled walk $\phi(W)$ with the same monomial $f(\phi(W)) = f(W)$, and moreover $W = \phi(\phi(W))$ holds. After this argument, we can let $\mathcal{C}^*_\ell \subseteq \mathcal{C}_\ell$ be the family of labeled walks in $\mathcal{C}_\ell$ where all labeled vertices have different colors, and we know that $f(\mathcal{C}^*_\ell) = f(\mathcal{C}_\ell)$. Therefore, we can focus only to constructing $\phi : \mathcal{C}^*_\ell \to \mathcal{C}^*_\ell$.

Now, the first approach (which does not work) would be to adapt the strategy of Björklund, Husfeldt, and Taslaman for our purposes. The essence of their strategy is the following. Since walks from $\mathcal{C}^*_\ell$ do not have labeled digons and because there is no k-colored (s,t)-path of length $\leq \ell$, it is possible to show that every walk $W \in \mathcal{C}^*_\ell$ has a "loop", that is a subwalk vUv starting and ending in the same vertex v, and so that U is not a palindrome. Then $\phi(W)$ is the walk W' obtained from W by reversing U. This approach does not work directly in our case. The reason is that a labeled vertex could also occur several times in a walk as unlabeled. Because of that, reversing a subwalk can result in a walk with a labeled digon, and thus ϕ could map W outside the family $\mathcal{C}^*_\ell$. For example, for a walk $abcd\hat{a}b$ (here $\hat{a}$ is a labeled vertex), reversing $a\overleftarrow{bcd\hat{a}b}$ results in walk $adcb\hat{a}b$ with labeled digon $b\hat{a}b$. A natural "patch" for that type of walks is to not reverse but to apply a new type of operation of swapping a label from one occurrence of a vertex to another occurrence of it. For example, swapping a label for $abcd\hat{a}b$ would result in $\hat{a}bcdab$. This results in a different labeled walk contributing the same monomial $f(W)$ to the polynomial. See Figure 1 for an illustration of the above examples.

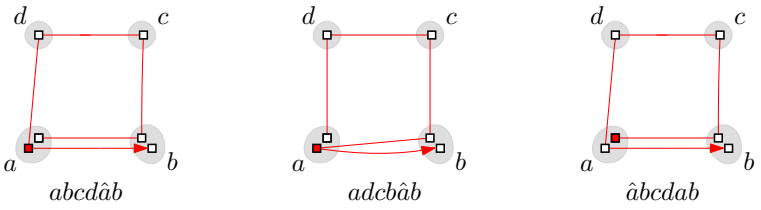

Figure 1: An illustration of the walks $abcd\hat{a}b$, $a\overleftarrow{bcd\hat{a}b} = adcb\hat{a}b$, and $\hat{a}bcdab$. The grey bags correspond to vertices of the graph. The squares are copies of the corresponding bag-vertex and these together with the red path illustrate the order the vertices appear in the walk. Red squares correspond to labeled vertices.

However, the new operation of swapping a label brings us new problems. First of all, swapping a label could again result in a labeled digon. For example, swapping a label for walk $\hat{a}bcac$ results in walk $abc\hat{a}c$ with labeled digon $c\hat{a}c$. An attempt to "patch" this by using a "mixed" strategy—when possible, swap a label, otherwise reverse—does not work either. For example, for walk $W = \hat{a}bcac$ we cannot label swap (that will result in a labeled digon $c\hat{a}c$), hence we reverse. Thus we obtain walk $W' = \phi(W) = \hat{a}\overleftarrow{bcac} = \hat{a}cbac$. For W', swapping a label for a is a valid operation, thus $\phi(W') = acb\hat{a}c$, but then we would have that $\phi(\phi(W)) \neq W$. See Figure 2 for an illustration of the above example.

At this moment, the situation becomes desperate: the more patches we introduce, the more

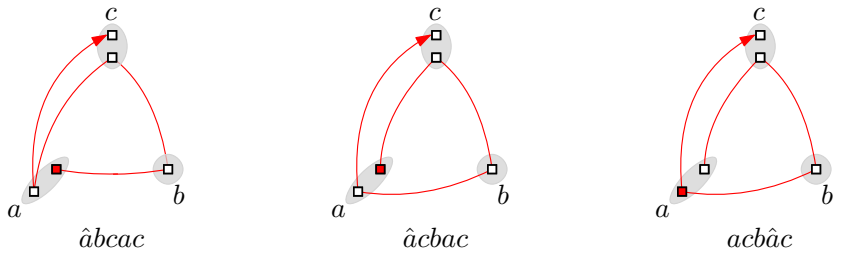

Figure 2: An illustration of the walks $\hat{a}bcac$, $\hat{a}\overleftarrow{bc}ac = \hat{a}cbac$, and $acb\hat{a}c$. The grey bags correspond to vertices of the graph. The squares are copies of the corresponding bag-vertex and these together with the red path illustrate the order the vertices appear in the walk. Red squares correspond to labeled vertices.

issues appear and the whole construction falls apart. Moreover, on top of that, one has to define the mapping ϕ recursively in order to deal with palindromic loops, making the situation even more complicated. We find it a bit surprising, that in the end a combination of label swaps and reverses allows the construction of the required mapping ϕ. To make it happen, we use quite an involved strategy to identify what labels can be swapped and what subwalks could be reversed, and in what order. Whole Section 4.4 is devoted to defining this strategy (for the more general setting of (S,T)-linkages) and to the proof of its correctness.

To evaluate the polynomial $f(\mathcal{C}_\ell)$, we apply quite standard dynamic programming techniques. In particular, the polynomial can be evaluated in $2^k n^{\mathcal{O}(1)}$ time by dynamic programming over walks, where we store the length of the walk, the last two vertices of the walk, the subset of labels used so far (causing the 2^k factor), and whether the last vertex is labeled. This is similar to the dynamic programming for T-CYCLE [4], with the difference only in that it is chosen in the dynamic programming which vertices of the walk are labeled, and that instead of a subset of T we store the subset of the labels.

To extend the algorithm from a single (s,t)-path to an (S,T)-linkage of order p, we define a family $\mathcal{C}_\ell$ of *labeled walkages* and a polynomial $f(\mathcal{C}_\ell)$ over them. We note that by a simple reduction we can assume that $|S| = |T| = p$, and that S and T are disjoint. A labeled walkage of order p and total length ℓ is a p-tuple $\mathcal{W} = (W^1, \ldots, W^p)$ of labeled walks W^i, whose sum of the lengths is ℓ. The family $\mathcal{C}_\ell$ contains labeled walkages $\mathcal{W}$ with the following properties: They have order p, total length ℓ, the starting vertices are ordered according to a total order on $V(G)$, ending vertices are distinct (each vertex in T is an ending vertex of exactly one walk in $\mathcal{W}$), the labeling is bijective (each label from $[k]$ is used exactly once), and no walk in $\mathcal{W}$ contains a labeled digon.

The monomial $f(\mathcal{W})$ is then defined as

$$f(\mathcal{W}) = \prod_{i=1}^{p} f(W^i),$$

and the polynomial $f(\mathcal{C}_\ell)$ as

$$f(\mathcal{C}_\ell) = \sum_{\mathcal{W} \in \mathcal{C}_\ell} f(\mathcal{W}).$$

The definitions are analogous to the single path case, in particular we recover the previously explained single path case by setting $p = 1$. The proof that if there exists a k-colored (S,T)-linkage of order p and total length ℓ then $f(\mathcal{C}_\ell)$ is non-zero is directly analogous to the one path case. Also the proof that we can consider the smaller family $\mathcal{C}^*_\ell \subseteq \mathcal{C}_\ell$ where all labeled vertices have different colors is analogous.

However, to prove that if there is no k-colored (S,T)-linkages of order p and total length $\leq \ell$ then $f(\mathcal{C}^*_\ell)$ is identically zero we need new cancellation arguments beyond the previous cycle reversal and label swap arguments. In particular, none of the previously considered arguments can be applied if we have a labeled walkage $\mathcal{W} = (W^1, W^2)$ of order two, where both W^1 and W^2 are labeled paths that intersect. In this case, the new argument is that we could swap the suffixes of W^1 and W^2 starting from the intersection point. For example, for a walkage $\mathcal{W} = (abc\hat{d}e, stcuv)$, we define $\phi(\mathcal{W}) = (abcuv, stc\hat{d}e)$. The property that the walks in $\mathcal{W}$ have different ending vertices is crucial here to ensure that $\phi(\mathcal{W}) \neq \mathcal{W}$.

However, also with this suffix swap cancellation argument we run into problems. In particular, the first challenge is that the suffix swap could create labeled digons, for example when $\mathcal{W} = (ab\hat{c}de, stcbu)$ both of the walks are paths, but swapping the suffix after c would create a labeled digon. In this situation we can instead use the label swap operation on c, from the first walk to the second, but of course this will add again even more complications. In the end, we manage to extend the strategy of ϕ from paths to linkages, but it makes the definition of ϕ even more complicated (see Definition 3, the path case uses the case groups A and C, while the linkage case needs the addition of case groups B and D).

The dynamic programming for (S,T)-linkage is similar to the (s,t)-path, extending the walks in the walkage one walk at the time. It requires two new fields to store, the index of the walk that we are currently extending, and the subset of the ending vertices T that have been already used. Storing the used ending vertices causes the additional factor 2^p in the time complexity (as we can assume that $|T| = p$).

2.3 Outline of Theorem 2

Recall that the main difference to Theorem 1 is that Theorem 2 provides a deterministic algorithm that, moreover, works on directed graphs. The price is, however, that this algorithm is only suitable for the special case of finding an (S,T)-linkage of length at least k, and the time complexity as a function of k and p is higher. Theorem 2 thus requires a completely different toolbox: the algorithm is based on ideas of random separation. Our result can be seen as a generalization of earlier works on finding paths and cycles of length at least k, the closest one being the result of Fomin et al. [19] on finding an (s,t)-cycle of length at least k. Note that their result is stated for undirected graphs, and that the problem of finding an (S,T)-linkage of order 2 and length at least k is equivalent to the problem of finding an (s,t)-cycle of length at least k (up to increasing k by 2) on undirected graphs.

Similarly to the earlier results, the case where the target (S,T)-linkage is of length close to k can be covered by a standard application of color-coding [1]. The difficulty is that the length of the (S,T)-linkage can be arbitrarily larger than k. While because of that it would be intractable to highlight the target (S,T)-linkage as a whole, it is still possible to apply random separation to give distinct colors to k-length segments at the end of each path in the (S,T)-linkage. The main hurdle is then to argue that at least in one color we can pick a finishing segment as an arbitrary shortest path of length k, without intersecting any other path in the optimal solution. Afterwards, finding the desired (S,T)-linkage is easy, as the length requirement is already satisfied; one only needs to find a suitable connection to complete the (S,T)-linkage, which exists as witnessed by the optimal solution.

Lemma 15 encapsulates the novel combinatorial result allowing the approach above, strongly generalizing a similar basic idea that appeared in [19] for two undirected paths. To give an intuition behind the lemma (see also Figure 3), observe first that the problem of finding an (S,T)-linkage of order p and of total length at least k is equivalent to the problem of finding an (s,t)-*linkage* of order

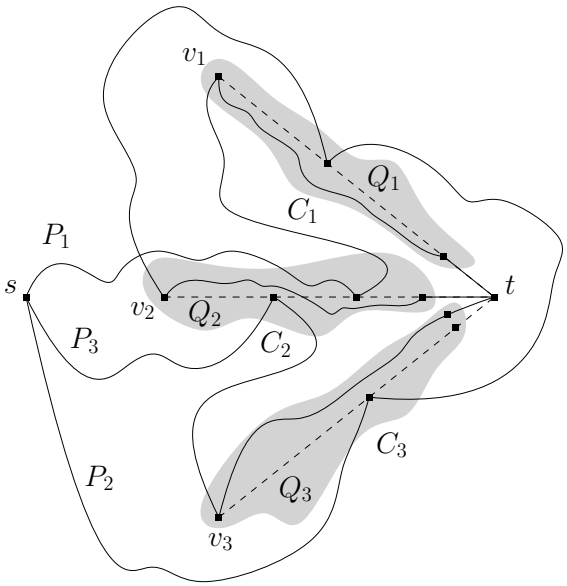

(a) The given digraph G, edge directions are implicit along the paths P_1–P_3, and Q_1–Q_3.

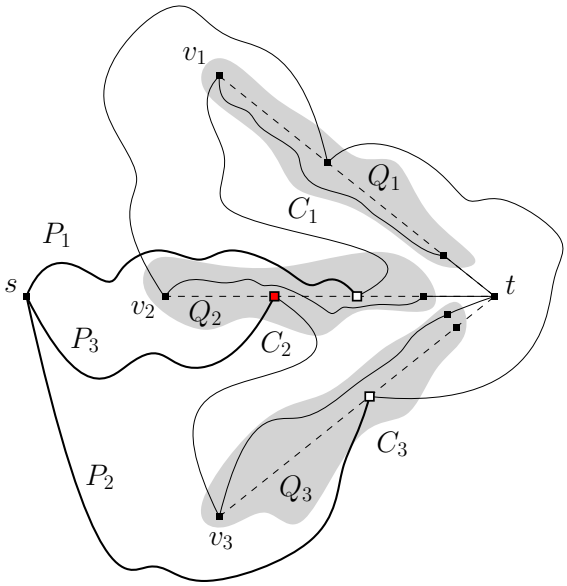

(b) Starting state S^1: two tokens belong to Q_2 so **Push** is applied to t_3^1.

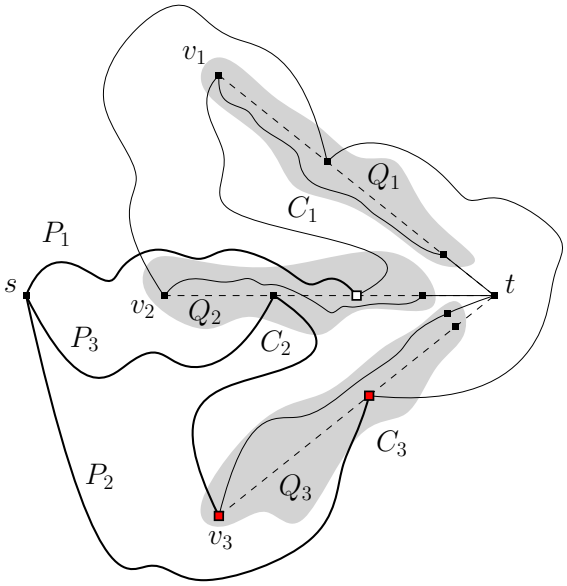

(c) In S^2, $t_3^2 = v_3$ so **Clear** is applied to Q_3, moving t_3^2 to t and t_2^2 next along P_2.

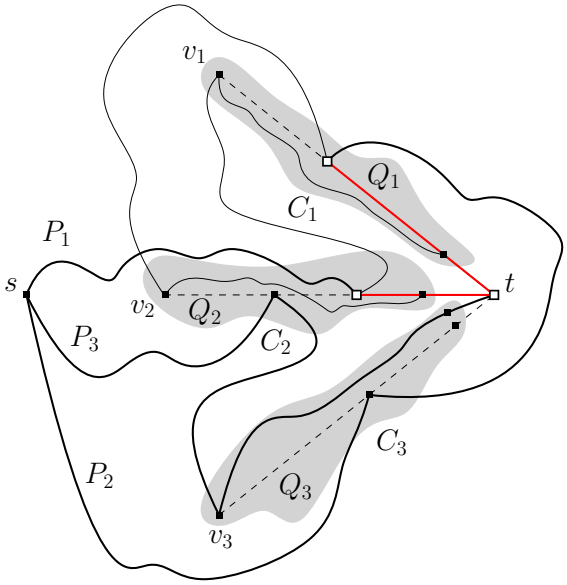

(d) Finishing state S^3: P_2 is continued via Q_1, P_1 via Q_2, and P_3 is preserved to obtain the solution.

Figure 3: Illustration to the proof of Lemma 16. Large empty squares mark tokens in the current state; those with red filling are moved by the next rule.

p and of total length at least $k+2$, where an (s,t)-linkage of order p consists of p internally-disjoint (s,t)-paths, for some $s,t \in V(G)$. Now let (s,t)-paths $P_1, \ldots, P_q$ come from the shortest solution, an (s,t)-linkage of order p and the smallest total length which is at least k. Let the sets $C_1, \ldots, C_q$ be the result of random separation applied to k-length suffixes of $P_1 - \{t\}, \ldots, P_q - \{t\}$, i.e., for each $i \in [q]$, the k-length suffix of $P_i - \{t\}$ is contained in C_i. The algorithm of Theorem 2 seeks to find a solution where for some $i \in [q]$, v_i is the k-th vertex of P_i from t, and Q_i is a k-length shortest path from v_i to t inside C_i, by guessing $v_i \in C_i$ and taking an arbitrary path Q_i of the

form above. The solution is then any collection of an (s, v_i)-path and $p-1$ (s,t)-paths that do not intersect each other and Q_i, together with Q_i. If Q_i does not intersect the (s, v_j)-prefix of P_j, for each $j \in [q]$, then the paths $P_1, \ldots, P_q$ certify that the desired collection of paths exists. Now comes Lemma 15: it claims, roughly, that if this is not the case for all $i \in [q]$, then there is a shorter (s,t)-linkage given by prefixes of $P_1, \ldots, P_q$ and suffixes of $Q_1, \ldots, Q_q$ (introducing another color to the random separation makes sure that the total length of the prefixes is still at least k), which is a contradiction.

The proof of Lemma 15 can be imagined as the following token sliding game. First, we put a token on each P_i, at the first place of intersection with some Q_j. Then we move the tokens by applying two rules, **Push** and **Clear**. If two tokens end up on the same Q_j for some $j \in [q]$, we move the farthest of them from t further along its path P_i, until it hits another $Q_{j'}$; this is called **Push**. As for the **Clear**, if at any step h the current token t_i^h of the path P_i reaches the vertex v_i, we forfeit this path: the token is moved to t, which corresponds to the i-th path of the shorter solution being exactly P_i, and all other tokens on Q_i are moved next along their paths similarly to the rule **Push**. Moreover, every future application of any rule will not place a token on Q_i, skipping it to the next Q_j that is still active. Clearly, this game is finite, as tokens are only being slid further along their paths. The main claim of Lemma 15 is that when the game is over, there is at least one remaining active token, these tokens are one per a path in $\{Q_j\}_{j\in[q]}$ (since **Push** is not applicable), and that all corresponding paths P_i can be simultaneously extended each along its own Q_j instead of taking their original routes, without intersections (since in **Push** we always keep the closest token to t). This is a shorter solution since a token of P_i, if active, is inside some Q_j at distance less than k from t, and the prefix of P_i up to this token is shorter than the prefix of P_i up to v_i.

Another challenge the proof of Theorem 2 faces, is that while the random separation approach is well-known, it is normally applied to separating two, rarely three (e.g. [19]), sets. We, on the other hand, need to apply random separation to p sets simultaneously, while making sure that it can be derandomized. To this end, in Lemma 16 we devise in a deterministic way a family of functions that models random separation of p sets of size at most k each. The size of this family is bounded by $p^{\mathcal{O}(kp)} \log n$, which matches the inverse probability (up to the $\log n$ factor) of coloring the universe in p colors uniformly at random so that each set receives its own color. The construction is based on perfect hash families [39].

3 Preliminaries

In this section, we introduce basic notation and state some auxiliary results.

3.1 Basic definitions and preliminary results

We use $\mathbb{Z}_{\geq 1}$ to denote the set of positive integers and $\mathbb{Z}_{\geq 0}$ the set of non-negative integers. Also, given integers p, q such that $p < q$, we use $[p, q]$ to denote the set $\{p, p+1, \ldots, q\}$ and, if $p \geq 1$, we use $[p]$ to denote the set $\{1, \ldots, p\}$.

Parameterized Complexity. We refer to the book of Cygan et al. [14] for introduction to the area. Here we only briefly mention the notions that are most important to state our results. A *parameterized problem* is a language $L \subseteq \Sigma^* \times \mathbb{N}$, where Σ^* is a set of strings over a finite alphabet Σ. An input of a parameterized problem is a pair (x, k), where x is a string over Σ and $k \in \mathbb{N}$ is a *parameter*. A parameterized problem is *fixed-parameter tractable* (or FPT) if it can be solved

in time $f(k) \cdot |x|^{\mathcal{O}(1)}$ for some computable function f. The complexity class FPT contains all fixed-parameter tractable parameterized problems.

Graphs. We use standard graph-theoretic terminology and refer to the textbook of Diestel [17] for missing notions. We consider only finite graphs, and the considered graphs are assumed to be undirected if it is not explicitly said to be otherwise. For a graph G, $V(G)$ and $E(G)$ are used to denote its vertex and edge sets, respectively. Throughout the paper we use $n = |V(G)| = |G|$ and $m = |E(G)|$ if this does not create confusion. For a graph G and a subset $X \subseteq V(G)$ of vertices, we write $G[X]$ to denote the subgraph of G induced by X. For a vertex v, we denote by $N_G(v)$ the *(open) neighborhood* of v, i.e., the set of vertices that are adjacent to v in G. For $X \subseteq V(G)$, $N_G(X) = \left(\bigcup_{v \in X} N_G(v)\right) \setminus X$. The *degree* of a vertex v is $d_G(v) = |N_G(v)|$. If G is a digraph, $N_G^+(v)$ denotes the *out-neighborhood* of v, i.e., the set of vertices that are adjacent to v in G via an arc from v, and $N_G^-(v)$ is the *in-neighborhood*, defined symmetrically for arcs going to v. We may omit subscripts if the considered graph is clear from a context.

A walk W of length ℓ in G is a sequence of vertices $v_1, v_2, \ldots, v_\ell$, where $v_i v_{i+1} \in E(G)$ for all $1 \leq i < \ell$. The vertices v_1 and v_ℓ are the *endpoints* of W and the vertices $v_2, \ldots, v_{\ell-1}$ are the *internal* vertices of W. A path is a walk where no vertex is repeated. For a path P with endpoints s and t, we say that P is an (s,t)-path. A cycle is a path with the additional property that $v_\ell v_1 \in E(G)$ and $\ell \geq 3$.

DeMillo-Lipton-Schwartz-Zippel lemma. Our strategy involves the use the DeMillo-Lipton-Schwartz-Zippel lemma for randomized polynomial identity testing.

Lemma 1 ([44, 50])**.** *Let $p(x_1, \ldots, x_n)$ be a non-zero polynomial of total degree d over a field $\mathbb{F}$, and let S be a subset of $\mathbb{F}$. If each x_i is independently assigned an uniformly random value from S, then $p(x_1, \ldots, x_n) = 0$ with probability at most $d/|S|$.*

3.2 Hardness results

We conclude this section by showing the NP-hardness of finding a k-colored (s,t)-path on directed graphs, for any $k \geq 2$, and the optimality of the time complexity of Theorem 1 assuming the Set Cover Conjecture (SeCoCo) of Cygan et al. [13].

