

Interference-free Walks in Time: Temporally Disjoint Paths ^{*}

Nina Klobas^{1†}, George B. Mertzios^{1†}, Hendrik Molter^{2†}, Rolf Niedermeier^{3†} and Philipp Zschoche^{3†}

¹ Department of Computer Science, Durham University, Upper Mountjoy Campus, Stockton Rd, Durham, DH1 3LE, England, United Kingdom.

² Department of Industrial Engineering and Management, Ben-Gurion University of the Negev, David Ben Gurion Blvd, Beer Sheva, 84105, Israel.

³ Algorithmics and Computational Complexity, Faculty IV, TU Berlin, Ernst-Reuter-Platz 7, Berlin, 10587, Germany.

Contributing authors: nina.klobas@durham.ac.uk;
george.mertzios@durham.ac.uk; molterh@post.bgu.ac.il;
rolf.niedermeier@tu-berlin.de; zschoche@tu-berlin.de;

[†]These authors contributed equally to this work.

Abstract

We investigate the computational complexity of finding *temporally disjoint* paths or walks in temporal graphs. There, the edge set changes over discrete time steps and a temporal path (resp. walk) uses edges that appear at monotonically increasing time steps. Two paths (or walks) are temporally disjoint if they never visit the same vertex at the same time; otherwise, they interfere. This reflects applications in robotics, traffic routing, or finding safe pathways in dynamically changing networks. On the one extreme, we show that on general graphs the problem is computationally hard. The “walk version” is W[1]-hard when parameterized by the number of walks. However, it is polynomial-time solvable for any constant number of walks. The “path version”

^{*}A preliminary version of this paper appeared at the Proceedings of IJCAI 2021 [1]. This revised version now contains all proof details.




Fig. 1: A temporal graph where a label of an edge reflects at which time it is available. There are two temporally disjoint (s, z) -paths P_1 and P_2 , where P_1 uses the solid (orange) edges and P_2 the dashed (blue) edges. Here, P_1 visits v before P_2 .

remains NP-hard even if we want to find only two temporally disjoint paths. On the other extreme, restricting the input temporal graph to have a path as underlying graph, quite counter-intuitively, we find NP-hardness in general but also identify natural tractable cases.

Keywords: Temporal graph, temporal path, NP-complete problem, parameterized complexity, polynomial-time algorithm.

1 Introduction

Computing (vertex-)disjoint paths in a graph is a cornerstone problem of algorithmic graph theory and many applied network problems. It was among the early problems that were shown to be NP-complete [2]. One of the deepest achievements in discrete mathematics, graph minor theory [3, 4], as well as the development of the theory of parameterized complexity analysis [5] are tightly connected to it. The problem is known to be solvable in quadratic time if the number of paths is constant, that is, it is fixed-parameter tractable when parameterized by the number of paths [6]. Besides being of fundamental interest in (algorithmic) graph theory, finding disjoint paths has many applications and there exist numerous variations of the problem. In AI and robotics scenarios, for instance, multi-agent path finding is an intensively studied, closely related problem [7, 8].

Coming from the graph-algorithmic side, we propose a new view on finding disjoint paths (and walks), that is, we place the problem into the world of temporal graphs. We add a “new dimension” to the classic, static graph scenario by generalizing to a setting where the edges of a graph may appear and disappear over (discrete) time. In our model, we consider two paths (or walks) to be disjoint if they do not use the same vertex at the same point of time. Consider Fig. 1 for an example. Moreover, the path finding also has to take into account that edges are not permanently available, reflecting dynamic aspects of many real-world scenarios such as routing in traffic or communication networks, or the very dynamic nature of social networks. We intend to initiate studies on this natural scenario. Doing so, we focus on two extreme cases for the underlying graphs, namely the (underlying) graph structure being completely unrestricted or it being restricted to just a path graph. For these

opposite poles, performing (parameterized) computational complexity studies, we present surprising discoveries. Before coming to these, we discuss (excerpts of) the large body of related work.

Related work. As mentioned above, the literature on (static) disjoint paths and on multi-agent path finding is very rich. Hence, we only list a small fraction of the relevant related work. In context of graph-algorithmic work, the polynomial-time (in-)approximability of the NP-hard maximization version has been studied [9]. Variants of the basic problem studied include bounds on the path length [10] or relaxing on the disjointness of paths [11–14].