We start with the hardness for directed graphs.

Proposition 1. *For any integers $k, \ell \geq 2$, it is* NP*-complete to decide, given a directed graph G, a coloring $c\colon V(G) \to [\ell]$, and two vertices s and t, whether G has a k-colored (s,t)-path.*

Proof. We show the claim for $k = \ell = 2$ as it is straightforward to generalize the proof for other values of k and ℓ. We reduce from the Disjoint Paths problem on directed graphs. The task of this problem is, given a (directed) graph G and k pairs of terminal vertices (s_i, t_i) for $i \in \{1, \ldots, k\}$, decide whether G has vertex-disjoint (s_i, t_i)-paths for $i \in \{1, \ldots, k\}$. This problem is well-known to be NP-complete on directed graphs even if $k = 2$ [23]. Consider an instance $(G, (s_1, t_1), (s_2, t_2))$ of Disjoint Paths, where G is a directed graph. We assume that the terminal vertices are pairwise distinct. We construct the directed graph G' from G by adding a vertex w and arcs (t_1, w) and (w, s_2). Note that G has vertex-disjoint (s_1, t_1) and (s_2, t_2)-paths if and only if G' has an (s_1, t_2)-path containing w. We define the coloring c by setting $c(w) = 1$ and defining $c(v) = 2$ for all $v \in V(G') \setminus \{w\}$. Clearly, G' has a 2-colored (s_1, t_2)-path if and only if G' has an (s_1, t_2)-path containing w. This immediately implies NP-hardness. □

Then, we show that Theorem 1 is optimal assuming the Set Cover Conjecture.

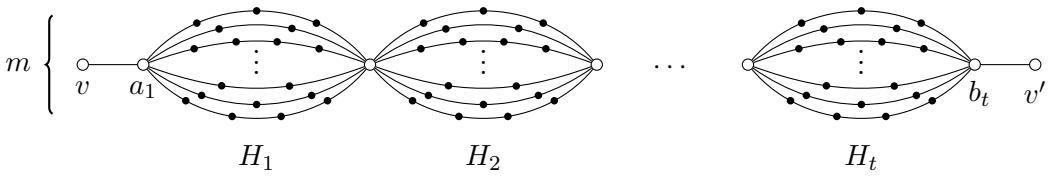

Figure 4: Construction of the graph G.

Proposition 2. *If there is a $(2-\varepsilon)^k n^{\mathcal{O}(1)}$ time algorithm for finding a k-colored path in a k-colored graph for some $\varepsilon > 0$, then there is a $(2-\varepsilon)^n (mn)^{\mathcal{O}(1)}$ time algorithm for* Set Cover.

Proof. In the Set Cover problem, we are given a universe U of n elements, a collection $\mathcal{S}$ of m subsets of U, and an integer t and we ask whether there is a collection $\mathcal{S}' \subseteq \mathcal{S}$ of size t such that for every $u \in U$, there is a set $S \in \mathcal{S}'$ such that $u \in S$.

Given an instance $(U, \mathcal{S}, t)$ of Set Cover, where $|U| = n$ and $\mathcal{S} = \{S_1, \ldots, S_m\}$, we construct a graph G as follows. We first construct the graph H by considering two vertices a and b and adding m internally vertex-disjoint (a,b)-paths $P_{S_1}, \ldots, P_{S_m}$, where for every $i \in [m]$, the vertices in P_{S_i} are bijectively mapped to the elements of S_i. We call a the *source* of H and b the *sink* of H. We finally construct a graph G that is obtained by considering t copies $H_1, \ldots, H_t$ of H, for each $i \in [t-1]$, identifying the sink b_i of H_i with the source a_{i+1} of H_{i+1}, and adding two new vertices v and v' of degree one, adjacent to a_1 and b_t respectively. See Figure 4 for an illustration of the construction of graph G. Note that $t \leq m$ and $|V(G)| = (mn)^{\mathcal{O}(1)}$.

Assuming an ordering $u_1, \ldots, u_n$ of U, for each $i \in [n]$, we assign color i to all vertices of G that correspond to u_i, color $n+1$ and $n+2$ to v and v', and color $n+3$ to all vertices in $V(G) \setminus \{v, v'\}$ that do not correspond to members of U. Observe that $(U, \mathcal{S}, t)$ is a yes-instance of Set Cover if and only if there is an $n+3$-colored path in G. Therefore, a $(2-\varepsilon)^k n^{\mathcal{O}(1)}$ time algorithm for finding a k-colored path in a k-colored n-vertex graph implies the existence of a $(2-\varepsilon)^n (mn)^{\mathcal{O}(1)}$ time algorithm for finding a set cover of size t in a universe U of size n with a collection $\mathcal{S}$ of m subsets of U. □

4 Randomized algorithm for colored (S,T)-linkages

In this section we prove the main result, i.e., Theorem 4. Recall that Theorem 1 is a special case of Theorem 4.

Let G be an n-vertex graph, p an integer, and $S, T \subseteq V(G)$. An (S,T)-linkage of order p is a set $\mathcal{P}$ of $p = |\mathcal{P}|$ vertex-disjoint paths between S and T. We denote by $V(\mathcal{P})$ the vertices in the paths of $\mathcal{P}$. The length of an (S,T)-linkage is the total number $|V(\mathcal{P})|$ of vertices in the paths. Let $c : V(G) \to [n]$ an arbitrary coloring of G, and $\mathtt{we} : V(G) \to \mathbb{Z}_{\geq 1}$ a weight function. For positive integers k and w, we say that an (S,T)-linkage $\mathcal{P}$ is (k,w)-colored if there exists a set $X \subseteq V(\mathcal{P})$ with $|X| = k$, all vertices of X have different colors, and $\mathtt{we}(X) = \sum_{v \in X} \mathtt{we}(v) = w$. We give a $2^{p+k} n^{\mathcal{O}(1)} w$ time algorithm for the problem of finding a minimum length (k,w)-colored (S,T)-linkage of order p (Theorem 4).

We will assume that $|S| = |T| = p$, and S and T are disjoint, as the general case can be reduced to this case by the following reduction: We add p vertices $s_1, \ldots, s_p$ with $N(s_i) = S$ and p vertices $t_1, \ldots, t_p$ with $N(t_i) = T$, all with the same new color and weight equal to $k \cdot \max_{v \in V(G)} \mathtt{we}(v) + 1$. Then, we can set $S = \{s_1, \ldots, s_p\}$ and $T = \{t_1, \ldots, t_p\}$, and solve the problem with k increased by one and w increased by $k \cdot \max_{v \in V(G)} \mathtt{we}(v) + 1$. Because we can assume that the original weights

are at most $w+1$, this increases the target weight w by a factor $\mathcal{O}(k)$, and therefore does not increase the time complexity of the algorithm.

4.1 Labeled walks and walkages

In this subsection we define labeled walks and labeled walkages.

Labeled walks. Let ℓ be an integer. A *walk of length* ℓ in G is a sequence of vertices $v_1, \ldots, v_\ell$ of G, where $v_i v_{i+1} \in E(G)$ for all $1 \leq i < \ell$. A *labeled walk of length* ℓ is a pair of sequences $W = ((v_1, v_2, \ldots, v_\ell), (r_1, r_2, \ldots, r_\ell))$, where $v_1, \ldots, v_\ell$ is a walk of length ℓ, and $r_1, \ldots, r_\ell$ is a sequence of integers from $[0, k]$, indicating a labeling. The interpretation of the labeling is that $r_i = 0$ indicates that the index i is unlabeled and $r_i \neq 0$ indicates that the index i is labeled with the label $r_i \in [k]$. A labeled walk is *injective* if each label from $[k]$ appears in it at most once. Most of the labeled walks that we treat in the algorithm have length at least one, but the definition allows also an empty labeled walk of length zero. The set of vertices *collected* by W is $R(W) = \{v_i \mid r_i \neq 0\}$, i.e., the set of vertices that occur at labeled indices. The set of edges of W is $E(W) = \{v_i v_{i+1} \mid 1 \leq i < \ell\}$. An index i in a labeled walk of length ℓ is a *digon* if $1 < i < \ell$ and $v_{i-1} = v_{i+1}$ (see Figure 5 for an illustration). An index i in a labeled walk is a *labeled digon* if it is a digon and $r_i \neq 0$.

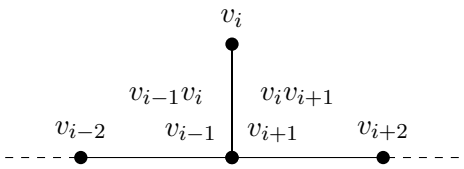

Figure 5: An example of a labeled walk W with a digon i.

Labeled walkages. A *labeled walkage of order* p is a tuple $\mathcal{W} = (W^1, \ldots, W^p)$, where each $W^i = ((v_1^i, \ldots, v_{\ell_i}^i), (r_1^i, \ldots, r_{\ell_i}^i))$ is a labeled walk of length $\ell_i \geq 1$. The *length* of $\mathcal{W}$ is $\sum_{i=1}^p \ell_i$. The *set of edges* of $\mathcal{W}$ is $E(\mathcal{W}) = \bigcup_{i=1}^p E(W^i)$. The set of vertices *collected* by $\mathcal{W}$ is $R(\mathcal{W}) = \bigcup_{i=1}^p R(W^i)$. The weight $\mathtt{we}(\mathcal{W})$ of $\mathcal{W}$ is the sum of the weights of the labeled vertices, i.e., $\mathtt{we}(\mathcal{W}) = \sum_{i=1}^p \sum_{j \in [\ell_i] | r_j^i \neq 0} \mathtt{we}(v_j^i)$. Note that the weight of a vertex can be counted more than once if the vertex occurs as labeled more than once. A labeled walkage is *injective* if each label from $[k]$ appears in it at most once, and *bijective* if each label from $[k]$ appears in it exactly once. Note that every labeled walk in an injective labeled walkage is injective.

The set of *ending* vertices of a labeled walkage $\mathcal{W}$ of order p is $\mathcal{T}(\mathcal{W}) = \{v_{\ell_i}^i \mid i \in [p]\}$. The tuple of *starting* vertices of $\mathcal{W}$ is $\mathtt{start}(\mathcal{W}) = (v_1^1, \ldots, v_1^p)$. Let $<$ be a total order on $V(G)$. A labeled walkage is *ordered* if $\mathtt{start}(\mathcal{W})$ is ordered according to $<$, i.e., $v_1^i < v_1^{i+1}$ holds for all $1 \leq i < p$. The asymmetry that the starting vertices are an ordered tuple while the ending vertices are an unordered set is essential for our algorithm. A *labeled linkage* is a labeled walkage where every vertex occurs at most once, i.e., the walks are vertex-disjoint paths.

We also define *semiproper* and *proper* labeled walkages. The intuition here is that, in Section 4.2, we define a polynomial over semiproper walkages (see also Definition 1). Then, walkages that are semiproper but not proper will be handled by using standard techniques and therefore we can focus on proper walkages. Dealing with proper walkages will be the most technical part of the proof. A labeled walkage is *semiproper* if it is injective, no walk in it contains labeled digons, and the ending vertices of the walkage are distinct, i.e., $v_{\ell_i}^i \neq v_{\ell_j}^j$ for $i \neq j$. A labeled walkage $\mathcal{W}$ is

proper if it is semiproper and all of its labeled indices correspond to vertices of different colors, i.e., $|c(R(\mathcal{W}))| = |\{(i,j) \mid r^i_j \neq 0\}|$. Note that being proper implies that no vertex is labeled twice, and note that if $\mathcal{W}$ is bijective and proper then $|c(R(\mathcal{W}))| = k$.

4.2 Algorithm

We assume that there is a total order $<$ on $V(G)$, and for a set $S \subseteq V(G)$ we denote by $\texttt{ordv}(S)$ the tuple containing the elements of S ordered according to $<$. Note that G contains a (k,w)-colored (S,T)-linkage of order p and length ℓ if and only if there is a bijective proper ordered labeled linkage $\mathcal{W}$ with order p, length ℓ, weight $\texttt{we}(\mathcal{W}) = w$, tuple of starting vertices $\texttt{start}(\mathcal{W}) = \texttt{ordv}(S)$, and set of ending vertices $\mathcal{T}(\mathcal{W}) = T$. We define a family of labeled walkages that includes all such labeled linkages, but relaxes the condition of being a linkage to walkage, and the condition of being proper to semiproper.

For each integer ℓ, we define a family of labeled walkages $\mathcal{C}_\ell$ of length ℓ.

Definition 1 (Family $\mathcal{C}_\ell$)**.** *Let ℓ a positive integer. The family $\mathcal{C}_\ell$ consists of the bijective semiproper ordered labeled walkages $\mathcal{W}$ with order p, length ℓ, weight $\texttt{we}(\mathcal{W}) = w$, tuple of starting vertices $\texttt{start}(\mathcal{W}) = \texttt{ordv}(S)$, and set of ending vertices $\mathcal{T}(\mathcal{W}) = T$.*

Definition of the polynomial. Let $q = 2^{3+\lceil \log_2 n \rceil}$ and keep in mind that $\mathrm{GF}(q)$ is a finite field of characteristic 2 and order $q \geq 8n$. Next, we define a polynomial over $\mathrm{GF}(q)$ that will be evaluated at a random point by our algorithm. For each edge $uv \in E(G)$ we associate a variable $f_e(uv)$, for each vertex $v \in V(G)$ we associate a variable $f_v(v)$, and for each color-label-pair $(x,y) \in [n] \times [k]$ we associate a variable $f_c(x,y)$. For a labeled walk $W = ((v_1,\ldots,v_\ell),(r_1,\ldots,r_\ell))$ we associate the monomial

$$f(W) = \prod_{i=1}^{\ell-1} f_e(v_i v_{i+1}) \cdot \prod_{i \in [\ell] | r_i \neq 0} f_v(v_i) \cdot f_c(c(v_i), r_i).$$

For a labeled walkage $\mathcal{W} = (W^1,\ldots,W^p)$ we associate the monomial

$$f(\mathcal{W}) = \prod_{i=1}^{p} f(W^i).$$

For a family $\mathcal{F}$ of labeled walkages we associate the polynomial

$$f(\mathcal{F}) = \sum_{\mathcal{W} \in \mathcal{F}} f(\mathcal{W}).$$

In particular, because the walkages in $\mathcal{C}_\ell$ are bijective, every monomial in the polynomial $f(\mathcal{C}_\ell)$ has degree $\ell - p + 2k$, being a product of $\ell - p$ variables corresponding to the edges of the walkage, k variables corresponding to the labeled vertices, and k variables corresponding to the color-label-pairs.

Algorithm for finding a (k,w)-colored (S,T)-linkage. Our algorithm for finding a (k,w)-colored (S,T)-linkage of order p works as follows. Starting with $\ell = p$, we evaluate the polynomial $f(\mathcal{C}_\ell)$ at a random point x over $\mathrm{GF}(q)$, for increasing values of ℓ. If $f(\mathcal{C}_\ell)(x) \neq 0$, we return that G contains a (k,w)-colored (S,T)-linkage of order p, and moreover that the shortest (k,w)-colored (S,T)-linkage of order p has length ℓ. Otherwise, we continue increasing ℓ until $\ell = n+1$ in which case we return that G does not contain a (k,w)-colored (S,T)-linkage of order p.

For the proof of correctness of the algorithm, in Section 4.3 we show that with probability at least 1/2 this algorithm returns the length of the shortest (k,w)-colored (S,T)-linkage of order p, and never returns a length shorter than the shortest (k,w)-colored (S,T)-linkage of order p.

Proof of time complexity of the algorithm. Next we prove the time complexity of the algorithm. The evaluation of the polynomial is done using dynamic programming. This is a standard application of dynamic programming over walks while keeping track of the set of labels used so far, the weight of the labeled vertices, and the set of ending vertices used. We prove that it can be performed in time $2^{p+k}n^{\mathcal{O}(1)}w$.

Lemma 2. *Let S,T be disjoint subsets of $V(G)$ of size $|S|=|T|=p$, $c:V(G)\to[n]$ a coloring of G, $\mathtt{we}:V(G)\to\mathbb{Z}_{\geq 1}$ a weight function, $\ell\leq n$ an integer, k,w integers, and $q=2^{3+\lceil\log_2 n\rceil}$. The polynomial $f(\mathcal{C}_\ell)$ can be evaluated at a random point over GF(q) in time $2^{p+k}n^{\mathcal{O}(1)}w$.*

Proof. We associate random values over $\mathrm{GF}(q)$ to all variables $f_v(v)$, $f_e(uv)$, and $f_c(x,y)$, and from now denote by $f_v(v)$, $f_e(uv)$, and $f_c(x,y)$ these associated values, and by extension for a walkage $\mathcal{W}$ denote by $f(\mathcal{W})$ the value associated to the monomial $f(\mathcal{W})$ and for a family of walkages $\mathcal{F}$ denote by $f(\mathcal{F})$ the value associated to the polynomial $f(\mathcal{F})$. Now, the task is to compute $f(\mathcal{C}_\ell)$.

Informally, we will compute $f(\mathcal{C}_\ell)$ by dynamic programming over partial walkages, growing the walkages one labeled walk at a time in the order specified by $\mathtt{ordv}(S)$.

Denote $\mathtt{ordv}(S)=(s_1,s_2,\ldots,s_p)$ and for any $t\in[p]$ denote by $\mathtt{pre}_t(S)$ the length-t prefix of $\mathtt{ordv}(S)$. For every integer $t\in[p]$, integer $l\in[\ell]$, set $L\subseteq[k]$ of labels, set $T'\subseteq T$ of ending vertices, weight $w'\in[0,w]$, vertices $x,y\in V(G)$, and integer $o\in\{0,1\}$, we define

$$D(t,l,L,T',w',x,y,o)=f(\mathcal{F}(t,l,L,T',w',x,y,o)),$$

where we define $\mathcal{F}(t,l,L,T',w',x,y,o)$ to be the family of labeled walkages $\mathcal{W}=(W^1,\ldots,W_t)$, where for each $i\in[t]$, $W^i=((v^i_1,\ldots,v^i_{\ell_i}),(r^i_1,\ldots,r^i_{\ell_i}))$, that satisfy the following properties:

1. Each labeled walk W^i in $\mathcal{W}$ has length at least 2 and does not contain labeled digons,
2. $\mathcal{W}$ has order t and ordered tuple of starting vertices $\mathtt{start}(\mathcal{W})=\mathtt{pre}_t(S)$,
3. $\mathcal{W}$ has length l,
4. $\mathcal{W}$ is injective and the set of used labels is L,
5. the set of ending vertices of all but the last walk in $\mathcal{W}$ is $\mathcal{T}((W^1,\ldots,W^{t-1}))=T'$,
6. $\mathcal{W}$ has weight $\mathtt{we}(\mathcal{W})=w'$,
7. the last vertex of the last walk in $\mathcal{W}$ is $v^t_{\ell_t}=x$,
8. the second last vertex of the last walk in $\mathcal{W}$ is $v^t_{\ell_t-1}=y$, and
9. if $o=0$, then $r_{t,\ell_t}=0$, otherwise $r_{t,\ell_t}\neq 0$.

In other words, t specifies the number of walks, l specifies the length, L specifies the used labels, T' specifies the used ending vertices, w' specifies the weight, x specifies the last vertex of the last walk, y specifies the second last vertex of the last walk, and o specifies whether the last vertex of the last walk is labeled. Note that it can be without loss of generality assumed that each walk has length at least 2 because S and T are disjoint.

Then, we define also a shorthand that for $t \in [p]$, $l \in [\ell]$, $L \subseteq [k]$, $T' \subseteq T$, and $w' \in [0, w]$,

$$D(t, l, L, T', w') = \sum_{x \in T'} \sum_{y \in N(x)} \sum_{o \in \{0,1\}} D(t, l, L, T' \setminus \{x\}, w', x, y, o),$$

which intuitively denotes the polynomial corresponding to a "completed" walkage of t walks with length l, used labels L, used ending vertices T', and weight w'.

Now it holds that

$$f(\mathcal{C}_\ell) = D(p, \ell, [k], T, w),$$

and therefore computing $f(\mathcal{C}_\ell)$ can be done by computing all of the values $D(t, l, L, T', w', x, y, o)$ by dynamic programming.

Next we specify this computation by dynamic programming. All values that we do not specify here will be set to zero. First, to initialize, we define a special value $D(0, 0, \emptyset, \emptyset, 0) = 1$ corresponding to a family of walkages containing one empty walkage.

Next, we describe computing the states where $o = 0$, i.e., the last vertex is not labeled, for all $t \in [p]$, $l \in [\ell]$, $L \subseteq [k]$, $T' \subseteq T$, $w' \in [0, w]$, $x \in V(G)$, and $y \in N(x)$, assuming that all the states with smaller l have already been computed. There are four cases, corresponding to the four lines of Equation (1). In the first case the walk W_t has length at least three, its second last vertex y is not labeled, and we are extending the walkage by adding one not labeled vertex x to W_t. Second case is the same, but the second last vertex y is labeled and thus we have to ensure to not create a labeled digon. Third case is the case that we are extending the walkage by adding one more labeled walk, consisting of two vertices y, x, where $y = s_t$, neither of them labeled. Fourth case is like the third, but the first vertex $y = s_t$ of the new walk is labeled. Recall the notation that $[y = s_t] = 1$ if $y = s_t$ holds, and 0 otherwise.

$$\begin{aligned}
&D(t, l, L, T', w', x, y, 0) = f_e(xy) \\
&\cdot \Bigg(\sum_{z \in V(G)} D(t, l-1, L, T', w', y, z, 0) \\
&+ \sum_{z \in V(G) \setminus \{x\}} D(t, l-1, L, T', w', y, z, 1) \\
&+ [y = s_t] \cdot D(t-1, l-2, L, T', w') \\
&+ [y = s_t] \cdot \sum_{r \in L} f_v(y) \cdot f_c(c(y), r) \cdot D(t-1, l-2, L \setminus \{r\}, T', w' - \mathtt{we}(y)) \Bigg).
\end{aligned} \tag{1}$$

Then, we describe computing the states where $o = 1$, i.e., the last vertex is labeled, for all $t \in [p]$, $l \in [\ell]$, $L \subseteq [k]$, $T' \subseteq T$, $w' \in [0, w]$, $x \in V(G)$, and $y \in N(x)$, assuming that all of the states with smaller l have already been computed. There are again four cases, analogously to Equation (1).

$$D(t,l,L,T',w',x,y,1) = \sum_{r\in L} f_v(x)\cdot f_c(c(x),r)\cdot f_e(xy)$$
$$\cdot\Bigg(\sum_{z\in V(G)} D(t,l-1,L\setminus\{r\},T',w'-\mathtt{we}(x),y,z,0)$$
$$+\sum_{z\in V(G)\setminus\{x\}} D(t,l-1,L\setminus\{r\},T',w'-\mathtt{we}(x),y,z,1) \tag{2}$$
$$+[y=s_t]\cdot D(t-1,l-2,L\setminus\{r\},T',w'-\mathtt{we}(x))$$
$$+[y=s_t]\cdot\sum_{r'\in L\setminus\{r\}} f_v(y)\cdot f_c(c(y),r')\cdot D(t-1,l-2,L\setminus\{r,r'\},T',w'-\mathtt{we}(x)-\mathtt{we}(y))\Bigg).$$

This completes the description of the dynamic programming, which shows that each of the states $D(t,l,L,T',w',x,y,o)$ can be computed in $n^{\mathcal{O}(1)}$ time given the values of the states with smaller l. As there are $p\cdot\ell\cdot 2^k\cdot 2^p\cdot(w+1)\cdot n\cdot n\cdot 2=\mathcal{O}(p2^{p+k}n^3w)$ states, the algorithm works in time $2^{p+k}n^{\mathcal{O}(1)}w$. □

As the algorithm can be implemented by $\mathcal{O}(n)$ applications of Lemma 2, the algorithm has time complexity $2^{p+k}n^{\mathcal{O}(1)}w$. Recovering the solution can be done by a factor of $\mathcal{O}(n^2)$ more applications.

4.3 Correctness

To prove the correctness of the algorithm, we show that

(a) the polynomial $f(\mathcal{C}_\ell)$ is non-zero if G contains a (k,w)-colored (S,T)-linkage of order p and length ℓ and

(b) the polynomial $f(\mathcal{C}_\ell)$ is the identically zero polynomial if the graph does not contain a (k,w)-colored (S,T)-linkage of order p and length $\leq\ell$.

Because $f(\mathcal{C}_\ell)$ has degree $\ell-p+2k\leq 3n\leq q/2$, it follows from Lemma 1 and (a) that if G contains a (k,w)-colored (S,T)-linkage of order p and length ℓ, then evaluating $f(\mathcal{C}_\ell)$ at a random point of $\mathrm{GF}(q)$ has probability at least $1/2$ to be non-zero. From (b) it follows that if G does not contain a (k,w)-colored (S,T)-linkage of order p and length $\leq\ell$, then evaluating $f(\mathcal{C}_\ell)$ at a random point is guaranteed to be zero. This establishes that the algorithm is correct with probability at least $1/2$, with one-sided error.

The part (a) is relatively easy to prove (Lemma 3). To prove (b), we first show that the monomials in $f(\mathcal{C}_\ell)$ corresponding to non-proper labeled walkages cancel out (Lemma 4). This argument is based on the now-standard technique of bijective labeling based cancellation introduced in [2]. The remaining part of the proof of (b) is much more complicated and is the main technical challenge. It is based on the technical Lemma 5, whose proof is postponed to Section 4.4.

We start with (a).

Lemma 3. *If G has a (k,w)-colored (S,T)-linkage of order p and length ℓ, then $f(\mathcal{C}_\ell)$ is non-zero.*

Proof. Consider a (k,w)-colored (S,T)-linkage $\mathcal{P}$ of order p and length ℓ. Let $X\subseteq V(\mathcal{P})$ be the set of vertices with $|X|=k$, different colors, and weight $\mathtt{we}(X)=w$. We can turn $\mathcal{P}$ into a proper labeled linkage $\mathcal{W}$ of order p, length ℓ, weight w, where $\mathtt{start}(\mathcal{W})=\mathtt{ordv}(S)$ and $\mathcal{T}(\mathcal{W})=T$, by

ordering the paths based on their starting vertices and assigning the labels $[k]$ arbitrarily to the vertices X when $\mathcal{W}$ intersects X.

Therefore $\mathcal{W} \in \mathcal{C}_\ell$, so it remains to prove that $\mathcal{W}$ is the only labeled walkage in $\mathcal{C}_\ell$ that corresponds to the monomial $f(\mathcal{W})$, which then implies that the monomial $f(\mathcal{W})$ occurs in the polynomial $f(\mathcal{C}_\ell)$ with coefficient 1, implying that $f(\mathcal{C}_\ell)$ is non-zero.

Notice that from $f(\mathcal{W})$, from the edge variables f_e we can recover the edges $E(\mathcal{W})$ of $\mathcal{W}$, from the vertex variables f_v we can recover the labeled vertices X, and because vertices in X have different colors, from the color-label pair variables f_c we can recover how the labels correspond to the labeled vertices. Therefore as the ordering of the paths is fixed by $\texttt{ordv}(S)$ and every vertex appears in $\mathcal{W}$ at most once, we have that $\mathcal{W}$ is the unique element of $\mathcal{C}_\ell$ that corresponds to the monomial $f(\mathcal{W})$. □

Then, we deal with non-proper walkages in $\mathcal{C}_\ell$. Let $\mathcal{C}^*_\ell \subseteq \mathcal{C}_\ell$ denote the family of proper labeled walkages in $\mathcal{C}_\ell$, i.e., the labeled walkages in $\mathcal{C}_\ell$ where all labeled indices have vertices of different colors.

Lemma 4. *It holds that* $f(\mathcal{C}^*_\ell) = f(\mathcal{C}_\ell)$.

Proof. We will show that there is a function $\phi : \mathcal{C}_\ell \setminus \mathcal{C}^*_\ell \to \mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$ that is an f-invariant fixed-point-free involution, i.e., for all $\mathcal{W} \in \mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$ it holds that (1) $f(\phi(\mathcal{W})) = f(\mathcal{W})$, (2) $\phi(\mathcal{W}) \neq \mathcal{W}$, and (3) $\phi(\phi(\mathcal{W})) = \mathcal{W}$. This implies that the set $\mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$ can be partitioned into pairs $\{\mathcal{W}, \phi(\mathcal{W})\}$ with $f(\mathcal{W}) = f(\phi(\mathcal{W}))$, and therefore every monomial corresponding to a labeled walkage in $\mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$ occurs in $f(\mathcal{C}_\ell)$ an even number of times, and therefore they cancel out because f is over a field of characteristic 2.

The function ϕ is defined as follows. Let $\mathcal{W} = (W^1, \ldots, W^p)$ be a labeled walkage in $\mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$, where $W^i = ((v^i_1, \ldots, v^i_{\ell_i}), (r^i_1, \ldots, r^i_{\ell_i}))$. Because $\mathcal{W}$ is semiproper but not proper, there exists two different labeled indices that have a vertex of the same color, i.e., pairs (i, a) and (j, b) with $i, j \in [p]$, $a \in [\ell_i]$, $b \in [\ell_j]$, $(i, a) \neq (j, b)$, $c(v^i_a) = c(v^j_b)$, $r^i_a \neq 0$, and $r^j_b \neq 0$. Let $(i, a), (j, b)$ be the lexicographically smallest such pair. We set $\phi(\mathcal{W})$ to be the labeled walkage obtained from $\mathcal{W}$ after swapping r^i_a with r^j_b.

First, we observe that $\phi(\mathcal{W}) \in \mathcal{C}_\ell$. Indeed, it cannot make a bijective walkage into non-bijective, and as it does not change the sequence of vertices of $\mathcal{W}$ or which indices are labeled, it cannot make a semiproper walk into non-semiproper, or change the order, the length, the weight, the tuple of starting vertices, or the set of ending vertices. Also $\phi(\mathcal{W})$ is not proper, i.e., $\phi(\mathcal{W}) \in \mathcal{C}_\ell \setminus \mathcal{C}^*_\ell$, because the vertices v^i_a and v^j_b are still labeled and have the same color.

To see why $f(\phi(\mathcal{W})) = f(\mathcal{W})$, note that, since ϕ does not change the vertices, it also does not change the edge variables f_e of the monomial, it does not change which vertices are labeled so it does not change the vertex variables f_v of the monomial, and because the vertices v^i_a and v^j_b have the same color the color-label-pair variables f_c of the monomial are also not changed.

Also, we have that $\phi(\mathcal{W}) \neq \mathcal{W}$, since the fact that $\mathcal{W}$ is bijective implies that $r^i_a \neq r^j_b$. Also, $\phi(\phi(\mathcal{W})) = \mathcal{W}$ because the swapping does not change which indices are labeled, and therefore does not change the lexicographically smallest pair of labeled indices with the same colors. □

As a result of Lemma 4, we can work with $f(\mathcal{C}^*_\ell)$ instead of $f(\mathcal{C}_\ell)$.

The most complicated part of the correctness proof will be to show part (b), that is, if there is no (k, w)-colored (S, T)-linkage of order p and length at most ℓ, then $f(\mathcal{C}^*_\ell)$ (and, thus by Lemma 4, $f(\mathcal{C}_\ell)$) is an identically zero polynomial. Most of this proof will be presented in Section 4.4, but we introduce here the statement the lemma that we will prove in Section 4.4. For this, we define *barren labeled walkages*.

Definition 2 (Barren labeled walkage)**.** *A labeled walkage $\mathcal{W}$ of length ℓ is* barren *if there exists no labeled linkage $\mathcal{W}'$ with starting vertices $\texttt{start}(\mathcal{W}') = \texttt{start}(\mathcal{W})$, set of ending vertices $\mathcal{T}(\mathcal{W}') = \mathcal{T}(\mathcal{W})$, set of collected vertices $R(\mathcal{W}') = R(\mathcal{W})$, length $\leq \ell$ and edges $E(\mathcal{W}') \subseteq E(\mathcal{W})$.*

In other words, a labeled walkage $\mathcal{W}$ of length ℓ is barren if its edges form a subgraph of G where no labeled *linkage* $\mathcal{W}'$ of length at most ℓ can have the same sets of starting vertices, ending vertices, and collected vertices as $\mathcal{W}$. Intuitively, this means that the labeled walkage $\mathcal{W}$ can not be "untangled" to give a corresponding labeled linkage. In particular, observe that because the "untangling" preserves the set of collected vertices, i.e., $R(\mathcal{W}') = R(\mathcal{W})$, if no (k, w)-colored (S, T)-linkages of order p and length at most ℓ exists, then all labeled walkages in $\mathcal{C}^*_\ell$ are barren.

Next, we state the main technical lemma for establishing the correctness of our algorithm. Section 4.4 is devoted to its proof.

Lemma 5. *Let G be a graph and let $\mathcal{B}$ the set of all proper barren labeled walkages in G. There exists a function $\phi : \mathcal{B} \to \mathcal{B}$ so that for all $\mathcal{W} \in \mathcal{B}$, the function ϕ satisfies that*

1. *$\phi(\phi(\mathcal{W})) = \mathcal{W}$ (ϕ is involution),*
2. *$\phi(\mathcal{W}) \neq \mathcal{W}$ (ϕ is fixed-point-free),*
3. *$f(\phi(\mathcal{W})) = f(\mathcal{W})$ (ϕ preserves the monomial),*
4. *$\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$ (ϕ preserves the set of ending vertices), and*
5. *$\texttt{start}(\phi(\mathcal{W})) = \texttt{start}(\mathcal{W})$ (ϕ preserves the ordered tuple of starting vertices).*

The main reason for defining the function ϕ for all proper barren labeled walkages instead of just barren walkages in $\mathcal{C}^*_\ell$ is that ϕ will be defined recursively, and in the recursion we will anyway need to handle all proper barren labeled walkages.

Now, the proof of (b) is an easy consequence of Lemma 5.

Lemma 6. *If G has no (k, w)-colored (S, T)-linkage of order p and length $\leq \ell$, then $f(\mathcal{C}^*_\ell)$ is an identically zero polynomial.*

Proof. First, because G has no (k, w)-colored (S, T)-linkage of order p and length $\leq \ell$, all labeled walkages in $\mathcal{C}^*_\ell$ are barren, i.e., $\mathcal{C}^*_\ell \subseteq \mathcal{B}$.

We show that if $\mathcal{W} \in \mathcal{C}^*_\ell$, then $\phi(\mathcal{W}) \in \mathcal{C}^*_\ell$. By definition, $\phi(\mathcal{W})$ is proper. By (3), ϕ preserves the set of labeled vertices and moreover because the labeled vertices have different colors it preserves also the label-vertex mapping, and therefore $\phi(\mathcal{W})$ is bijective and has weight w. By (4) and (5), ϕ preserves the set of ending vertices and the ordered tuple of starting vertices. By (3), ϕ also preserves the length ℓ, as the order of $\mathcal{W}$ is preserved by (5). Therefore the restriction $\phi \restriction_{\mathcal{C}^*_\ell}$ is a function $\phi \restriction_{\mathcal{C}^*_\ell} : \mathcal{C}^*_\ell \to \mathcal{C}^*_\ell$.

Then, by (1-3), $\phi \restriction_{\mathcal{C}^*_\ell}$ is an f-invariant fixed-point-free involution on $\mathcal{C}^*_\ell$, implying that the set $\mathcal{C}^*_\ell$ can be partitioned into pairs $\{\mathcal{W}, \phi(\mathcal{W})\}$ with $f(\mathcal{W}) = f(\phi(\mathcal{W}))$, and therefore for every monomial $f(\mathcal{W})$, there is an even number of labeled walkages $\mathcal{W} \in \mathcal{C}^*_\ell$ corresponding to it, and therefore because $f(\mathcal{C}^*_\ell)$ is over a field of characteristic 2, it is identically zero. □

4.4 Proof of Lemma 5

In this subsection we prove Lemma 5 by explicitly defining the function ϕ and then showing that it has all of the required properties.

In order to define ϕ we first introduce some notation for manipulating labeled walks and labeled walkages. Let $W = ((v_1, \ldots, v_\ell), (r_1, \ldots, r_\ell))$ be a labeled walk. For indices a, b with $1 \leq a \leq b \leq \ell$, we denote by $W[a, b]$ the labeled subwalk between a and b, inclusive, i.e., the labeled walk $W[a, b] = ((v_a, \ldots, v_b), (r_a, \ldots, r_b))$. If $a > b$, then $W[a, b]$ denotes an empty labeled walk.

The involution ϕ will use three types of operations: reversing a subwalk, swapping a label from one occurrence of a vertex to another occurrence of it (possibly in a different walk), and swapping suffixes of two walks.

The subwalk reversal operation is defined as follows. Let W be a labeled walk of length ℓ and a, b indices with $1 \leq a \leq b \leq \ell$. The walk obtained from W by reversing the subwalk between a and b, inclusive, including the labels, is denoted by $W\overleftarrow{[a, b]}$. For example, if $W = ((v_1, v_2, v_3, v_4), (0, 1, 0, 2))$, then $W\overleftarrow{[2, 3]} = ((v_1, v_3, v_2, v_4), (0, 0, 1, 2))$. A labeled walk W is a *palindrome* if $W = W\overleftarrow{[1, \ell]}$ holds, i.e., the labeled walk is the same in reverse. Note that $W\overleftarrow{[a, b]} = W$ holds if and only if $W[a, b]$ is a palindrome and that a subwalk $W[a, b]$ can be a palindrome only if its length is odd or if it is the empty walk. We will use the following lemma about palindromic subwalks of labeled walks, and in particular the reason to forbid labeled digons is to make this lemma true. Recall that any labeled walk in a proper labeled walkage is injective and does not contain labeled digons. Recall also that $R(W[a + 1, b - 1]) = \emptyset$ if and only if W has no labels in the subwalk $W[a + 1, b - 1]$.

Lemma 7. *Let $W = ((v_1, \ldots, v_\ell), (r_1, \ldots, r_\ell))$ be an injective labeled walk of length ℓ that does not contain labeled digons, and let $a, b \in [\ell]$. If $v_a = v_b$ and $W[a + 1, b - 1]$ is a palindrome, then $R(W[a + 1, b - 1]) = \emptyset$.*

Proof. First, because $W[a + 1, b - 1]$ is injective and palindrome, the only vertex of $W[a + 1, b - 1]$ that can be labeled is the middle vertex. However, a label cannot occur at the middle vertex of a palindrome with more than one vertex because it would be a labeled digon. If $W[a + 1, b - 1]$ has exactly one vertex, then again this vertex cannot be labeled because $v_a = v_b$ and W does not contain labeled digons. □

The label swap operation is defined as follows. Let $\mathcal{W} = (W^1, \ldots, W^p)$ be a labeled walkage of order p, where for each $i \in [p]$ the walkage W^i is denoted by $((v^i_1, \ldots, v^i_{\ell_i}), (r^i_1, \ldots, r^i_{\ell_i}))$. Let (i, a), (j, b) be pairs with $i, j \in [p]$, $a \in [\ell_i]$, $b \in [\ell_j]$, $v^i_a = v^j_b$, and exactly one of r^i_a and r^j_b equal to zero (i.e. one of them unlabeled and one labeled). The labeled walkage obtained from $\mathcal{W}$ by swapping r^i_a with r^j_b is denoted by $\mathcal{W} \frown^{i,j}_{a,b}$. Note that because $r^i_a \neq r^j_b$, it holds that $\mathcal{W} \frown^{i,j}_{a,b} \neq \mathcal{W}$. See Figure 6 for an illustration of the label swap operation.

The suffix swap operation is defined as follows. Let (i, a) and (j, b) be pairs with $i, j \in [p]$, $a \in [\ell_i + 1]$, $b \in [\ell_j + 1]$, and $i \neq j$. The labeled walkage obtained from $\mathcal{W}$ by swapping the suffix of W^i starting at index a with the suffix of W^j starting at index b is denoted by $\mathcal{W} \leftrightarrow^{i,j}_{a,b}$. Note that here we allow that $a = \ell_i + 1$ or $b = \ell_j + 1$, with the interpretation that this corresponds to the empty suffix. Clearly, if both $a = \ell_i + 1$ and $b = \ell_j + 1$, then this operation does not do anything, but otherwise if $\mathcal{W}$ is a proper labeled walkage, applying this operation will in fact always result in a different walkage because of the different ending vertices condition. See Figure 7 for an illustration of the suffix swap operation.

If W^1 and W^2 are labeled walks so that the last vertex of W^1 is adjacent to the first vertex of W^2, then $W^1 \circ W^2$ denotes the concatenation of W^1 and W^2. If $\mathcal{W} = (W^1, \ldots, W^p)$ is a labeled walkage and W is a labeled walk, then $W \diamond \mathcal{W}$ denotes the labeled walkage $(W \circ W^1, \ldots, W^p)$ and $W \sqcup \mathcal{W}$ denotes the labeled walkage $(W, W^1, \ldots, W^p)$.

Next we define the function ϕ of Lemma 5. We will provide some intuition about ϕ right after the definition, and Figures 8 to 16 demonstrate different cases of it. The definition of ϕ will be recursive, using induction by the length of the walkage.

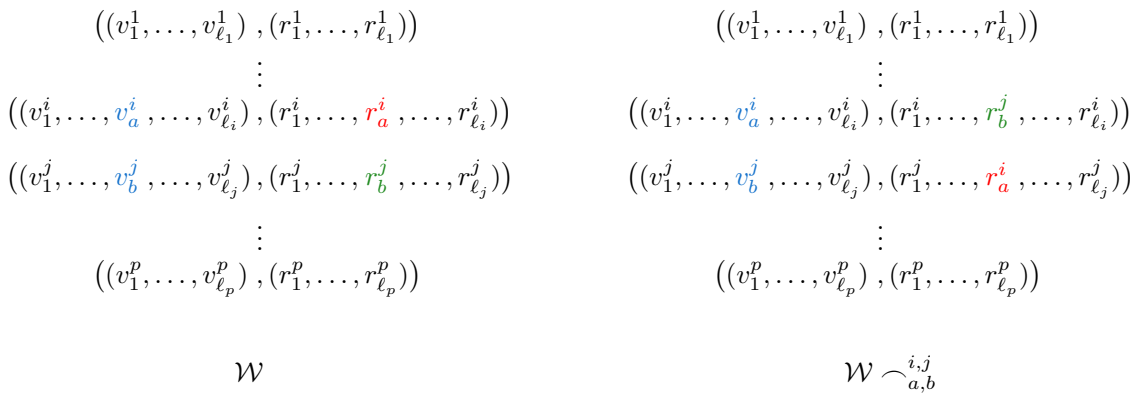

Figure 6: An illustration of the label swap operation. On the left: a labeled walkage $\mathcal{W} = (W^1, \ldots, W^p)$, and pairs (i, a), (j, b) with $i, j \in [p]$, $a \in [\ell_i]$, $b \in [\ell_j]$, $v^i_a = v^j_b$, and exactly one of r^i_a and r^j_b equal to zero. Note that we allow $i = j$. On the right: the labeled walkage $\mathcal{W} \frown^{i,j}_{a,b}$.

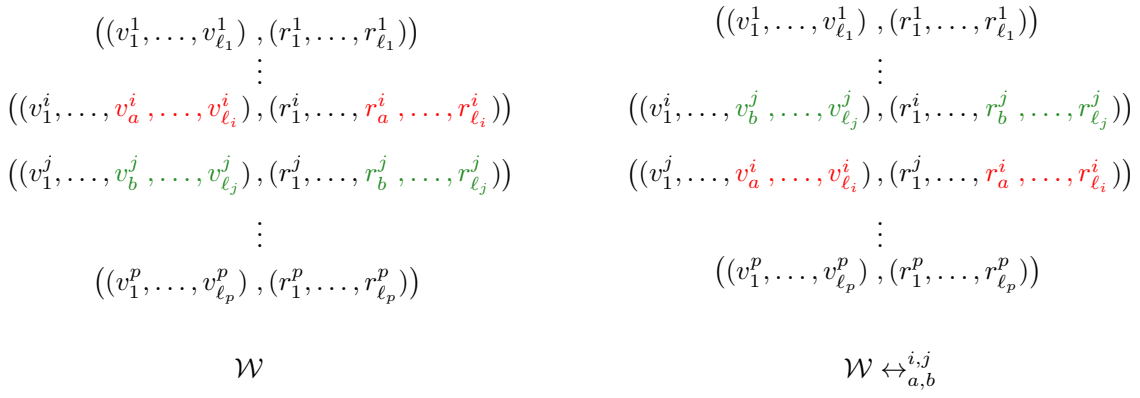

Figure 7: An illustration of the suffix swap operation. On the left: a labeled walkage $\mathcal{W} = (W^1, \ldots, W^p)$, and pairs (i, a), (j, b) with $i, j \in [p]$, $a \in [\ell_i + 1]$, $b \in [\ell_j + 1]$, and $i \neq j$. On the right: the labeled walkage $\mathcal{W} \leftrightarrow^{i,j}_{a,b}$.

Definition 3 (The function ϕ)**.** *Let $\mathcal{W} = (W^1, \ldots, W^p)$ be a proper barren labeled walkage of order p. For each $i \in [p]$, denote $W^i = ((v^i_1, \ldots, v^i_{\ell_i}), (r^i_1, \ldots, r^i_{\ell_i}))$. The value $\phi(\mathcal{W})$ is defined, in some cases recursively, by selecting the first matching case from the following list:*

A. if the vertex v^1_1 occurs only once in $\mathcal{W}$:

1. *if $\ell_1 \geq 2$, then $\phi(\mathcal{W}) = W^1[1,1] \diamond \phi(W^1[2, \ell_1], W^2, \ldots, W^p)$.*
2. *otherwise (i.e., $\ell_1 = 1$), $\phi(\mathcal{W}) = W^1 \sqcup \phi(W^2, \ldots, W^p)$.*

B. if the vertex v^1_1 occurs in at least three different walks W^i:
There must be at least two different walks W^i that contain v^1_1 but do not contain it as labeled. Let i, j be the two smallest indices so that both W^i and W^j contain v^1_1 but do not contain it as labeled. Let a be the index of the first occurrence of v^1_1 in W^i and b be the index of the first occurrence of v^1_1 in W^j. Now, $\phi(\mathcal{W}) = \mathcal{W} \leftrightarrow^{i,j}_{a,b}$.

C. if the vertex v^1_1 occurs only in the walk W^1:
By the case (A), the vertex v^1_1 occurs multiple times in W^1. Let b be the index of the last

occurrence of v_1^1 in W^1 and a be the index of the second last occurrence of v_1^1 in W^1. Note that $a = 1$ if v_1^1 occurs only twice in W^1, and note also that $1 \le a \le b-2$.

1. *if $r_1^1 = r_b^1 = 0$:*
   (a) *if $W^1[2, b-1]$ is not a palindrome, then $\phi(\mathcal{W}) = (W^1[\overleftarrow{2, b-1}], W^2, \ldots, W^p)$.*
   (b) *otherwise, if $b < \ell_1$, then $\phi(\mathcal{W}) = W^1[1, b] \diamond \phi(W^1[b+1, \ell_1], W^2, \ldots, W^p)$.*
   (c) *otherwise (i.e., $b = \ell_1$), $\phi(\mathcal{W}) = W^1 \sqcup \phi(W^2, \ldots, W^p)$.*
2. *if the index b is not a digon in W^1, then $\phi(\mathcal{W}) = \mathcal{W} \frown_{1,b}^{1,1}$.*
   Note: If neither case (1) nor (2) applies, then $r_1^1 \neq 0$.
3. *if $W^1[2, a-1]$ is not a palindrome, then $\phi(\mathcal{W}) = (W^1[\overleftarrow{2, a-1}], W^2, \ldots, W^p)$.*
   Note: If $a = 1$, then $W^1[2, a-1]$ is the empty walk which is a palindrome.
4. *if $v_{a+1}^1 = v_{b-1}^1$:*
   (a) *if $W^1[a+1, b-1]$ is not a palindrome, then $\phi(\mathcal{W}) = (W^1[\overleftarrow{a+1, b-1}], W^2, \ldots, W^p)$.*
   (b) *otherwise, $\phi(\mathcal{W}) = W^1[1, b] \diamond \phi(W^1[b+1, \ell_1], W^2, \ldots, W^p)$.*
   Note: Here $W^1[b+1, \ell_1]$ cannot be an empty walk because by case (C.2) b is a digon in W^1.

X. *otherwise, $\phi(\mathcal{W}) = W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$.*
   Note: The case C.X will form a "common case" with the case D.X.

D. *if the vertex v_1^1 occurs in exactly two different walks:*
Let i be the index of the another walk W^i in which v_1^1 occurs and let b be the index of the first occurrence of v_1^1 in W^i.

1. *if $r_1^1 = r_b^i = 0$, then $\phi(\mathcal{W}) = \mathcal{W} \leftrightarrow_{1,b}^{1,i}$.*
2. *if the index b is not a digon in W^i, then $\phi(\mathcal{W}) = \mathcal{W} \frown_{1,b}^{1,i}$.*
   Note: If neither case (1) nor (2) applies, then $r_1^1 \neq 0$.
3. *if v_1^1 occurs at least twice in W^i, then let c be the index of its second occurrence and $\phi(\mathcal{W}) = \mathcal{W} \leftrightarrow_{2,c+1}^{1,i}$.*
   Note: It can happen that one of the suffixes in this case is empty. However, both of them cannot be empty at the same time because W^1 and W^i have different ending vertices because $\mathcal{W}$ is proper.
   Note: In the remaining cases, v_1^1 occurs exactly once in W^i, and this occurrence is a digon at index b.
   Now, let a be the index of the last occurrence of v_1^1 in W^1 (if v_1^1 occurs only once in W^1, then $a = 1$).
4. *if $W^1[2, a-1]$ is not a palindrome, then $\phi(\mathcal{W}) = (W^1[\overleftarrow{2, a-1}], W^2, \ldots, W^p)$.*
   Note: If $a = 1$, then $W^1[2, a-1]$ is the empty walk which is a palindrome.
5. *if $a = \ell_1$, then $\phi(\mathcal{W}) = W^1 \sqcup \phi(W^2, \ldots, W^p)$.*
6. *if $v_{a+1}^1 = v_{b+1}^i$, then $\phi(\mathcal{W}) = \mathcal{W} \leftrightarrow_{a+1,b+1}^{1,i}$.*
   Note: By case (5) it holds that $a < \ell_1$ and by case (2) it holds that $b < \ell_i$.

X. *otherwise, $\phi(\mathcal{W}) = W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$.*
   Note: The case D.X will form a "common case" with the case C.X.

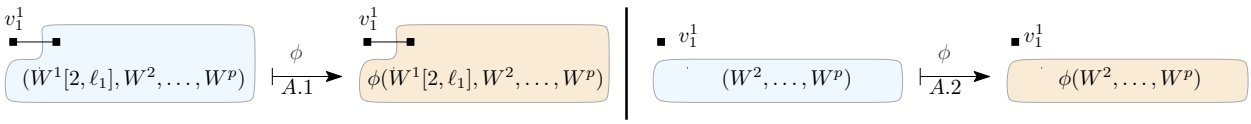

Figure 8: Examples of cases A.1 and A.2 of the definition of ϕ. The vertex v_1^1 can be either labeled or unlabeled.

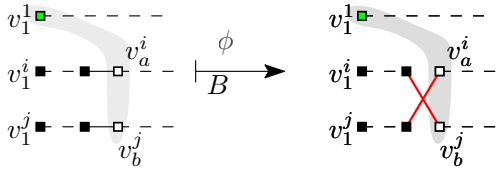

Figure 9: Example of case B of the definition of ϕ. All vertices inside the grey bag are different occurrences of the same vertex v_1^1 of the graph. The white vertices v_a^i and v_b^j are unlabeled, the black vertices could be labeled or unlabeled, and the green vertex v_1^1 is labeled in this specific example.

Intuition for ϕ. Before laboriously proving that ϕ indeed is a function from proper barren labeled walkages to proper barren labeled walkages satisfying the required properties, let us give some rough outline of ideas behind it. First, the general idea is that if the walkage $\mathcal{W}$ goes to a certain case, then the walkage $\phi(\mathcal{W})$ goes again to the same case, which then maps it back to $\mathcal{W}$. The only exception is that the cases C.X and D.X could map to each other.

Then, let us consider the cases relevant for a single walk, i.e., the case A.1 and the cases under C. Here, the intuition of case A.1 is to just move forward in the walk: we don't care much about what ϕ does to the rest of the walk because it must preserve the vertex right after v_1^1, and attaching v_1^1 to the front will not create a digon because v_1^1 occurs only at one index. Then, case C.1.a is the standard loop reversal case, which is safe because neither index 1 nor b is labeled. The case C.1.b (and C.1.c) corresponds to ignoring a palindromic subwalk, which can be safely done by Lemma 7. Then, case C.2 is the standard label swap case, which is safe because the index b is a not digon (note that the index 1 is never a digon). The cases C.1–C.2 are in some sense the "easy cases", while the cases C.3–C.X require more analysis of the remaining situation and quite unintuitive design. First, if neither C.1 nor C.2 applies, we know that the index 1 is labeled and the index b is digon. The purpose of case C.3 is to, in some sense reduce to a situation where we pretend that the vertex v_1^1 occurs only at indices 1, a, and b, as the walk between 1 and a is an irrelevant palindromic loop. Then, case C.4 handles a corner condition which would prevent case C.X from working. The case C.X ignores the palindromic loop between 1 and a, leaving the only occurrence of the vertex v_1^1 in the rest of the walk to be at the digon b, which in some sense makes it "harmless" in that the recursive calls will never need to analyse the vertex v_b^1 again as the first vertex.

The intuition for the case of multiple walks is as follows. First, the case A.2 is just an analogue of A.1 when the first walk has length 1. Then, if the vertex v_1^1 occurs multiple times, we consider three different cases: v_1^1 occurs in at least three walks, v_1^1 occurs in one walk, and v_1^1 occurs in two walks. Here, the three walks case B is quite easy, as we can just consider two of the walks where v_1^1 is not labeled, circumventing all issues with labeled digons. When v_1^1 occurs in only one walk we go to the one walk case C. Then, when v_1^1 occurs in two walks W^1 and W^i, the intuition of cases under D is that we concatenate W^1 with reversed W^i, with some special marker in between, and then apply the single walk cases under C for this concatenation. Here, in the case D.X this can change whether v_1^1 occurs in two walks or a single walk, and therefore it is necessary to have the

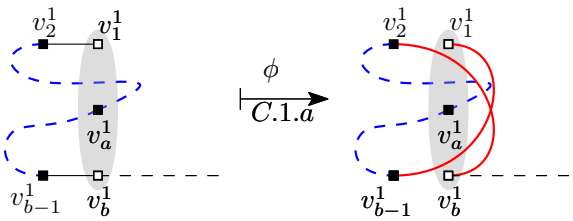

Figure 10: Example of case C.1.a of the definition of ϕ. All vertices inside the grey bag are different occurrences of the same vertex v_1^1 of the graph. The vertices v_1^1 and v_b^1 are unlabeled and the black vertices v_2^1, v_a^1, and v_{b-1}^1 can be either labeled or unlabeled.

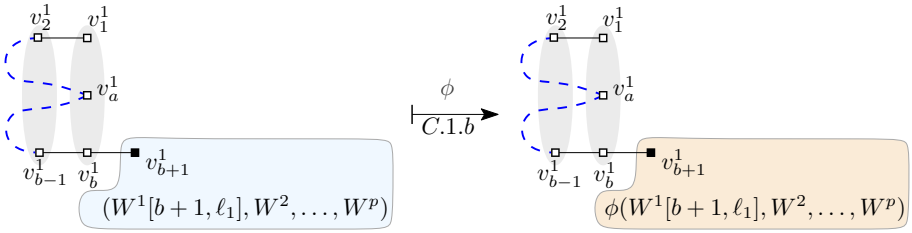

Figure 11: Example of case C.1.b of the definition of ϕ. All vertices inside the grey bags are the same vertex of the graph. By case C.1.a, the blue subwalk $W^1[2, b-1]$ is palindrome, and therefore by Lemma 7, the vertices in it are unlabeled.

common case of C.X and D.X, moreover taking care in the proof that moving back from C.X to D.X will be handled correctly.

Correctness proof for ϕ. We will then proceed to first show that ϕ is well-defined, then that ϕ maps proper barren labeled walkages to proper barren labeled walkages, and then that ϕ satisfies all of the properties stated in Lemma 5, with $\phi(\phi(\mathcal{W})) = \mathcal{W}$ being the most complicated of them to prove. The proof is long because we have to analyze most of the 18 cases one by one. However, most of the arguments in these proofs are relatively easy once the definition of ϕ is set. The main challenge in the proof was to come up with the right definition of ϕ.

Well-definedness of ϕ. In Definition 3, in several cases, namely A.1, A.2, C.1.b, C.1.c, C.4.b, C.X, D.5, and D.X, the function ϕ is defined recursively. A priori it is not even clear why the syntactic value $\phi(\mathcal{W})$ is even well-defined in these cases. It requires proof that in these cases the recursive argument is in the domain of ϕ, in particular that it is also a proper barren labeled walkage.

Next we show that the syntactic value $\phi(\mathcal{W})$ for proper barren labeled walkages $\mathcal{W}$ is well-defined. We remark that Lemma 8 does not yet show that $\phi(\mathcal{W})$ is a proper barren labeled walkage; it will require more efforts to prove (see Lemmas 9 to 11).

Lemma 8. *In case A.1 of Definition 3 it holds that $(W^1[2,\ell_1], W^2, \ldots, W^p)$ is a proper barren labeled walkage, in cases A.2, C.1.c, and D.5 it holds that $(W^2, \ldots, W^p)$ is a proper barren labeled walkage, in cases C.1.b, and C.4.b it holds that $(W^1[b+1,\ell_1], W^2, \ldots, W^p)$ is a proper barren labeled walkage, and in cases C.X and D.X it holds that $(W^1[a+1,\ell_1], W^2, \ldots, W^p)$ is a proper barren labeled walkage.*

Proof. In all cases, the labeled walkage used as the recursive argument is obtained from $\mathcal{W}$ by removing either the walk W^1 or a prefix of W^1. First we need to argue that the recursive argument is a labeled walkage. For this, the only thing to argue is that (1) the recursive argument contains at least one walk (i.e., $p \geq 2$ in cases A.2, C.1.c, and D.5) and that (2) all walks in the recursive

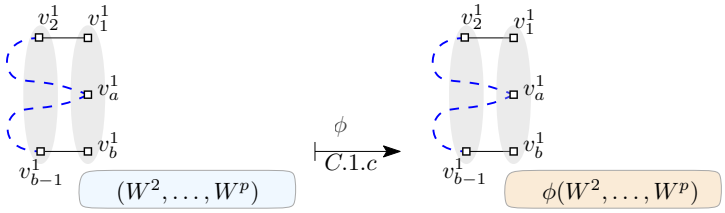

Figure 12: Example of case C.1.c of the definition of ϕ. All vertices inside the grey bags correspond to the same vertex of the graph. By case C.1.a, the blue subwalk $W^1[2, b-1]$ is palindrome, and therefore by Lemma 7, the vertices in it are unlabeled.

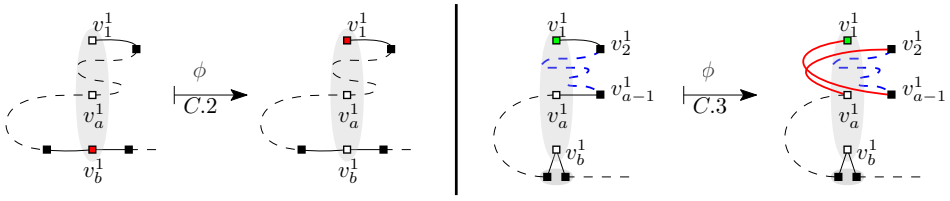

Figure 13: Example of cases C.2 and C.3 of the definition of ϕ. Vertices inside the same grey bags correspond to the same vertex of the graph. On the left part of the figure (case C.2), in the initial configuration the vertex v_1^1 is unlabeled and the vertex v_b^1 is labeled (with color red) and the application of ϕ in this case exchanges this label from v_b^1 to v_1^1. On the right part of the figure, v_b^1 is a digon and therefore it is unlabeled and by cases C.1 and C.2, v_1^1 has to be labeled (depicted in green).

argument are non-empty (i.e. $\ell_1 \geq 2$ in case A.1, $b < \ell_1$ in cases C.1.b and C.4.b, and $a < \ell_1$ in cases C.X and D.X). The other properties of labeled walkages are clearly satisfied when removing either W^1 or a prefix of W^1.

The above conditions are satisfied directly by definition in cases A.1 and C.1.b. In case C.4.b, $b < \ell_1$ holds by the fact that (due to case C.2) the index b is a digon in W^1. In case C.X, recall that a is the index of the second last occurrence of v_1^1 in W^1, so $a < \ell_1$. In case D.X, we have that $a < \ell_1$ by case D.5. For the remaining cases A.2, C.1.c, and D.5, observe the following. If $p = 1$ would hold, then $\mathcal{W} = (W^1)$. Then, since in all these three cases $v_{\ell_1}^1 = v_1^1$ and $W^1[2, \ell_1]$ cannot contain labels (in A.2 trivially, in C.1.c by $r_1^1 = r_b^1 = 0$ and Lemma 7, and in D.5 by cases D.1, D.2, and D.4 combined with Lemma 7), it should hold that $(W^1[1, 1])$ is a labeled linkage consisting only of one walk with one vertex that would contradict the fact that $\mathcal{W}$ is barren by Definition 2.

It is clear by definition of a proper labeled walkage that removing a walk or a prefix of a walk maintains that the walkage is proper. To complete the proof, it remains to show case by case that the labeled walkages used as recursive arguments are barren.

In all of the cases the proofs will follow the same template: For the sake of contradiction we suppose that the labeled walkage $\mathcal{W}'$ used as a recursive argument is not barren, then consider the labeled linkage $\mathcal{W}''$ that witnesses that $\mathcal{W}'$ is not barren, and then use $\mathcal{W}''$ to construct a labeled linkage that shows that $\mathcal{W}$ is not barren, obtaining a contradiction. We spell out these steps in detail for the case A.1, and in less detail for subsequent cases.

Case A.1. For the sake of contradiction, suppose that $\mathcal{W}' = (W^1[2, \ell_1], W^2, \ldots, W^p)$ is not barren. Then by the definition of barren, there exists a labeled linkage $\mathcal{W}''$ with $\texttt{start}(\mathcal{W}'') = \texttt{start}(\mathcal{W}')$, $\mathcal{T}(\mathcal{W}'') = \mathcal{T}(\mathcal{W}')$, $R(\mathcal{W}'') = R(\mathcal{W}')$, length $\leq \ell - 1$, and edges $E(\mathcal{W}'') \subseteq E(\mathcal{W}')$. By the assumptions of case A.1, the labeled walkage $\mathcal{W}'$ does not contain v_1^1, so the labeled linkage $\mathcal{W}''$ cannot contain

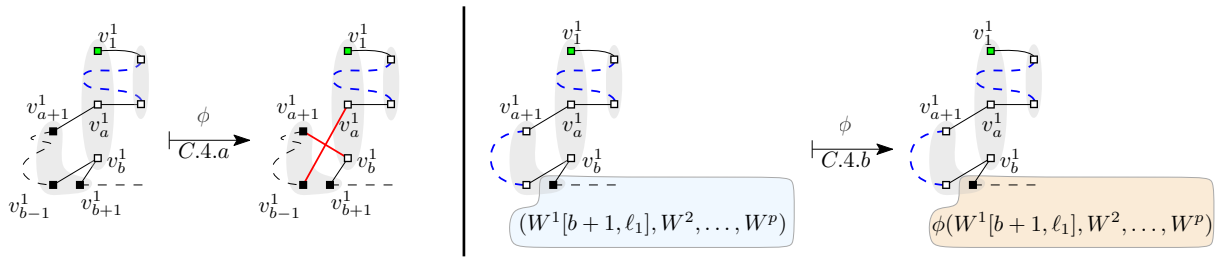

Figure 14: Examples of cases C.4.a and C.4.b of the definition of ϕ. By case C.2, b is a digon on W^1 and by case C.3, $W^1[2, a-1]$ is a palindrome. For both case C.4.a and case C.4.b, we have that $v^1_{a+1} = v^1_{b-1}$ (v^1_{a+1}, v^1_{b-1}, and v^1_{b+1} are in the same grey bag). If $W^1[a+1, b-1]$ is not a palindrome, then we are in case C.4.a (on the left), while if $W^1[a+1, b-1]$ is a palindrome, we are in case C.4.b. (on the right).

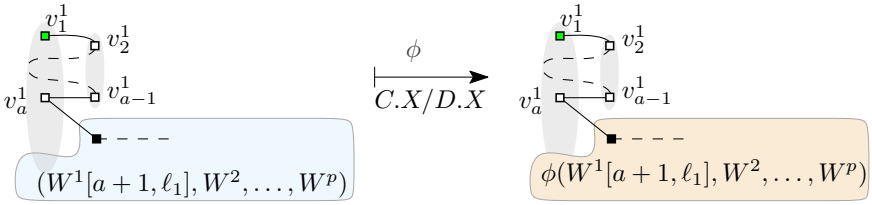

Figure 15: Examples of cases C.X and D.X of the definition of ϕ.

v_1^1 because the edge property $E(\mathcal{W}'') \subseteq E(\mathcal{W}')$ ensures that v_1^1 cannot occur in a walk of length more than one, and the start vertex property $\texttt{start}(\mathcal{W}'') = \texttt{start}(\mathcal{W}')$ ensures that v_1^1 cannot occur in a walk of length one. By the start vertex property, it holds that the first vertex of the first walk in $\mathcal{W}''$ is v_2^1. Therefore, $W^1[1,1] \diamond \mathcal{W}''$ is a labeled linkage. Because $v_1^1 v_2^1 \in E(\mathcal{W})$ and $E(\mathcal{W}'') \subseteq E(\mathcal{W}') \subseteq E(\mathcal{W})$, we have that $E(W^1[1,1] \diamond \mathcal{W}'') \subseteq E(\mathcal{W})$. Also, observe that because of $R(\mathcal{W}'') = R(\mathcal{W}')$, it holds that $R(W^1[1,1] \diamond \mathcal{W}'') = R(\mathcal{W})$. Similarly, we observe that $\texttt{start}(W^1[1,1] \diamond \mathcal{W}'') = \texttt{start}(\mathcal{W})$, $\mathcal{T}(W^1[1,1] \diamond \mathcal{W}'') = \mathcal{T}(\mathcal{W})$, and the length of $\mathcal{W}''$ is at most ℓ. Therefore, $W^1[1,1] \diamond \mathcal{W}''$ is a labeled linkage that according to Definition 2 contradicts the fact that $\mathcal{W}$ is barren.

Case A.2. Again, suppose that $\mathcal{W}' = (W^2, \ldots, W^p)$ is not barren, and consider the witness $\mathcal{W}''$. Because v_1^1 does not occur in $\mathcal{W}''$, it holds that $W^1 \sqcup \mathcal{W}''$ is a labeled linkage that contradicts that $\mathcal{W}$ is barren.

Case C.1.b. Suppose that $\mathcal{W}' = (W^1[b+1, \ell_1], W^2, \ldots, W^p)$ is not barren and consider the witness $\mathcal{W}''$. As, by definition of b in case C, v_1^1 occurs in $\mathcal{W}$ only in the subwalk $W^1[1,b]$, it cannot occur in $\mathcal{W}''$. Therefore $W^1[1,1] \diamond \mathcal{W}''$ is a labeled linkage. It is easy to observe that $\texttt{start}(W^1[1,1] \diamond \mathcal{W}'') = \texttt{start}(\mathcal{W})$, $\mathcal{T}(W^1[1,1] \diamond \mathcal{W}'') = \mathcal{T}(\mathcal{W})$, and $E(W^1[1,1] \diamond \mathcal{W}'') \subseteq E(\mathcal{W})$. Also, Lemma 7 implies that $R(W^1[1,b]) = \emptyset$ and therefore $R(W^1[1,1] \diamond \mathcal{W}'') = R(\mathcal{W})$, therefore contradicting that $\mathcal{W}$ is barren.

Case C.1.c. This case is similar as the previous, in particular, Lemma 7 implies that $R(W^1) = \emptyset$. Therefore, if we assume that $\mathcal{W}' = (W^2, \ldots, W^p)$ is not barren and we take the labeled linkage $\mathcal{W}''$ that witnesses that $\mathcal{W}'$ is not barren, we can construct a labeled linkage $W^1[1,1] \sqcup \mathcal{W}''$ that contradicts the fact that $\mathcal{W}$ is barren.

Case C.4.b. Assume that $\mathcal{W}' = (W^1[b+1, \ell_1], W^2, \ldots, W^p)$ is not barren and take the labeled linkage $\mathcal{W}''$ that witnesses that $\mathcal{W}'$ is not barren. As v_1^1 occurs in $\mathcal{W}$ only in the subwalk $W^1[1,b]$, it cannot occur in $\mathcal{W}''$. Therefore $W^1[1,1] \diamond \mathcal{W}''$ is a labeled linkage. Note that in this case $W^1[2, a-1]$ is a palindrome, the index a of the walk W^1 is not labeled because the index 1 is

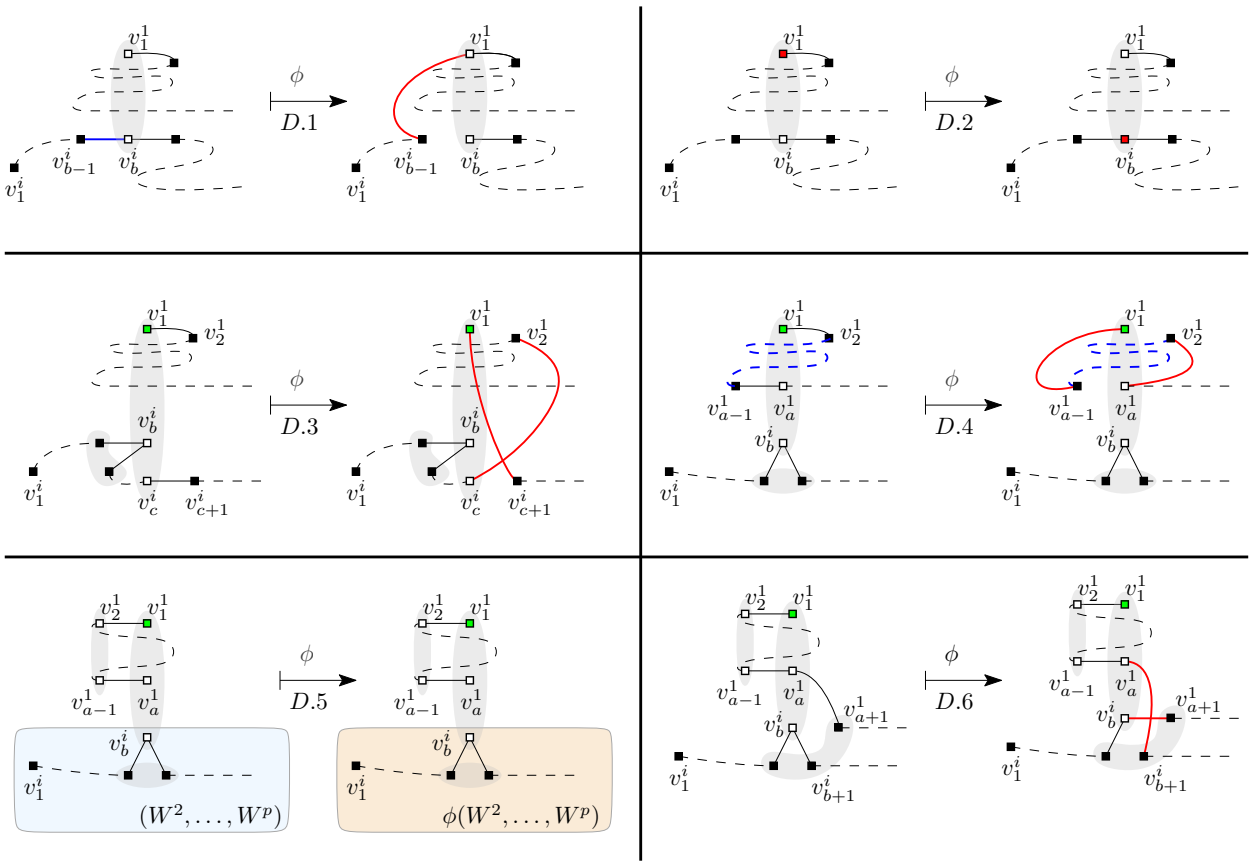

Figure 16: Example of cases D.1-D.6 of the definition of ϕ. In each grey bag, all vertices inside the corresponding bag are the same vertex of the graph. Labeled vertices are depicted in red and green. White vertices correspond to unlabeled vertices and black vertices can be either labeled or unlabeled.

labeled and $\mathcal{W}$ is proper, and $W^1[a+1, b-1]$ is a palindrome, and the index b of W^1 is not labeled. Therefore, by Lemma 7, $R(W^1[1, b]) = \{v_1^1\}$, and therefore $R(W^1[1, 1] \diamond \mathcal{W}'') = R(\mathcal{W})$, and therefore we contradict the fact that $\mathcal{W}$ is barren.

Case C.X. In this case, the argument is less apparent because v_1^1 indeed occurs in $W^1[a+1, \ell_1]$. Again, we start by assuming that $\mathcal{W}' = (W^1[a+1, \ell_1], W^2, \ldots, W^p)$ is not barren, and take the labeled linkage $\mathcal{W}''$ that witnesses that $\mathcal{W}'$ is not barren. Now, note that v_1^1 occurs in $\mathcal{W}'$ only as a single digon in $W^1[a+1, \ell_1]$. Therefore, v_1^1 cannot occur as a starting or ending vertex in $\mathcal{W}''$. Also, there is only one edge in $E(\mathcal{W}')$ incident to v_1^1, which then prevents v_1^1 occuring at any position in $\mathcal{W}''$, because any position containing v_1^1 would have to a digon, but $\mathcal{W}''$ is labeled linkage and thus does not contain digons. Therefore, we construct a labeled linkage $W^1[1, 1] \diamond \mathcal{W}''$ and use the fact that $W^1[2, a-1]$ is palindrome with Lemma 7 to conclude that $R(W^1[1, 1] \diamond \mathcal{W}'') = R(\mathcal{W})$, and to finally observe that $W^1[1, 1] \diamond \mathcal{W}''$ satisfies also all the other needed properties to contradict the fact that $\mathcal{W}$ is barren.

Case D.5. Note that here we have that $r_1^1 \neq 0$, $v_1^1 = v_{\ell_1}^1$, and $W^1[2, \ell_1 - 1]$ is a palindrome. The arguments are similar to case C.X: The vertex v_1^1 does occur in the walkage $(W^2, \ldots, W^p)$, but it occurs in it only a single time, which is a digon in the walk W^i. So suppose that $\mathcal{W}' = (W^2, \ldots, W^p)$ is not barren, and consider the witness $\mathcal{W}''$. By similar arguments as in case C.X, v_1^1 cannot occur

in $\mathcal{W}''$. Therefore, we construct a labeled linkage $W^1[1,1] \sqcup \mathcal{W}''$ and use the fact that $W^1[2, a-1]$ is palindrome with Lemma 7 to conclude that $R(W^1[1,1] \sqcup \mathcal{W}'') = R(\mathcal{W})$, and to finally observe that $W^1[1,1] \sqcup \mathcal{W}''$ satisfies also all the other needed properties to contradict the fact that $\mathcal{W}$ is barren.

Case D.X. Note that here again, we have that $r_1^1 \neq 0$ and $W^1[2, a-1]$ is a palindrome, where a is the last occurrence of v_1^1 in W^1. The arguments are similar to case C.X: The vertex v_1^1 does occur in the walkage $(W^1[a+1, \ell_1], W^2, \ldots, W^p)$, but it occurs in it only a single time, which is a digon in the walk W^i. So again suppose that $\mathcal{W}' = (W^1[a+1, \ell_1], W^2, \ldots, W^p)$ is not barren, and consider the witness $\mathcal{W}''$. Again by arguments of C.X we have that v_1^1 cannot occur in $\mathcal{W}''$. Therefore, we again construct a labeled linkage $W^1[1,1] \diamond \mathcal{W}''$ and use the fact that $W^1[2, a-1]$ is palindrome with Lemma 7 to conclude that $R(W^1[1,1] \diamond \mathcal{W}'') = R(\mathcal{W})$, and to finally observe that $W^1[1,1] \diamond \mathcal{W}''$ satisfies also all the other needed properties to contradict that $\mathcal{W}$ is barren. □

The next three lemmas establish that $\phi(\mathcal{W})$ is a proper barren labeled walkage. In addition, Lemma 9 shows that ϕ satisfies the properties $f(\phi(\mathcal{W})) = f(\mathcal{W})$, $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$, and $\texttt{start}(\phi(\mathcal{W})) = \texttt{start}(\mathcal{W})$.

Lemma 9. *Let $\mathcal{W}$ be a proper barren labeled walkage. It holds that $\phi(\mathcal{W})$ is a labeled walkage, $\texttt{start}(\phi(\mathcal{W})) = \texttt{start}(\mathcal{W})$, $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$, and $f(\phi(\mathcal{W})) = f(\mathcal{W})$.*

Proof. We prove the lemma by induction on the length of the walkage $\mathcal{W}$. Here, all of the cases should be easy to verify, so the arguments we provide will be terse.

Cases A.1 works directly by induction, in particular, we can use the induction assumptions that

- $\phi(W^1[2, \ell_1], W^2, \ldots, W^p)$ is a labeled walkage,
- $\texttt{start}(\phi(W^1[2, \ell_1], W^2, \ldots, W^p)) = \texttt{start}((W^1[2, \ell], W^2, \ldots, W^p))$,
- $\mathcal{T}(\phi(W^1[2, \ell_1], W^2, \ldots, W^p)) = \mathcal{T}((W^1[2, \ell], W^2, \ldots, W^p))$, and
- $f(\phi(W^1[2, \ell_1], W^2, \ldots, W^p)) = f((W^1[2, \ell], W^2, \ldots, W^p))$,

to prove the same properties for $\mathcal{W}$. In particular, we use the start vertex property to ensure that the first vertex of the first walk of $\phi(W^1[2, \ell_1], W^2, \ldots, W^p)$ is v_2^1, and therefore the first edge of the first walk of $W^1[1,1] \diamond \phi(W^1[2, \ell_1], W^2, \ldots, W^p)$ is $v_1^1 v_2^1$.

Case A.2 works similarly to A.1. Case B works by the property that $v_a^i = v_b^j$, in particular observing that if both of the suffixes are non-empty, the suffix swap operation preserves the set of ending vertices $\mathcal{T}(\mathcal{W})$. Case C.1.a works because $v_1^1 = v_b^1$, and cases C.1.b and C.1.c work by similar induction as A.1. Case C.2 works because $v_1^1 = v_b^1$, in particular, even though the index of the label in the walk changes, the vertex variable or the label-color pair variable do not change because the vertex does not change. Case C.3 works because $v_1^1 = v_a^1$. Case C.4.a works because $v_a^1 = v_b^1$ and case C.4.b works by similar induction as A.1. Case C.X works again by induction. Case D.1 works because $v_1^1 = v_b^i$. Case D.2 works because $v_1^1 = v_b^i$, by the same argument as case C.2.

In case D.3, all other conditions work directly by $v_1^1 = v_c^i$, but we should pay attention to the ending vertices condition $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$, because it can happen that one of the suffixes is empty. As observed already in the definition, observe that at most one of the suffixes can be empty because $v_1^1 = v_c^i$ and W^1 and W^i have different ending vertices because $\mathcal{W}$ is proper. First, if $\ell_1 = 1$, then the ending vertex of W^i becomes $v_c^i = v_1^1 = v_{\ell_1}^1$, and the ending vertex of W^1 becomes $v_{\ell_i}^i$, so the condition holds. Second, if $\ell_i = c$, then the ending vertex of W^1 becomes $v_1^1 = v_c^i = v_{\ell_i}^i$, and the ending vertex of W^i becomes $v_{\ell_1}^1$, so the condition holds.

Case D.4 works because $v_1^1 = v_a^1$. Case D.5 works by induction. Case D.6 works because $v_{a+1}^1 = v_{b+1}^i$ and the suffixes are guaranteed to be non-empty. Case D.X works by induction. □

Note that because $f(\mathcal{W})$ and $\texttt{start}(\mathcal{W})$ determine $R(\mathcal{W})$, $E(\mathcal{W})$, and the length of $\mathcal{W}$ uniquely, Lemma 9 implies that $\phi(\mathcal{W})$ is a barren walkage (because $\mathcal{W}$ is barren). It remains to prove that $\phi(\mathcal{W})$ is proper, and to prove that $\phi(\mathcal{W})$ is proper the only remaining thing to prove is that $\phi(\mathcal{W})$ does not contain labeled digons (Lemma 11). In particular, the property $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$ guarantees that the ending vertices of $\phi(\mathcal{W})$ are distinct, and $f(\phi(\mathcal{W})) = f(\mathcal{W})$, which implies $R(\phi(\mathcal{W})) = R(\mathcal{W})$ guarantees that the labeled vertices of $\phi(\mathcal{W})$ have different colors (because $\mathcal{W}$ is proper).

We will make use of the following lemma that follows directly from Lemma 9.

Lemma 10. *If a vertex occurs exactly once in $\mathcal{W}$ and this occurrence is a digon, then this vertex also occurs exactly once in $\phi(\mathcal{W})$ and this occurrence is also a digon with the same adjacent vertices.*

Proof. Suppose that a vertex v occurs exactly once in $\mathcal{W}$ and this occurrence is a digon. Therefore, it cannot be a starting or ending vertex in $\mathcal{W}$. By Lemma 9, it holds that $\texttt{start}(\phi(\mathcal{W})) = \texttt{start}(\mathcal{W})$ and $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$ and therefore v cannot be a starting or ending vertex neither in $\phi(\mathcal{W})$. Then, since by Lemma 9, we have that $f(\phi(\mathcal{W})) = f(\mathcal{W})$, implying $E(\phi(\mathcal{W})) = E(\mathcal{W})$, the vertex v must have the exactly same adjacent vertices in $\phi(\mathcal{W})$ as in $\mathcal{W}$. □

Then we prove that $\phi(\mathcal{W})$ has no labeled digons.

Lemma 11. *Let $\mathcal{W}$ be a proper barren labeled walkage. The labeled walkage $\phi(\mathcal{W})$ has no labeled digons.*

Proof. We prove the lemma by induction on the length of the walkage $\mathcal{W}$.

Case A.1. In this case, as $\phi(W^1[2,\ell_1], W^2, \ldots, W^p)$ has no labeled digons by induction, the only potential place for a labeled digon could be index 2 of W^1. However, because v_1^1 occurs only once in $\mathcal{W}$ and therefore does not occur in $\phi(W^1[2,\ell_1], W^2, \ldots, W^p)$, we have that the index 2 of W^1 cannot become a digon.

Case A.2. Trivially by induction.

Case B. Here, by the definition of v_a^i and v_b^j in case B, both v_a^i and v_b^j are unlabeled and $v_a^i = v_b^j$ holds, so if $\mathcal{W} \leftrightarrow_{a,b}^{i,j}$ would contain a labeled digon then also $\mathcal{W}$ would.

Case C.1.a. Here, the indices 1 and b of W^1 are not labeled so they cannot become labeled digons. For indices in $[2, b-1]$, note that if $i \in [2, b-1]$ would be a labeled digon in $W^1[\overleftarrow{2, b-1}]$, then $b+1-i$ would have been a labeled digon in W^1.

Case C.1.b. Potential places for labeled digons here are incides b and $b+1$ at W^1. However, b is not labeled so no labeled digon can be at b, and because v_b^1 does not occur in $(W^1[b+1,\ell_1], W^2, \ldots, W^p)$, it cannot occur in $\phi(W^1[b+1,\ell_1], W^2, \ldots, W^p)$ and therefore $b+1$ cannot become labeled digon.

Case C.1.c. Trivially by induction.

Case C.2. The index 1 of W^1 is not digon by definition of digon, and the index b is not digon by definition of this case, so no labeled digons are created.

Case C.3. Here, the index 1 cannot be a digon by definition, and the index a of W^1 has $r_a^1 = 0$ by case C.2, so they cannot become labeled digons. For indices in $[2, a-1]$, the same argument as in case C.1.a applies.

Case C.4.a. Neither index a nor b is labeled so they cannot become labeled digons, and for indices in $[a+1, b-1]$ the same argument as in case C.1.a applies.

Case C.4.b. Same argument as in C.1.b, and using that by case C.2, it holds that $r_b^1 = 0$.

Case C.X. Here, the index a of W^1 cannot become a labeled digon because it is not labeled. For the index $a+1$ the argument is more complicated: First note that the vertex $v_a^1 = v_b^1 (= v_1^1)$ occurs only once in $(W^1[a+1, \ell_1], W^2, \ldots, W^p)$ (as v_b^1). Also, by case C.2, b is a digon in W^1. Therefore, by Lemma 10, the vertex $v_b^1 = v_a^1$ occurs in $\phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$ only once, and this occurrence is a digon adjacent with vertex $v_{b-1}^1 = v_{b+1}^1$. Because by Lemma 9 ϕ preserves the starting vertices, and by case C.4 it holds that $v_{a+1}^1 \neq v_{b-1}^1$, it holds that the occurrence of v_b^1 in $\phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$ cannot be as the second vertex of the first walk. Therefore, the index $a+1$ of the first walk cannot be a labeled digon in $W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$.

Case D.1. As $r_1^1 = r_b^i = 0$ and $v_1^1 = v_b^i$, it holds that if $\mathcal{W} \leftrightarrow_{1,b}^{1,i}$ would contain a labeled digon then also $\mathcal{W}$ would.

Case D.2. The index 1 of W^1 cannot become a digon by definition of digon. The index b of W^i cannot become a digon by definition of this case.

Case D.3. Because $v_1^1 = v_a^i$, in this case if $\mathcal{W} \leftrightarrow_{2,c+1}^{1,i}$ would have a labeled digon at index 2 of W^1 or at index $c+1$ of W^i, then this same labeled digon would have existed also in $\mathcal{W}$. Then, no labeled digon can be created to index 1 of W^1 by definition of digon or to index c of W^i because $r_1^1 \neq 0$ and therefore $r_c^i = 0$ by case D.2.

Case D.4. The index 1 of W^1 cannot become a digon by definition, and the index a cannot become a labeled digon because it is not labeled because the index 1 is labeled. For indices in $[2, a-1]$, the same argument as in case C.1.a applies.

Case D.5. Trivially by induction.

Case D.6. In this case, by cases D.1 and D.2 we have that $r_1^1 \neq 0$ and $r_b^i = 0$. The former implies that $r_a^1 = 0$. Therefore, the index a of W^1 nor the index b of W^i cannot become labeled digons because they cannot become labeled. Then, if the index $a+1$ of W^1 would be a labeled digon in $\mathcal{W} \leftrightarrow_{a+1,b+1}^{1,i}$, then the index $b+1$ of W^i would have been a labeled digon in $\mathcal{W}$ because $v_a^1 = v_b^i$ (and symmetrically for $b+1$ of W^i).

Case D.X. Here, a similar argument as in C.X works: By Lemma 10 we know that v_b^i occurs only as a digon surrounded by $v_{b-1}^i = v_{b+1}^i$ in $\phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$. Therefore, because ϕ maintains the starting vertices and $v_{a+1}^1 \neq v_{b+1}^i$, it holds that the index $a+1$ of W^1 cannot be a digon in $W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$. The index a of W^1 cannot become a labeled digon because it is not labeled. □

The function ϕ is an involution. Now we have shown that ϕ is a function $\phi : \mathcal{B} \to \mathcal{B}$, where $\mathcal{B}$ is the set of all barren proper labeled walkages, and that $f(\phi(\mathcal{W})) = f(\mathcal{W})$, $\mathcal{T}(\phi(\mathcal{W})) = \mathcal{T}(\mathcal{W})$, and $\texttt{start}(\phi(\mathcal{W})) = \texttt{start}(\mathcal{W})$ hold. Next we show that ϕ is an involution on $\mathcal{B}$, i.e., $\phi(\phi(\mathcal{W})) = \mathcal{W}$ holds.

Lemma 12. *For any proper barren labeled walkage $\mathcal{W}$ it holds that $\phi(\phi(\mathcal{W})) = \mathcal{W}$.*

Proof. We use induction on the length of walkage $\mathcal{W}$. The structure of the proof is to show that in all cases except C.X and D.X, the walk $\phi(\mathcal{W})$ goes to the same case of Definition 3 as $\mathcal{W}$. Then, the cases C.X and D.X are treated together.

Case A.1. In both $\mathcal{W}$ and $\phi(\mathcal{W})$ it holds that the v_1^1 occurs only once in the walkage and ϕ does not change the first vertex, so $\phi(\phi(\mathcal{W})) = \mathcal{W}$ holds by induction.

Case A.2. In this case the first walk of $\mathcal{W}$ is the same as the first walk of $\phi(\mathcal{W})$, so $\phi(\phi(\mathcal{W})) = \mathcal{W}$ holds by induction.

Case B. In this case, the function ϕ preserves the set of walks in which v_1^1 occurs, and moreover preserves the walk in which v_1^1 occurs as labeled (if it occurs as labeled in any walk). Therefore the walkage $\phi(\mathcal{W})$ also goes to case B, and in case B the indices i, j selected for $\phi(\mathcal{W})$ are the same as selected for $\mathcal{W}$. The suffix swap operation also does not change the indices of the first occurrences of v_1^1 in W^i and W^j, so the indices a and b selected are the same. Then, the lemma follows by observing that $\mathcal{W} \leftrightarrow_{a,b}^{i,j} \leftrightarrow_{a,b}^{i,j} = \mathcal{W}$.

At this point, let us observe that $\mathcal{W}$ goes to cases C.1-4 if and only if $\phi(\mathcal{W})$ goes to cases C.1-4. This is because all of these cases maintain that v_1^1 occurs only in the walk W^1, but multiple times in the walk W^1. In particular, in the recursive cases this is maintained by the fact that v_1^1 does not occur in the recursive argument.

We also observe that $\mathcal{W}$ goes to cases D.1-6 if and only if $\phi(\mathcal{W})$ goes to cases D.1-6. This is because all of these cases maintain that v_1^1 occurs in exactly two different walks. In the cases D.1-4 and D.6 this is easy to observe since these are not recursive, and in the case D.5 this follows from the fact that the walk W^1 is not changed and that in this case v_1^1 occurs exactly once in $(W^2, \ldots, W^p)$.

Case C.1. Note that going to the cases C.1.a, C.1.b, and C.1.c depends only on the last occurrence b of v_1^1, and the labels r_1^1 and r_b^1. None of these cases change these, so $\phi(\phi(\mathcal{W})) = \mathcal{W}$ holds in case C.1.a because $W^1[\overleftarrow{2, b-1}] = W^1$ and reversing a subwalk does not change whether it is a palindrome, and by induction in cases C.1.b and C.1.c.

Case C.2. The order of the vertices in the walks and the fact that exactly one of r_1^1 and r_b^1 is labeled is maintained, so if $\mathcal{W}$ goes to case C.2 then also $\phi(\mathcal{W})$ goes to C.2. Then, $\phi(\phi(\mathcal{W})) = \mathcal{W}$ because $\mathcal{W} \frown_{1,b}^{1,1} \frown_{1,b}^{1,1} = \mathcal{W}$.

Case C.3. It holds that $a < b$, so therefore reversing $W^1[2, a-1]$ does not change the fact that b is a digon in W^1. It also does not change the index a of the second last occurrence of v_1^1, nor the index b of the last occurrence of v_1^1, nor the fact that $r_1^1 \neq 0$, nor the fact that $W^1[2, a-1]$ is not a palindrome.

Case C.4.a. Because $v_{a+1}^1 = v_{b-1}^1$, reversing $W^1[a+1, b-1]$ does not change the fact that b is a digon in W^1. Reversing $W^1[a+1, b-1]$ also does not change the index a of the second last occurrence of v_1^1, nor the index b of the last occurrence of v_1^1, nor the fact that $r_1^1 \neq 0$, nor the fact that $W^1[2, a-1]$ is a palindrome.

Case C.4.b. Going to the case C.4.b depends only on the subwalk $W^1[1, b]$ and on the vertex with index $b+1$ in W^1 (whether the index b is a digon). Clearly, ϕ does not change the subwalk $W^1[1, b]$ in this case. The vertex with index $b+1$ is not changed because by Lemma 9 the starting vertices are preserved by $\phi(W^1[b+1, \ell_1], W^2, \ldots, W^p)$, so the lemma holds by induction.

Before moving to the cases C.X and D.X, we handle the cases D.1-D.5.

Case D.1. In this case, the operation $\mathcal{W} \leftrightarrow_{1,b}^{1,i}$ does not change the two walks in which v_1^1 occurs, nor it changes the fact that the first occurrence of v_1^1 in W^i is at index b, nor that $r_1^1 = r_b^i = 0$. The lemma follows from the fact that $\mathcal{W} \leftrightarrow_{1,b}^{1,i} \leftrightarrow_{1,b}^{1,i} = \mathcal{W}$.

Case D.2. This case does not change the vertices of the walks, so it is maintained that v_1^1 occurs only in W^1 and W^i. It also does not change the fact that exactly one of r_1^1 and r_b^i is labeled or the fact the index b is not a digon in W^i, so the lemma follows from the fact that $\mathcal{W} \frown_{1,b}^{1,i} \frown_{1,b}^{1,i} = \mathcal{W}$.

Case D.3. This case does not change the index b of the first occurrence of v_1^1 in W^i or the index c of the second occurrence of v_1^1 in W^i, and neither does it change the fact that $r_1^1 \neq 0$. Because $c > b+1$, it also does not change that the index b is a digon in W^i. The lemma follows from the fact that $\mathcal{W} \leftrightarrow_{2,c+1}^{1,i} \leftrightarrow_{2,c+1}^{1,i} = \mathcal{W}$.

Case D.4. This case does not change the index b of the first occurrence of v_1^1 in W^i, nor that $r_1^1 \neq 0$, nor that b is a digon in W^i, nor that v_1^1 occurs only once in W^i. It also does not change the index a of the last occurrence of v_1^1 in W^i, or the fact that $W^1[2, a-1]$ is a palindrome so the lemma holds.

Case D.5. The vertex v_1^1 occurs exactly once in $(W^2, \ldots, W^p)$, so it is also maintained that v_1^1 occurs in exactly two walks. The walk W^1 is not changed, so it is maintained that $r_1^1 \neq 0$ and therefore $\phi(\mathcal{W})$ does not go to case D.1. By Lemma 10 we have that $\phi(\mathcal{W})$ does not go to case D.2, and again because v_1^1 occurs only once in $(W^2, \ldots, W^p)$ we have that $\phi(\mathcal{W})$ does not go to case D.3. As W^1 is not changed we have that $\phi(\mathcal{W})$ does not go to case D.4. Then, as case D.5 does not change the walk W^1, it is maintained that $a = \ell_1$, so $\phi(\mathcal{W})$ goes to case D.5.

Case D.6. This case does not change that $r_1^1 \neq 0$, so $\phi(\mathcal{W})$ does not go to case D.1. It also does not change the index b of the first occurrence of v_1^1 in W^i, and because $v_{a+1}^1 = v_{b+1}^i$, it does not change that b is a digon in W^i, and therefore $\phi(\mathcal{W})$ does not go to case D.2. Because a is the last occurrence of v_1^1 in W^1, it also does not change that v_1^1 occurs in W^i only once, so $\phi(\mathcal{W})$ does not go to case D.3. It also does not change the index a of the last occurrence of v_1^1 in W^1 or the subwalk $W^1[2, a-1]$, so $\phi(\mathcal{W})$ does not go to cases D.4. or D.5. Therefore, $\phi(\mathcal{W})$ goes to case D.6 with the same values of a, b, and i, so $\phi(\phi(\mathcal{W})) = \mathcal{W}$ holds because $\mathcal{W} \leftrightarrow^{1,i}_{a+1,b+1} \leftrightarrow^{1,i}_{a+1,b+1} = \mathcal{W}$.

Case C.X. We aim to prove that if $\mathcal{W}$ goes to case C.X, then $\phi(\mathcal{W})$ also goes to case C.X or to the case D.X, with the same value of the index a, and therefore $\phi(\phi(\mathcal{W})) = \mathcal{W}$ will hold by induction, as in both cases ϕ is defined as $\phi(\mathcal{W}) = W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$. First, it is maintained that v_1^1 occurs more than once, but in at most two walks, because v_1^1 occurs only once in $(W^1[a+1, \ell_1], W^2, \ldots, W^p)$. Therefore, $\phi(\mathcal{W})$ does not go to case A or B.

Suppose that v_1^1 occurs in $\phi(\mathcal{W})$ only in the walk W^1, i.e., goes to case C. We will show that $\phi(\mathcal{W})$ goes to case C.X. It is maintained that $r_1^1 \neq 0$, so $\phi(\mathcal{W})$ does not go to case C.1. Then, by Lemma 10 it is maintained that the last occurrence of v_1^1 must be a digon so $\phi(\mathcal{W})$ does not go to case C.2. Also, this case does not change the index a of the second last occurrence of v_1^1 nor the walk $W^1[2, a-1]$, so it is maintained that $W^1[2, a-1]$ is a palindrome and therefore $\phi(\mathcal{W})$ does not go to case C.3. Then, to argue that $\phi(\mathcal{W})$ does not go to case C.4, observe that because ϕ maintains the starting vertex, the vertex at the position $a+1$ is maintained. The vertices around the digon at the last occurrence of v_1^1 are maintained by Lemma 10, so therefore if $\mathcal{W}$ does not go to case C.4 then also $\phi(\mathcal{W})$ does not go to case C.4. Therefore $\phi(\mathcal{W})$ goes to case C.X, and as the walk $W^1[1, a]$ is maintained, it goes to this case with the same value of a, so the lemma holds by induction.

Then, suppose that v_1^1 occurs in $\phi(\mathcal{W})$ in two walks W^1 and W^i, i.e., goes to case D. We will show that $\phi(\mathcal{W})$ goes to case D.X. It is maintained that $r_1^1 \neq 0$, so $\phi(\mathcal{W})$ does not go to case D.1. Then, by Lemma 10 it must be that v_1^1 occurs in W^i only once and as a digon, and therefore $\phi(\mathcal{W})$ does not go to case D.2 or D.3. Now, it will hold that the index a of the last occurrence of v_1^1 in the walk W^1 of $\phi(\mathcal{W})$ is the same as the index a of the last occurrence of v_1^1 in the walk W^1 of $\mathcal{W}$. Therefore, the subwalk $W^1[1, a]$ will be the same in $\mathcal{W}$ and $\phi(\mathcal{W})$, and therefore $\phi(\mathcal{W})$ will not go to case D.4 because $\mathcal{W}$ did not go to case C.3. Then, because ϕ cannot turn a non-empty walk into an empty walk, it is maintained that the length of W^1 is more than a, so $\phi(\mathcal{W})$ cannot go to case D.5. Then, for case D.6 we again note that ϕ maintains the vertex at position $a+1$, and that the digon around the occurrence of v_1^1 outside of $W^1[1, a]$ is maintained by Lemma 10. Therefore, $\phi(\mathcal{W})$ goes to case D.X with the same value of a, so the lemma holds by induction.

Case D.X. We will show that if $\mathcal{W}$ goes to case D.X, then $\phi(\mathcal{W})$ also goes to case D.X or to the case C.X, with the same value of the index a, and therefore $\phi(\phi(\mathcal{W})) = \mathcal{W}$ will hold by induction, as in both cases ϕ is defined as $\phi(\mathcal{W}) = W^1[1, a] \diamond \phi(W^1[a+1, \ell_1], W^2, \ldots, W^p)$. First, it is maintained

that v_1^1 occurs more than once, so therefore $\phi(\mathcal{W})$ does not go to case A. Then, as v_1^1 occurs only once in $(W^1[a+1,\ell_1], W^2, \ldots, W^p)$, it can occur in at most two walks in $\phi(\mathcal{W})$, so $\phi(\mathcal{W})$ cannot go to case B.

Suppose that v_1^1 occurs in $\phi(\mathcal{W})$ only in the walk W^1, i.e., goes to case C. We will show that $\phi(\mathcal{W})$ goes to case C.X. It is maintained that $r_1^1 \neq 0$, so $\phi(\mathcal{W})$ does not go to case C.1. Then, because $W^1[1,a]$ contains all other occurrences of v_1^1 in $\mathcal{W}$ except the occurrence in W^i, it must be that now the last occurrence of v_1^1 in W^1 of $\phi(\mathcal{W})$ corresponds to the occurrence of v_1^1 in W^i of $\mathcal{W}$, in particular, by Lemma 10 the last occurrence of v_1^1 in W^1 of $\phi(\mathcal{W})$ must be a digon, and therefore $\phi(\mathcal{W})$ does not go to case C.2. By the same reasoning, it also must be that the the index a of the last occurrence of v_1^1 in W^1 of $\mathcal{W}$ is the same as the index a of the second last occurrence of v_1^1 in W^1 of $\phi(\mathcal{W})$, and therefore the subwalk $W^1[1,a]$ of $\phi(\mathcal{W})$ in case C is the same as the subwalk $W^1[1,a]$ of $\mathcal{W}$ in case D. Then, it follows that because $\mathcal{W}$ did not go to case D.4, $\phi(\mathcal{W})$ does not go to case C.3. Then, as the case D.X does not change the vertex at the index $a+1$ of W^1, it holds that the vertex at the index $a+1$ of W^1 is the same in $\mathcal{W}$ and $\phi(\mathcal{W})$. Also, by Lemma 10 the vertices around the digon of the last occurrence of v_1^1 are the same in $\mathcal{W}$ and $\phi(\mathcal{W})$, so $\phi(\mathcal{W})$ does not go to case C.4 because $\mathcal{W}$ did not go to case D.5. Therefore, $\phi(\mathcal{W})$ must go to case C.X, and we already reasoned that the index a is the same as for $\mathcal{W}$ in the case D.X, so the lemma holds by induction.

Suppose that v_1^1 occurs in $\phi(\mathcal{W})$ in two different walks, i.e., goes to case D. We will show that $\phi(\mathcal{W})$ goes to case D.X. First, it is maintained that $r_1^1 \neq 0$, so $\phi(\mathcal{W})$ does not go to case D.1. Then, we note that the walk W^i that contains the other occurrence of v_1^1 may be different for $\phi(\mathcal{W})$ than $\mathcal{W}$. However, it is maintained that v_1^1 occurs only once outside of W^1, and by Lemma 10 that the other occurrence is a digon and the vertices around this digon are maintained. Therefore, $\phi(\mathcal{W})$ does not go to case D.2, nor to the case D.3. Now, the index a of the last occurrence of v_1^1 in W^1 will be the same for $\phi(\mathcal{W})$ and $\mathcal{W}$ because $\phi(\mathcal{W})$ does not change the subwalk $W^1[1,a]$. Therefore, it is maintained that $W^1[2,a-1]$ is a palindrome, and therefore $\phi(\mathcal{W})$ does not go to case D.4. Then, because ϕ cannot turn a non-empty walk into an empty walk, it is maintained that the length of W^1 is more than a, so $\phi(\mathcal{W})$ cannot go to case D.5. Then, by the start vertex property of ϕ, the vertex at index $a+1$ of W^1 is also maintained, and by Lemma 10 the vertices around the digon of the other occurrence of v_1^1 are maintained, so $\phi(\mathcal{W})$ does not go to case D.6. Therefore, $\phi(\mathcal{W})$ goes to the case D.X, with the same value of a, and therefore the lemma holds by induction. □

Finally, we show that ϕ is fixed-point-free.

Lemma 13. *For any proper barren labeled walkage $\mathcal{W}$ it holds that $\phi(\mathcal{W}) \neq \mathcal{W}$.*

Proof. We prove this by induction on the length ℓ of the walkage. In the recursive cases A.1, A.2, C.1.b, C.1.c, C.4.b, C.X, D.5, and D.X this holds directly by induction. In cases B, D.1, D.3, and D.6, ϕ changes the suffixes of two walks, and at least one of the suffixes is non-empty. Because $\mathcal{W}$ is proper, the ending vertices of all walks in $\mathcal{W}$ are different, so the ending vertex of at least one of the walks involved in the suffix swap is changed (in fact the ending vertices of both of the walks change, but it is not necessary for this proof). In cases C.1.a, C.3, C.4.a, and D.4, ϕ reverses a non-palindromic subwalk so $\phi(\mathcal{W}) \neq \mathcal{W}$. In cases C.2 and D.2 ϕ changes a label from one position to another, so $\phi(\mathcal{W}) \neq \mathcal{W}$. □

This completes the proof of Lemma 5.

5 From colored graphs to frameworks

In this section, we extend our results from weighted colored graphs to weighted frameworks, in particular, we prove Theorem 5 (recall that Theorem 5 implies Theorem 3), and then discuss even further extensions to frameworks where the matroid is not necessarily represented over a finite field.

5.1 Frameworks

We recall definitions related to frameworks.

Matroids. We refer to the textbook of Oxley [41] for the introduction to Matroid Theory.

Definition 4. *A pair $M = (V, \mathcal{I})$, where V is a* ground set *and $\mathcal{I}$ is a family of subsets of V, called* independent sets of M*, is a* matroid *if it satisfies the following conditions, called* independence axioms*:*

(I1) $\emptyset \in \mathcal{I}$,

(I2) *if $X \subseteq Y$ and $Y \in \mathcal{I}$ then $X \in \mathcal{I}$,*

(I3) *if $X, Y \in \mathcal{I}$ and $|X| < |Y|$, then there is $v \in Y \setminus X$ such that $X \cup \{v\} \in \mathcal{I}$.*

An inclusion maximal set of $\mathcal{I}$ is called a *base*. We use $V(M)$ and $\mathcal{I}(M)$ to denote the ground set and the family of independent sets of M, respectively.

Let $M = (V, \mathcal{I})$ be a matroid. We use 2^V to denote the set of all subsets of V. A function $r\colon 2^V \to \mathbb{Z}_{\geq 0}$ such that for every $X \subseteq V$,

$$r(X) = \max\{|Y| \colon Y \subseteq X \text{ and } Y \in \mathcal{I}\}$$

is called the *rank function* of M. The *rank of* M, denoted $r(M)$, is $r(V)$; equivalently, the rank of M is the size of any base of M. A matroid $M' = (V, \mathcal{I}')$ is a *k-truncation* of $M = (V, \mathcal{I})$ if for every $X \subseteq V$, $X \in \mathcal{I}'$ if and only if $X \in \mathcal{I}$ and $|X| \leq k$.

We work with several particular types of matroids. A *uniform* matroid is defined by the ground set V and its rank r; every subset S of V of size at most r is independent. *Partition* matroids are the matroids that can be written as disjoint sums of uniform matroids. *Transversal* matroids arise from graphs. For a bipartite graph $G = (U \cup B, E)$ with all edges between U and B, we can define a matroid $M = (V, \mathcal{I})$ such that a set $S \subseteq V$ is independent if there exists a matching in G such that every vertex in S is an endpoint of a matching edge.

Matroid representations. Let $M = (V, \mathcal{I})$ be a matroid and let $\mathbb{F}$ be a field. An $r \times n$-matrix A is a *representation of M over $\mathbb{F}$* if there is a bijective correspondence f between V and the set of columns of A such that for every $X \subseteq V$, $X \in \mathcal{I}$ if and only if the set of columns $f(X)$ consists of linearly independent vectors of $\mathbb{F}^r$. Equivalently, A is a representation of M if M is isomorphic to the *column* matroid of A, that is, the matroid whose ground set is the set of columns of the matrix and the independence of a set of columns is defined as the linear independence. If M has a such a representation, then M is *representable* over $\mathbb{F}$ and it is also said M is a *linear* (or *$\mathbb{F}$-linear*) matroid. We can assume that the number of rows $r = r(M)$ for a matrix representing M [37].

Whenever we consider a linear matroid, it is assumed that its representation is given and the size of M is $\|M\| = \|A\|$, that is, the bit-length of the representation matrix. Notice that given a representation of a matroid, deciding whether a set is independent demands a polynomial number of field operations. In particular, if the considered field is a finite or is the field of rationals, we can verify independence in time that is a polynomial in $\|M\|$.

Frameworks. A framework is a pair (G, M), where $M = (V, \mathcal{I})$ is a matroid whose ground set is the set of vertices of G, i.e., $V(M) = V(G)$. A weighted framework is a triple $(G, M, \mathtt{we})$, where (G, M) is a framework and $\mathtt{we} : V(G) \to \mathbb{Z}_{\geq 1}$ is a weight function. An (S, T)-linkage $\mathcal{P}$ in a weighted framework $(G, M, \mathtt{we})$ is (k, w)-ranked if $V(\mathcal{P})$ contains a set $X \subseteq V(\mathcal{P})$ with $X \in \mathcal{I}$, size $|X| = k$, and weight $\mathtt{we}(X) = w$. When discussing algorithms for (weighted) frameworks, we explicitly specify how M is represented.

5.2 From colored graphs to frameworks

In this section we prove Theorem 5. We reduce the more general cases of matroids to the case of Theorem 4.

We start by giving our algorithm for the special case when the rank of M is bounded by k, in particular when M is represented as a $k \times n$ matrix.

Lemma 14. *There is a randomized algorithm, that given a weighted framework $(G, M, \mathtt{we})$, where G is an n-vertex graph and M is represented as a $k \times n$ matrix over a finite field of order q, sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, in time $2^{p+\mathcal{O}(k^2 \log q)} n^{\mathcal{O}(1)} w$ either finds a (k, w)-ranked (S, T)-linkage of order p and of minimum total length, or determines that $(G, M, \mathtt{we})$ has no (k, w)-ranked (S, T)-linkages of order p.*

Proof. The matrix has at most q^k distinct column vectors so we can guess the k column vectors forming the independent set X of size k that we are looking for with at most q^{k^2} guesses. By inserting $|S|$ new vertices with neighborhoods equal to S, all-zero column vector, and weight 1, we can assume that the vectors of the starting vertices S will never correspond to the guessed basis. Then, we $k+1$-color the graph, assigning the color $k+1$ to the vertices of the set S and other vertices whose column vectors are not in the guessed basis, and the colors $[k]$ to the other vertices according to which of the k guessed column vectors they correspond to. We also assign the weight of all vertices whose column vector is not in the guessed basis to be 1.

Then, $(G, M, \mathtt{we})$ has a (k, w)-ranked (S, T)-linkage of order p if and only if it has a $(k+1, w+1)$-colored (S, T)-linkage of order p. In particular, the extra color $k+1$ contributes weight one and one color more, and the selected set $X \subseteq V(\mathcal{P})$ without the extra color must correspond to an independent set of M. Therefore, we get an algorithm with time complexity $q^{k^2} 2^{p+k} n^{\mathcal{O}(1)} w = 2^{p+\mathcal{O}(k^2 \log q)} n^{\mathcal{O}(1)} w$. □

By extending Lemma 14 to matrices with a large number of rows by using randomized lossy truncation, we prove Theorem 5, which we restate here.

Theorem 5. *There is a randomized algorithm that given a weighted framework $(G, M, \mathtt{we})$, where G is an n-vertex graph and M is represented as a matrix over a finite field of order q, sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, in time $2^{p+\mathcal{O}(k^2 \log(q+k))} n^{\mathcal{O}(1)} w$ either finds a (k, w)-ranked (S, T)-linkage of order p and of minimum total length, or determines that $(G, M, \mathtt{we})$ has no (k, w)-ranked (S, T)-linkages of order p.*

Proof. Let A be a $r \times n$ matrix representing M. Our goal is to obtain a "lossy" representation of the k-truncation of M as a $k \times n$ matrix over a field of order $\mathcal{O}(q+k^2)$. In particular, a representation so that any independent set of M of size k is independent in the representation with probability $\geq 1/2$, and any dependent set of M is dependent in the representation. Then, we obtain the algorithm by applying the algorithm of Lemma 14. Note that $2^{p+\mathcal{O}(k^2 \log(q+k^2))} n^{\mathcal{O}(1)} w = 2^{p+\mathcal{O}(k^2 \log(q+k))} n^{\mathcal{O}(1)} w$.

We use two techniques from [37], increasing the order of the field and truncation. First, we make sure that the order of the field is at least $2k$ by choosing the least integer i such that $q^i \geq 2k$,

and going to the field of order q^i, as detailed in Proposition 3.2 of [37]. Now, we can assume that A is over a field of order at least $2k$ and at most $\mathcal{O}(q+k^2)$.

Then, we truncate the matroid by multiplying the matrix A by a random $k \times r$ matrix R, in particular we claim that the $k \times n$ matrix $B = RA$ is now the desired representation of the k-truncation of M. The analysis here is the same as in Proposition 3.7 of [37], but with a smaller field. In particular, let us consider a subset U of the ground set of M, and let A_0 be the $r \times |U|$ submatrix of A corresponding to S. Now, in B, the $k \times |U|$ submatrix corresponding to S will be the matrix $B_0 = RA_0$. The rank of B_0 is at most the rank of A_0, so if U is dependent in M it will be dependent in the representation by B. Then, assume that U is an independent set of M and $|U| = k$. Now, $\det RA_0$ can be considered as a degree-k non-zero polynomial whose variables are the kr random entries of R. Therefore, by Lemma 1, the probability that $\det RA_0 = 0$ is at most $k/2k$. □

With minor adjustments, Theorem 5 can be adapted for frameworks with matroids that are in general not representable over a finite field of small order. For example, uniform matroids, and more generally transversal matroids, are representable over a finite field, but the field of representation must be large enough. We first show how Theorem 1 can be applied in the case of transversal matroids.

Theorem 6. *There is a randomized algorithm that given a weighted framework $(G, M, \mathtt{we})$, where G is an n-vertex graph and M is a transversal matroid represented by the corresponding bipartite graph, sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, in time $2^{p+\mathcal{O}(k^2 \log k)} n^{\mathcal{O}(1)} w$ either finds a (k,w)-ranked (S,T)-linkage of order p and of minimum total length, or determines that $(G, M, \mathtt{we})$ has no (k,w)-ranked (S,T)-linkages of order p.*

Proof. We will construct a representation of the transversal matroid as a linear matroid over a finite field of order $\mathcal{O}(k)$, so that any independent set of M of size k is independent in the representation with probability $\geq 1/2$, and any dependent set of M is dependent in the representation. This yields the algorithm by then using Theorem 5.

Our construction is the same as the construction of [37], except by using a smaller field. We choose the least prime p with $p \geq 2k$ and work in the field of order p. Let the bipartition of the vertices of the bipartite graph be (A, B). We construct an $|B| \times |A|$ matrix, so that an entry of the matrix is a random element of the field if it corresponds to an edge, and zero otherwise. Now, the determinant of a submatrix is guaranteed to be zero if there is no corresponding matching, so any dependent set of M is dependent in the representation. Otherwise, the determinant of a $k \times k$ submatrix can be seen as a non-zero degree-k polynomial that was evaluated at a random point. Therefore, as $p \geq 2k$, by Lemma 1, the probability that it is non-zero is at least $1/2$. □

It is also possible to apply Theorem 5 in the situation when M is represented by an integer matrix over rationals with entries bounded by $n^{\mathcal{O}(k)}$.

Theorem 7. *There is a randomized algorithm that given a weighted framework $(G, M, \mathtt{we})$, where G is an n-vertex graph and M is represented as an integer matrix over rationals with entries bounded by $n^{\mathcal{O}(k)}$, sets of vertices $S, T \subseteq V(G)$, and integers p, k, w, in time $2^{p+\mathcal{O}(k^2 \log k)} n^{\mathcal{O}(1)} w$ either finds a (k,w)-ranked (S,T)-linkage of order p and of minimum total length, or determines that $(G, M, \mathtt{we})$ has no (k,w)-ranked (S,T)-linkages of order p.*

Proof. Let c be a constant so that the entries of the matrix are bounded by n^{ck}. We pick a random prime p among the first $2\log_2(k! n^{ck^2})$ primes, go to the finite field of order p by taking every entry modulo p, and then apply the algorithm of Theorem 1.

We first analyze the time complexity and then the correctness. By the prime number theorem, the prime p is bounded by

$$p = \mathcal{O}(\log k! n^{ck^2} \cdot \log\log k! n^{ck^2}) = \mathcal{O}(k^3 \log n \log\log n).$$

We can find such random prime in $n^{\mathcal{O}(1)}$ time by elementary methods. Then, the time complexity by using Theorem 1 will be $2^{\mathcal{O}(p+k^2 \log(k+k^3 \log n \log\log n)} n^{\mathcal{O}(1)}$. Denote $t(n) = \log n \log\log n$ and consider two cases. First, if $t(n) \leq k^5$, then the time complexity is $2^{\mathcal{O}(p+k^2 \log(k^8)} n^{\mathcal{O}(1)} = 2^{\mathcal{O}(p+k^2 \log k)} n^{\mathcal{O}(1)}$. Second, if $t(n) > k^5$, then the time complexity is $2^{\mathcal{O}(p+k^2 \log k^3 t(n))} n^{\mathcal{O}(1)} = 2^{\mathcal{O}(p+t(n)^{1/2} \log t(n))} n^{\mathcal{O}(1)} = 2^{\mathcal{O}(p)} \cdot 2^{\mathcal{O}(\log^{1/2} n \log^{\mathcal{O}(1)} \log n)} = 2^{\mathcal{O}(p)} n^{\mathcal{O}(1)}$.

Then, for the correctness we show that any dependent set of M is dependent in the representation and any independent set of M of size k is independent in the representation with probability $\geq 1/2$. Let A be a square submatrix of the original representation and A_p the corresponding submatrix in the presentation modulo p. Now, $\det A_p = \det A \mod p$. Therefore, all dependent sets stay dependent. Then, assume that A is a $k \times k$ submatrix corresponding to an independent set, i.e., $\det A \neq 0$. Now, the independent set can change into dependent only if $\det A$ is divisible by p. The value $\det A$ is bounded by $k! n^{ck^2}$, so there are at most $\log_2(k! n^{ck^2})$ primes dividing it. We chose p randomly among the first $2\log_2(k! n^{ck^2})$ primes, so with probability $\geq 1/2$ the prime p does not divide $\det A$. □

6 Deterministic algorithm for longest (S,T)-linkage

This section is dedicated to the proof of Theorem 2, which we restate next for convenience.

Theorem 2. *There is a deterministic algorithm that, given an n-vertex digraph G, two sets of vertices $S, T \subseteq V(G)$, an integer p, and an integer k, in time $p^{\mathcal{O}(kp)} n^{\mathcal{O}(1)}$ either returns a directed (S,T)-linkage of order p and of total length at least k, or determines that G has no directed (S,T)-linkage of order p and total length at least k.*

We start with showing the main combinatorial lemma behind the theorem. The lemma is illustrated in Figure 3.

Lemma 15. *Let G be a digraph and let $C_1, \ldots, C_q$ be disjoint sets in $V(G)$. For $s,t \in V(G)$ let $P_1, \ldots, P_q$ be internally-disjoint (s,t)-paths. For each $i \in [q]$, let $v_i \in V(P_i)$ be such that the suffix of P_i starting from v_i lies inside C_i, except for t. For $i \in [q]$, let Q_i be a path from v_i to t with all internal vertices in C_i. Then there exist internally-disjoint (s,t)-paths $P'_1, \ldots, P'_q$ such that P'_i is either (i) P_i or (ii) a composition of a prefix of P_i not containing any vertices of P_i beyond v_i, and a suffix of Q_j for some $j \in [q]$, and there is at least one path of type (ii) among $P'_1, \ldots, P'_q$.*

Proof. For each $i \in [q]$, denote the subpath of P_i from s to v_i by $P_i^{\rightarrow}$, and from v_i to t by $P_i^{\leftarrow}$. First, assume there exists $i \in [q]$ such that Q_i does not share a common internal vertex with any $P_j^{\rightarrow}$, $j \in [q]$. In this case, the solution is immediate: for each $j \neq i$, set $P'_j = P_j$, and set $P'_i = P_i^{\rightarrow} \circ Q_i$. The path $P_i^{\rightarrow}$ does not intersect any other P'_j internally since $P_1, \ldots, P_q$ are internally-disjoint, and by the assumption Q_i internally intersects neither $P_i^{\rightarrow}$ nor any other P'_j. So for the remaining part of the proof we assume that for each Q_i there exists $j \in [q]$ such that Q_i and $P_j^{\rightarrow}$ share a common internal vertex.

We now show the statement by analyzing a certain token sliding game. Intuitively, we put a token on each of the paths $P_1, \ldots, P_q$, originally on the place of the first intersection between P_i and some Q_j (see Figure 3b). Then we slide the tokens further along the paths according to certain

rules, until the tokens reach a state where no rules can be applied (Figures 3c and 3d). Our goal is to show that in this case we obtain the desired paths $P'_1, \ldots, P'_q$.

More formally, we define a *state* S as a tuple $(t_1, \ldots, t_q)$, where $t_i \in V(P_i)$ for each $i \in [q]$. The original state $S^1 = (t_1^1, \ldots, t_q^1)$ is defined as follows: for each $i \in [q]$, t_i^1 is the closest to s vertex along P_i that belongs to Q_j for some $j \in [q]$. The game then proceeds iteratively, constructing the state S^{h+1} from S^h for each h starting from $h = 1$ by applying one of the following rules. For $j \in [q]$, we shall refer to a path Q_j as *active* if $t_j^h \neq t$.

Clear Let $t_i^h = v_i$ for some $i \in [q]$. Then set $t_i^{h+1} = t$, and for each $j \in [q]$ such that $t_j^h \in V(Q_i) \setminus \{t\}$, set t_j^h to be the next vertex along the path P_j that belongs to an active $Q_{j'}$ for some $j' \in [q]$, here Q_i is not considered active. For all remaining $j \in [q]$, set $t_j^{h+1} = t_j^h$.

Push Let i and i', $i \neq i' \in [q]$, be such that both t_i^h and $t_{i'}^h$ belong to $Q_j - \{t\}$ for some $j \in [q]$; additionally, let t_i^h be the farthest of two from t along Q_j. Set t_i^h to be the next vertex along the path P_i that belongs to an active $Q_{j'}$ for some $j' \in [q]$. For all $i'' \in [q]$, $i'' \neq i$, set $t_{i''}^{h+1} = t_{i''}^h$.

As long as there is a possibility, a **Clear** rule is applied; **Push** is only applied if no **Clear** is available. If there are several options for applying the same rule, ties are breaking arbitrarily. We observe that every application of each rule moves at least one of the state vertices further along its respective path, and these vertices are never moved back. Thus, after a finite number of steps we reach a state where neither of the rules is applicable. Denote this state by S_i^T, our goal is to show the following.

Claim 6.1. *Let $S_i^T = (t_1^T, \ldots, t_q^T)$. There is $i \in [q]$ such that $t_i^T \neq t$.*

Before showing the proof of Claim 6.1 we make the following simple observation.

Claim 6.2. *For a state S^h, $h \in [T]$, for each $j \in [q]$, $t_j^h \in V(P_j)$, and t_j^h is either t or belongs to an active Q_i for some $i \in [q]$. Moreover, all of t_j^h that are not t, are distinct.*

Proof. By construction, the first part of the statement holds for the starting state S^1. Both rules either move vertices to an active Q_i further along its path, or directly to t. The second part follows immediately from the fact that the paths P_j are internally-disjoint. □

We first explain how Claim 6.1 implies the claim in the lemma. By Claim 6.1, there is at least one $i \in [q]$ such that $t_i^T \neq t$; let $I \subset [q]$ be the set of all indices with this property. By Claim 6.2, each vertex in $\{t_i^T\}_{i \in I}$ lies on an active Q_j for some $j \in [q]$, and these vertices are all distinct. Moreover, since the rule **Push** is not applicable, no two of these vertices share the same Q_j. Let $\pi : I \to [q]$ be the injection that maps $i \in I$ to the index j such that $t_i^T \in V(Q_j)$. We construct the desired family of paths as follows: for $i \in [q] \setminus I$, let P'_i be P_i, and for $j \in I$, let P'_j be a concatenation of the subpath of P_i from s to t_j^T (denoted $\widehat{P}_i$), and the subpath of $Q_{\pi(i)}$ from t_j^T to t (denoted $\widehat{Q}_i$). Observe also that for each $i \in I$, t_i^T is not v_i since the rule **Clear** is not applicable. Moreover, since I is non-empty, there is at least one path of type (ii) in the constructed family. It only remains to show that the paths $\{P'_i\}_{i \in [q]}$ are internally-disjoint.

For $i \in [q] \setminus I$, denote $\widehat{P}_i = P_i^{\rightarrow}$ and $\widehat{Q}_i = P_i^{\leftarrow}$. For $i \neq i' \in [q]$, s is the only intersection between $\widehat{P}_i$ and $\widehat{P}_{i'}$ since $\widehat{P}_i$ and $\widehat{P}_{i'}$ are proper prefixes of P_i and $P_{i'}$ respectively, and these paths are internally-disjoint by the assumption of the lemma. Observe that for each $i \in [q]$, the vertices of $\widehat{Q}_i$ except t lie in the set C_j, for some $j \in [q]$, and this correspondence between $\{\widehat{Q}_i\}_{i \in [q]}$ and

$\{C_j\}_{j\in[q]}$ is a bijection defined by π on I and by the identity permutation on $[q]\setminus I$. Since the sets $\{C_j\}_{j\in[q]}$ are disjoint, for any $i\neq i'\in[q]$, we get that the paths $\widehat{Q}_i$ and $\widehat{Q}_{i'}$ share the only common vertex t.

It remains to verify that for each $i\neq i'\in[q]$, $\widehat{P}_i$ shares no common vertices with $\widehat{Q}_{i'}$. For $i'\in[q]\setminus I$, $\widehat{Q}_{i'}=P_{i'}^{\leftarrow}$, which is a suffix of $P_{i'}$, and this path cannot intersect $\widehat{P}_i$ which is a prefix of P_i; thus in the following we assume $i'\in I$. Assume the contrary, then there is a vertex u on $\widehat{P}_i$ that belongs to $Q_{\pi(i')}$, and is located on $Q_{\pi(i')}$ closer to t than $t^T_{i'}$, which is the starting vertex of $\widehat{Q}_{i'}$. Observe that the rule **Clear** has never been applied to $Q_{\pi(i')}$, otherwise $Q_{\pi(i')}$ would not be active in S^T. Thus, there is a state S^h where $t^h_i=u$, since any application of the rules to $t^{h'}_i$ for any $h'\in[T]$ either leaves this vertex in place or moves it to the next vertex of P_i belonging to an active Q_j for some $j\in[q]$. Now observe that no application of the rule **Push** makes the closest vertex to t on $Q_{\pi(i')}-\{t\}$ among $\{t^h_j\}_{j\in[q]}$ farther, by definition of **Push**. However, we get that t^h_i is closer to t on $Q_{\pi(i')}$ than t^T_i, which is the only vertex of $\{t^T_j\}_{j\in[q]}$ on $Q_{\pi(i')}-\{t\}$. This is a contradiction to the assumption that $\widehat{P}_i$ and $\widehat{Q}_{i'}$ intersect.

Proof of Claim 6.1. Assume the contrary, that $t^T_i=t$ for each $i\in[q]$. Since the paths $\{P_i^{\leftarrow}\}_{i\in[q]}$ are non-empty, $T>1$. Thus, the state S^{T-1} is defined and S^T is obtained by applying a rule to S^{T-1}. First, observe that this rule could not have been **Push**, as it assumes there exist distinct t^{T-1}_i and $t^{T-1}_{i'}$ on $Q_j-\{t\}$ for some $i,i',j\in[q]$, and only moves t^{T-1}_i away while keeping $t^T_{i'}$ on $Q_j-\{t\}$. Therefore, **Clear** has been applied to S^{T-1}, replacing $t^{T-1}_i=v_j$ by $t^T_i=t$, for some $i,j\in[q]$. Now, we claim that there is another $i'\in[q]$, $i'\neq i$ such that $t^{T-1}_{i'}\in V(Q_j)\setminus\{t\}$. Indeed, by the starting assumption of the proof, there exists i' such that $P_{i'}^{\rightarrow}$ and Q_j share an internal vertex u. Since Q_j is active until the last step, and since the application of any rule moves $t^h_{i'}$ to the next vertex of $P_{i'}$ intersecting some active $Q_{j'}$, there exists a step $h\in[T-1]$ where $t^h_{i'}=u$. Before the step $T-1$, only the rule **Push** could have been applied to Q_j, and an application of this rule never makes the intersection $\{t^h_{i''}\}_{i''\in[q]}\cap(Q_j\setminus\{v_j,t\})$ empty. Therefore there exists $i''\in[q]$ such that $t^{T-1}_{i''}\in Q_j\setminus\{v_j,t\}$. This contradicts the assumption that $t^T_{i''}=t$ since the application of **Clear** to Q_j moves $t^{T-1}_{i''}$ to the next vertex on $P_{i''}$ on an active $Q_{j'}$; this will not take $t^T_{i''}$ farther than $v_{i''}$ along $P_{i''}$. □

□

Before we move to the proof of the main theorem, we note that the basic idea of random separation is to exploit random colorings of the vertex set. We, on the other hand, are first and foremost looking for a deterministic algorithm; the standard approach would be to enumerate a sufficiently "expressive" set of colorings, instead of trying a pre-set number of random colorings. Unfortunately, the existing results on derandomization of random separation algorithms cannot be applied directly, as normally random separation is considered for constant number of sets; most often two. Thus in the next lemma we directly construct a suitable family of functions by using the standard tool of perfect hash families, given by the classical result of Naor, Schulman, and Srinivasan [39] (we refer to [14, Chapter 5] for the detailed introduction to the concept). For integers n and k, an (n,k)-*perfect hash family* $\mathcal{F}$ is a family of functions from $[n]$ to $[k]$ such that for each set $S\subset[n]$ of size k there exists $f\in\mathcal{F}$ that acts on S injectively. We are now ready to state our derandomization lemma.

Lemma 16. *For an n-element set U and q integers $k_1,\ldots,k_q$, $\sum_{i=1}^q k_i=\ell$ there exists a family of functions $\mathcal{F}$ of size $q^{\mathcal{O}(\ell)}\log n$ mapping U to $\{1,\ldots,q\}$ with the following property. For any disjoint*

sets $A_1, \ldots, A_q \subset U$ with $|A_i| = k_i$ for $i \in [q]$, there exists a function $f \in \mathcal{F}$ such that $f(x) = i$ if $x \in A_i$. Moreover, $\mathcal{F}$ can be computed in time $q^{\mathcal{O}(\ell)} n \log n$.

Proof. First, construct an (n, ℓ)-perfect hash family $\mathcal{H}$ of size $e^\ell \ell^{\mathcal{O}(\log \ell)} \log n$ in time $e^\ell \ell^{\mathcal{O}(\log \ell)} n \log n$ by the result of Naor, Schulman, and Srinivasan [39]. For every $h \in \mathcal{H}$ and a partition $[t] = I_1 \cup \ldots \cup I_q$ such that $|I_i| = k_i$ for $i \in [q]$, add a function $f^h_{I_1,\ldots,I_q}$ to $\mathcal{F}$. The function acts as follows: for any $x \in U$, $f^h_{I_1,\ldots,I_q}(x) = i$ if $h(x) \in I_i$.

We now show that $\mathcal{F}$ defined above satisfies the conditions of the lemma. Fix the subsets A_1, $\ldots$, A_q of U, denote $A = A_1 \cup \ldots \cup A_q$, $|A| = \ell$. By the definition of an (n, ℓ)-perfect hash family, there exists $h \in \mathcal{H}$ such that the images $h(x)$ are distinct for all $x \in A$. For each $i \in [q]$, define I_i to be the set of indices that h assigns to A_i. By definition, $f^h_{I_1,\ldots,I_q}(x) = i$ if $h(x) \in I_i$, and $h(x) \in I_i$ if and only if $x \in A_i$. It only remains to bound the number of partitions $I_1, \ldots, I_q$.

Claim 6.3. *The number of partitions of $[\ell]$ into disjoint subsets $I_1, \ldots, I_q$ with $|I_i| = k_i$ for $i \in [q]$ is $\mathcal{O}(\ell \cdot q^\ell)$.*

Proof. The number of partitions is equal to the multinomial coefficient

$$\binom{\ell}{k_1, k_2, \ldots, k_q} = \frac{\ell!}{k_1! k_2! \cdots k_q!} \leq e\sqrt{\ell} \frac{\left(\frac{\ell}{e}\right)^t}{\left(\frac{k_1}{e}\right)^{k_1} \cdot \left(\frac{k_2}{e}\right)^{k_2} \cdots \left(\frac{k_q}{e}\right)^{k_q}} = \mathcal{O}\left(\ell \cdot \frac{\ell^\ell}{k_1^{k_1} \cdot k_2^{k_2} \cdots k_q^{k_q}}\right),$$

by Stirling's formula. We now argue that $k_1^{k_1} \cdot k_2^{k_2} \cdots k_q^{k_q} \geq \left(\frac{\ell}{q}\right)^\ell$, which immediately implies that the desired number of partitions is bounded by $\mathcal{O}\left(\ell \cdot \frac{\ell^\ell}{\left(\frac{\ell}{q}\right)^\ell}\right) = \mathcal{O}(\ell \cdot q^\ell)$. For that, we observe that the function $f(x) = x \log x$ is strictly convex on $x \geq 1$, since $(x \log x)' = (1 + \log x)$, thus $x \log x + y \log y \geq 2 \cdot \left(\frac{x+y}{2} \cdot \log \frac{x+y}{2}\right)$, and $x^x \cdot y^y \geq \left(\frac{x+y}{2}\right)^{2 \cdot \frac{x+y}{2}}$ for any $x, y \geq 1$, where the equality only holds if $x = y$. Now, consider the function $h(x_1, x_2, \ldots, x_q) = x_1^{x_1} \cdot x_2^{x_2} \cdots x_q^{x_q}$ defined on the polytope $K \subset \mathbb{R}^q$ bounded by $x_1 \geq 1$, $\ldots$, $x_q \geq 1$, $\sum_{i=1}^q x_i = \ell$. Since h is continuous on K and K is compact, h attains its minimum in K. Assume that h achieves its minimum on x_1, $\ldots$, $x_q \in K$ with $x_i \neq x_j$ for some $i, j \in [q]$. Then by the above, $h(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_q) > h(x_1, \ldots, \frac{x_i+x_j}{2}, \ldots, \frac{x_i+x_j}{2}, \ldots, x_q)$; the tuple on the right-hand side still belongs to K. Thus, such a x_1, $\ldots$, x_q cannot achieve the minimum, and h is minimized at the only point with equal coordinates, $(\ell/q, \ldots, \ell/q) \in K$. Since $(k_1, \ldots, k_q) \in K$, the claim is done. □

The bound on the size of $\mathcal{F}$ now follows directly from Claim 6.3: $|\mathcal{F}| = |\mathcal{H}| \cdot \mathcal{O}(\ell \cdot q^\ell) = q^{\mathcal{O}(\ell)} \cdot n^{\mathcal{O}(1)}$. □

Finally, with Lemma 15 and Lemma 16 at hand, we move to the proof of Theorem 2 itself.

Proof of Theorem 2. First, we observe that finding an (S, T)-linkage of order p and total length at least k is equivalent to finding an (s, t)-linkage of order p and total length at least $(k + 2)$, where moreover $s \neq t$ and s is not adjacent to t. Indeed, consider the digraph G' that is a copy of G with two new vertices s and t, where $N^+_{G'}(s) = S$ and $N^-_{G'}(t) = T$. Then, any directed (s, t)-linkage of order p and length $k + 2$ in G' induces a directed (S, T)-linkage of order p and length k in G by removing s and t, and vice versa. Thus for the rest of the proof we assume that the task is to find an (s, t)-linkage of order p and total size at least k, $s \neq t$, and s is not adjacent to t. We now describe two separate subroutines of our algorithm, tailored for different cases of the maximum length of the path in the target (s, t)-linkage. The **short case** succeeds if there is an (s, t)-linkage

where all paths have less than $2k$ internal vertices, and the **main case** succeeds otherwise (the proof of correctness follows after the description of the algorithm).

Short case. For each $i \in [p]$, we branch over the number of internal vertices k_i of the i-th path in the target linkage, $1 \leq k_i < 2k$. If $\sum_{i=1}^{p} k_i < k-2$, we disregard the choice of $\{k_i\}_{i=1}^{p}$ and proceed to the next branch. Otherwise, consider a function family $\mathcal{F}$ given by an invocation of Lemma 16 with $q = p$ and the current values of $k_1, \ldots, k_p$. Branch over the choice of $f \in \mathcal{F}$ and denote by $C_1, C_2, \ldots, C_p$ the vertices colored by the respective colors via f. For each $i \in [p]$, we use a deterministic algorithm for finding a directed (s,t)-path with exactly k_i internal vertices in the graph $G[C_i \cup \{s,t\}]$ in time $2^{Oh(k_i)} \cdot n^{\mathcal{O}(1)}$. The fastest-known such algorithm is the algorithm of Zehavi [48] running in time $\mathcal{O}(2.597^{k_i}) \cdot n^{\mathcal{O}(1)}$[3] If for some choice of $\{k_i\}_{i=1}^{p}$ and f the desired collection of paths is found, the algorithm returns it. If no branch succeeds, the algorithm reports a no-instance.

We now argue for correctness of the algorithm above. Since the paths are internally-disjoint by construction and $\sum_{i=1}^{p} k_i \geq k-2$, if the algorithm returns a collection of paths, they clearly form a solution. In the other direction, fix a solution induced by directed (s,t)-paths $P_1^*, \ldots, P_p^*$, where for each i the i-th path contains exactly $k_i < 2k$ internal vertices, and consider the respective branch of the algorithm above. If for each $i \in [p]$ the set C_i contains the internal vertices of the i-th path, the algorithm succeeds, as $G[C_i \cup \{s,t\}]$ contains an (s,t)-path with exactly k_i internal vertices. Denote by A_i the set of internal vertices of P_i^*, for each $i \in [p]$, Lemma 16 guarantees that there exists $f \in \mathcal{F}$ that colors each A_i in color i, concluding the proof of correctness in this case.

Algorithm 1: Main case of the algorithm in Theorem 2.

```
1  for q' = 1, ..., p do
2      Invoke Lemma 16 with q = p + 1, k_1 = ... = k_{q'} = k, k_{q'+1} = ... = k_p = 2k, and
         k_{p+1} = k, to obtain the function family F;
3      foreach f ∈ F do
4          Denote by C_0, C_1, ..., C_p the vertices colored by the respective colors via f;
5          for i = 1, ..., p do
6              foreach v_i ∈ C_i at distance k from t in G_i = G[C_i ∪ {t}] do
7                  Find a shortest (v_i, t)-path Q_i in G_i;
8                  Find paths P_1, P_2, ..., P_p in G − (V(Q_i) \ {v_i, t}), where P_i is an (s, v_i)-path
                     and for each j ≠ i, P_j is an (s, t)-path, such that no two paths share a
                     vertex except for s and t, and P_i does not contain t;
9                  if such paths P_1, P_2, ..., P_p exist then
10                     Set P_i = P_i ∘ Q_i;
11                     return the paths P_1, P_2, ..., P_p;
12                 end
13             end
14         end
15     end
16 end
```

Main case. The basic procedure is given in Algorithm 1. Observe that the task in Line 8 can be easily reduced to an instance of network flow, thus the whole procedure given in Lines 4

[3]While the result in [48] is stated for finding an arbitrary path of certain length, it could be easily adjusted to finding an (s,t)-path.

to 14 runs in polynomial time. If no iteration returns a collection of paths, the algorithm reports a no-instance.

It is easy to observe that if Algorithm 1 returns a collection of paths, then these paths constitute a solution to the given instance. Indeed, by construction, $P_1, \ldots, P_p$ are internally-disjoint (s,t)-paths. The length of Q_i is exactly k, since the path P_i contains Q_i as a subpath, the (s,t)-linkage given by the paths $P_1, \ldots, P_p$ is of order p and length at least k. It remains to verify that if there exists an (s,t)-linkage of order p where there is a path with at least $2k$ internal vertices, then our algorithm successfully returns a collection of paths for some choice of f. Denote the p-many (s,t)-paths that form a solution of minimum total length by $P_1^*, \ldots, P_p^*$, assuming that the paths are ordered from longest to shortest. In particular, $|V(P_1^*) \setminus \{s,t\}| \geq 2k$. Denote by q' the maximum index such that $|V(P_{q'}^*) \setminus \{s,t\}| \geq 2k$, $1 \leq q' \leq p$.

We say that the coloring $C_0, C_1, \ldots, C_p$ *agrees* with the paths $P_1^*, \ldots, P_p^*$ if the following holds:

(i) for each $i \in [q']$, the last k internal vertices of P_i^* belong to C_i,

(ii) the first k vertices of P_1^* belong to C_0,

(iii) for each $j \in [p] \setminus [q']$, all internal vertices of P_j^* belong to C_j.

For $i \in [q']$, denote by A_i the last k internal vertices of P_i^*; by A_{p+1} the first k vertices of P_1^*; for $j \in [p] \setminus [q']$, denote by A_j all internal vertices of P_j^* plus arbitrary vertices not yet part of any A_i so that $|A_j| = 2k$. By Lemma 16, there exists $f \in \mathcal{F}$ that induces a coloring $C_1, \ldots, C_{p+1}$ with $A_i \subset C_i$ for each $i \in [p+1]$, meaning that this coloring agrees with the solution $P_1^*, \ldots, P_p^*$. In the remainder of the proof, we argue that for this choice of f Algorithm 1 outputs a solution.

For $i \in [p]$, let v_i be the vertex of P_i^* at distance exactly k from t along the path. Since the coloring agrees with the solution, the last k internal vertices of P_i^* belong to C_i, including v_i. If v_i is at distance k from t in G_i, denote by Q_i the shortest (v_i,t)-path in G_i that the algorithm finds on Line 7; otherwise denote by Q_i an arbitrary shortest (v_i,t)-path in G_i. We now apply Lemma 15 to the paths $P_1^*, \ldots, P_q^*$ in the graph G with selected disjoint vertex subsets $C_1, \ldots, C_q$, with selected vertices $v_1, \ldots, v_q$and paths $Q_1, \ldots, Q_q$. By the lemma, there exist internally-disjoint (s,t)-paths $P_1', \ldots, P_q'$, such that for each $i \in [q]$, P_i' is either P_i^* or a concatenation of a prefix $\widehat{P}_i$ of P_1^* not extending beyond v_i, and a suffix $\widehat{Q}_i$ of Q_j for some $j \in [q]$; moreover, at least one of the paths is of the second type.

First, we claim that the paths $P_1', \ldots, P_q', P_{q+1}^*, \ldots, P_p^*$ together form an $(\{s\},\{t\})$-linkage of length at least k. Indeed, the paths $P_1^*, \ldots, P_p^*$ are internally-disjoint from the beginning; thus for each $q < i < j \leq p$, P_i^* and P_j^* do not share common internal vertices. Moreover, for each $i \in [q]$, $j \in [p] \setminus [q]$, P_i' and P_j^* are immediately internally-disjoint in case $P_i' = P_i^*$. In case P_i' is a concatenation of $\widehat{P}_i$ and $\widehat{Q}_i$, the path $\widehat{P}_i$ again does not share a vertex with P_j^* as a prefix of P_i^*, except for s. The suffix $\widehat{Q}_i - \{t\}$ on the other hand is fully contained in some $C_{i'}$, $i' \in [q]$, while $P_j^* - \{s,t\}$ is contained in C_j by the property (iii) of the coloring. Finally, for the length observe that the path P_1' contains the first k internal vertices of the path P_1^* since they belong to the set C_0 disjoint from C_i for any $i > 0$.

Consider now a path P_i' that is not P_i^* but a concatenation, $\widehat{P}_i \circ \widehat{Q}_i$. Since the paths $P_1^*, \ldots, P_p^*$ come from a solution of smallest total length, it cannot be that P_i' is shorter than P_i^*. On the other hand, the length of $\widehat{Q}_i$ is at most k, which is the length of the suffix of P_i^* from v_i to t, and $\widehat{P}_i$ is at most as long as the prefix of P_i^* from s to v_i. Thus it has to be that $\widehat{P}_i$ is exactly the prefix of P_i^* from s to v_i, and $\widehat{Q}_i$ is Q_i; additionally, Q_i is then of length exactly k, thus v_i is at distance k from t in G_i, and Q_i is the shortest path chosen by the algorithm in Line 7.

We now claim that the existence of (s,t)-linkage $P'_1, \ldots, P'_q, P_{q+1}, \ldots, P_p$ implies that Line 8 is executed successfully for the respective choice of v_i. Indeed, denote $P_i = \widehat{P}_i$, for each $i' \in [q]$, $i' \neq i$, denote $P_{i'} = P'_i$, and for each $j \in [p] \setminus [q]$, denote $P_j = P^*_j$. This collection of paths is of form required by the conditions in Line 8, and thus the algorithm of Line 8 returns a suitable collection of paths as well (not necessarily the same). This concludes the proof of correctness for **Main case**.

Finally, observe that the running time of both cases is dominated by the invocation of Lemma 16 with $q = p + 1$ and $\ell \leq 2k(p+1)$, resulting in the total running time of $p^{\mathcal{O}(kp)} \cdot n^{\mathcal{O}(1)}$. □

7 Conclusion

We conclude with several concrete open questions. The first question is about derandomizing Theorem 1, even for the case when $p = 1$. The algorithm in Theorem 1 is based on DeMillo-Lipton-Schwartz-Zippel lemma for polynomial identity testing, and therefore we do not expect to derandomize it using similar techniques [27, 38]. However, similarly to Theorem 2, we do not exclude that other methods could result in (maybe slower) deterministic algorithms. We are not aware of any *deterministic* and FPT in k algorithm for Maximum Colored Path.

The second question is about the Disjoint Paths problem. Here for a given set of pairs of terminal vertices $(s_1,t_1),\ldots,(s_r,t_r)$, the problem is to decide whether there are vertex-disjoint (s_i,t_i)-paths, $i \in [r]$. The problem is FPT parameterized by r by the seminal algorithm of Robertson and Seymour [43]. A natural extension of this problem would be on colorful graphs, where we want the disjoint paths to collect at least k colors. We do not know whether the colored variant of the problem is FPT parameterized by k even for 2-Disjoint Paths, that is for $r = 2$.

The third question concerns extending Theorem 3, where we demand matroid M to be represented as a matrix over a finite field of order q. The natural question here is whether there is an FPT algorithm for k-ranked (s,t)-path (and more generally, for k-ranked (S,T)-linkage of order p) in frameworks (G,M), where M is a linear matroid represented as a matrix over rationals. We also ask what is the complexity of this problem when M is given by an independence oracle. As was shown by Jensen and Korte [26], various matroid problems have unconditional complexity lower bounds asserting that they admit no algorithms where the number of oracle calls is bounded by a polynomial on the size of the matroid ground set. For example, this concerns the classical Matroid Parity problem that can be solved in polynomial time on linear matroids as it was shown by Lovász (see, e.g., [36]). It is natural to ask whether such a lower bound can be shown for k-ranked (s,t)-path.

The last concrete question is about Longest (s,t)-Path and Longest Cycle. Our algorithm implies the first $2^k n^{\mathcal{O}(1)}$ time algorithms for these problems, and the dependency on k in the time complexity of our algorithm is unlikely to be improved in the general colored case. However, it remains an interesting open problem whether Longest (s,t)-Path or Longest Cycle could be solved in time $(2-\varepsilon)^k n^{\mathcal{O}(1)}$ for some $\varepsilon > 0$, especially keeping in mind that k-Path admits an $1.66^k n^{\mathcal{O}(1)}$ time algorithm [3].

References

[1] N. Alon, R. Yuster, and U. Zwick, *Color-coding*, J. ACM, 42 (1995), pp. 844–856. 1, 3, 6, 11

[2] A. Björklund, *Determinant sums for undirected hamiltonicity*, SIAM J. Comput., 43 (2014), pp. 280–299. 1, 6, 7, 20

[3] A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto, *Narrow sieves for parameterized paths and packings*, J. Comput. Syst. Sci., 87 (2017), pp. 119–139. 1, 2, 6, 46

[4] A. Björklund, T. Husfeldt, and N. Taslaman, *Shortest cycle through specified elements*, in Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012, Y. Rabani, ed., SIAM, 2012, pp. 1747–1753. 3, 5, 6, 7, 8, 10

[5] A. Björklund, P. Kaski, and L. Kowalik, *Constrained multilinear detection and generalized graph motifs*, Algorithmica, 74 (2016), pp. 947–967. 6

[6] H. Broersma, X. Li, G. J. Woeginger, and S. Zhang, *Paths and cycles in colored graphs*, Australas. J Comb., 31 (2005), pp. 299–312. 3

[7] G. Călinescu, C. Chekuri, M. Pál, and J. Vondrák, *Maximizing a monotone submodular function subject to a matroid constraint*, SIAM J. Comput., 40 (2011), pp. 1740–1766. 4

[8] R. Cerulli, P. Dell'Olmo, M. Gentili, and A. Raiconi, *Heuristic approaches for the minimum labelling hamiltonian cycle problem*, Electron. Notes Discret. Math., 25 (2006), pp. 131–138. 3

[9] C. Chekuri and M. Pál, *A recursive greedy algorithm for walks in directed graphs*, in 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, 2005, pp. 245–253. 4

[10] N. Christofides, *Vehicle routing*, in The Traveling Salesman Problem, E. Lawler, J. Lenstra, A. Kan, and D. S. (Eds.), eds., Wiley, New York, 1985, pp. 431–448. 3

[11] J. Cohen, G. F. Italiano, Y. Manoussakis, N. K. Thang, and H. P. Pham, *Tropical paths in vertex-colored graphs*, J. Comb. Optim., 42 (2021), pp. 476–498. 3

[12] B. Couëtoux, E. Nakache, and Y. Vaxès, *The maximum labeled path problem*, Algorithmica, 78 (2017), pp. 298–318. 3

[13] M. Cygan, H. Dell, D. Lokshtanov, D. Marx, J. Nederlof, Y. Okamoto, R. Paturi, S. Saurabh, and M. Wahlström, *On problems as hard as CNF-SAT*, ACM Trans. Algorithms, 12 (2016), pp. 41:1–41:24. 2, 14

[14] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh, *Parameterized Algorithms*, Springer, 2015. 1, 13, 42

[15] M. Cygan, S. Kratsch, and J. Nederlof, *Fast hamiltonicity checking via bases of perfect matchings*, in Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC), ACM, 2013, pp. 301–310. 1

[16] M. Cygan, J. Nederlof, M. Pilipczuk, M. Pilipczuk, J. M. M. van Rooij, and J. O. Wojtaszczyk, *Solving connectivity problems parameterized by treewidth in single exponential time*, in Proceedings of the 52nd Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2011, pp. 150–159. 1

[17] R. Diestel, *Graph Theory, 4th Edition*, vol. 173 of Graduate texts in mathematics, Springer, 2012. 14

[18] E. Eiben and I. Kanj, *A colored path problem and its applications*, ACM Trans. Algorithms, 16 (2020), pp. 47:1–47:48. 4

[19] F. V. Fomin, P. A. Golovach, D. Sagunov, and K. Simonov, *Algorithmic extensions of Dirac's theorem*, CoRR, abs/2011.03619 (2020). 2, 6, 11, 13

[20] ——, *Algorithmic extensions of Dirac's theorem*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA22), SIAM, 2022, pp. 406–416. 2

[21] F. V. Fomin, D. Lokshtanov, F. Panolan, and S. Saurabh, *Efficient computation of representative families with applications in parameterized and exact algorithms*, J. ACM, 63 (2016), pp. 29:1–29:60. 1

[22] F. V. Fomin, D. Lokshtanov, F. Panolan, S. Saurabh, and M. Zehavi, *Long directed (*s, t*)-path: FPT algorithm*, Inf. Process. Lett., 140 (2018), pp. 8–12. 2

[23] S. Fortune, J. E. Hopcroft, and J. Wyllie, *The directed subgraph homeomorphism problem*, Theor. Comput. Sci., 10 (1980), pp. 111–121. 14

[24] M. Grohe, K. Kawarabayashi, D. Marx, and P. Wollan, *Finding topological subgraphs is fixed-parameter tractable*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC), ACM, 2011, pp. 479–488. 2

[25] R. Hassin, J. Monnot, and D. Segev, *Approximation algorithms and hardness results for labeled connectivity problems*, J. Comb. Optim., 14 (2007), pp. 437–453. 3

[26] P. M. Jensen and B. Korte, *Complexity of matroid property algorithms*, SIAM J. Comput., 11 (1982), pp. 184–190. 46

[27] V. Kabanets and R. Impagliazzo, *Derandomizing polynomial identity tests means proving circuit lower bounds*, computational complexity, 13 (2004), pp. 1–46. 46

[28] K. Kawarabayashi, *An improved algorithm for finding cycles through elements*, in 13th International Conference on Integer Programming and Combinatorial Optimization (IPCO), vol. 5035 of Lecture Notes in Computer Science, Springer, 2008, pp. 374–384. 3

[29] I. Koutis, *Faster algebraic algorithms for path and packing problems*, in Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP), vol. 5125, Springer, 2008, pp. 575–586. 1, 6

[30] I. Koutis and R. Williams, *Algebraic fingerprints for faster algorithms*, Commun. ACM, 59 (2016), pp. 98–105. 6

[31] L. Kowalik and J. Lauri, *On finding rainbow and colorful paths*, Theor. Comput. Sci., 628 (2016), pp. 110–114. 3

[32] N. Kumar, D. Lokshtanov, S. Saurabh, and S. Suri, *A constant factor approximation for navigating through connected obstacles in the plane*, in Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, (SODA), SIAM, 2021, pp. 822–839. 3

[33] A. Lingas and M. Persson, *A fast parallel algorithm for minimum-cost small integral flows*, Algorithmica, 72 (2015), pp. 607–619. 6, 7

[34] L. Lovász, *Flats in matroids and geometric graphs*, in Combinatorial surveys (Proc. Sixth British Combinatorial Conf., Royal Holloway Coll., Egham, 1977), 1977, pp. 45–86. 4

[35] L. Lovász, *Graphs and geometry*, vol. 65 of American Mathematical Society Colloquium Publications, American Mathematical Society, Providence, RI, 2019. 4

[36] L. Lovász and M. D. Plummer, *Matching theory*, vol. 121 of North-Holland Mathematics Studies, North-Holland Publishing Co., Amsterdam; North-Holland Publishing Co., Amsterdam, 1986. Annals of Discrete Mathematics, 29. 4, 46

[37] D. Marx, *A parameterized view on matroid optimization problems*, Theor. Comput. Sci., 410 (2009), pp. 4471–4479. 37, 38, 39

[38] J. A. Montoya and M. Müller, *Parameterized random complexity*, Theory Comput. Syst., 52 (2013), pp. 221–270. 46

[39] M. Naor, L. J. Schulman, and A. Srinivasan, *Splitters and near-optimal derandomization*, in Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS 1995), IEEE, 1995, pp. 182–191. 13, 42, 43

[40] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, *An analysis of approximations for maximizing submodular set functions - I*, Math. Program., 14 (1978), pp. 265–294. 4

[41] J. Oxley, *Matroid theory*, vol. 21 of Oxford Graduate Texts in Mathematics, Oxford University Press, Oxford, second ed., 2011. 37

[42] F. Panolan, S. Saurabh, and M. Zehavi, *Parameterized algorithms for list k-cycle*, Algorithmica, 81 (2019), pp. 1267–1287. 3, 6

[43] N. Robertson and P. D. Seymour, *Graph minors XIII. The disjoint paths problem*, J. Comb. Theory, Ser. B, 63 (1995), pp. 65–110. 46

[44] J. T. Schwartz, *Fast probabilistic algorithms for verification of polynomial identities*, J. ACM, 27 (1980), pp. 701–717. 8, 14

[45] M. Wahlström, *Abusing the Tutte matrix: An algebraic instance compression for the K-set-cycle problem*, in 30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany, N. Portier and T. Wilke, eds., vol. 20 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013, pp. 341–352. 3

[46] R. Williams, *Finding paths of length k in $O^*(2^k)$ time*, Inf. Process. Lett., 109 (2009), pp. 315–318. 1, 6

[47] H.-C. Wirth, *Multicriteria approximation of network design and network upgrade problems*, PhD thesis, Universität Würzburg, 2001. 3

[48] M. Zehavi, *Mixing color coding-related techniques*, in Algorithms - ESA 2015, N. Bansal and I. Finocchi, eds., Berlin, Heidelberg, 2015, Springer Berlin Heidelberg, pp. 1037–1049. 44

[49] M. Zehavi, *A randomized algorithm for long directed cycle*, Inf. Process. Lett., 116 (2016), pp. 419–422. 1, 2

[50] R. Zippel, *Probabilistic algorithms for sparse polynomials*, in Symbolic and Algebraic Computation, EUROSAM '79, An International Symposiumon Symbolic and Algebraic Computation, Marseille, France, June 1979, Proceedings, E. W. Ng, ed., vol. 72 of Lecture Notes in Computer Science, Springer, 1979, pp. 216–226. 8, 14