In directed graphs, finding two disjoint paths is already NP-hard [15], whereas in directed acyclic graphs the problem is solvable in polynomial time for every fixed number of paths [16].

As to multi-agent path finding, we remark that it has been intensively researched (with several possible definitions) in the last decade in the AI and robotics communities [7, 8, 17–19]. Timing issues (concurrency of moving agents) and the various objective functions of the agents play a fundamental role here; also a high variety of conflict scenarios is studied.¹ The scenario we study in this work can be interpreted as a basic variant of multi-agent path planning, now translated into the world of temporal graphs.

In algorithmic graph theory, edge-colored graphs have also been studied. Edge-colored graphs are essentially multilayer (or multiplex) graphs where the fundamental difference to temporal graphs is that there is no order on the graph snapshots (also referred to as layers). Here, path-finding scenarios are motivated, for example, by applications in social and optical (routing) networks [20–22].

Finally, as to temporal graphs, note that several prominent graph problems have been studied in this fairly new framework. This includes path problems [23–30] and in particular another model of vertex-disjoint temporal paths [31], where two temporal paths are considered vertex-disjoint if they do not visit the same vertex. The problem of finding two such paths is NP-hard [31]. Note that the major difference to our model is that we allow two *temporally* disjoint paths to visit the same vertex as long as they do not both visit that vertex at the same *time*. Apart from path-related problems, in the previous years also various non-path temporal problems have been introduced and studied, such as temporal separation problems [32, 33], vertex coloring [34], matching [35], vertex cover [36, 37], transitive orientation [38], and betweenness [39].

Our contributions. Our results can be grouped into two parts. First, studying temporal graphs where the underlying graph (which is obtained by making all temporal edges permanent) is unrestricted, we show in Section 3 that finding walks instead of paths turns out to be computationally easier. More specifically, finding temporally disjoint walks is W[1]-hard with respect to the number of walks but can be solved in polynomial time if this number is constant (that is,

¹Also see the multi-agent path planning webpage: <http://mapf.info/>

TEMP. DISJOINT	PATHS	WALKS
unrestricted underlying graph	NP-hard for $ S = 2$	W[1]-hard wrt. $ S $ XP wrt. $ S $
temporal line or tree	NP-hard FPT wrt. $ S $	<i>open</i>
temporal line & S contains only pairs of extremal points	poly-time	

Table 1: Overview computational complexity of TEMPORALLY DISJOINT (PATHS/WALKS). Here, S is the multiset of source-sink pairs. Temporal line means that the underlying graph is a path.

in the language of parameterized algorithmics, we develop an XP algorithm), whereas finding temporally disjoint paths already turns out to be NP-hard for two paths. Second, restricting the input to be a temporal line or a temporal tree (i.e., the underlying graph to be a path or a tree, respectively), we prove in Section 4 that the problem is NP-hard (for both paths and walks). However, we also provide a fixed-parameter tractability result with respect to the number of paths. For the special case where, in an input temporal line, the given multiset of source-sink pairs only contains pairs of the extremal points of the temporal line, we provide a polynomial-time algorithm. We survey our results in Table 1.

2 Preliminaries and problem definition

We denote by \mathbb{N} and \mathbb{N}_0 the natural numbers excluding and including 0, respectively. An interval on \mathbb{N}_0 from a to b is denoted by $[a, b] := \{i \in \mathbb{N}_0 \mid a \leq i \leq b\}$ and $[a] := [1, a]$.

Static graphs. An undirected graph $G = (V, E)$ consists of a set V of vertices and a set $E \subseteq \{\{v, w\} \mid v, w \in V, v \neq w\}$ of edges. For a graph G , we also denote by $V(G)$ and $E(G)$ the vertex and edge set of G , respectively. For a vertex set $W \subseteq V$, the induced subgraph $G[W]$ is defined as the graph $(W, \{\{v, w\} \in E \mid v, w \in W\})$. A path $P = (V, E)$ is a graph with a set $V(P) = \{v_1, \dots, v_k\}$ of distinct vertices and edge set $E(P) = \{\{v_i, v_{i+1}\} \mid 1 \leq i < k\}$ (we often represent path P by the tuple (v_1, v_2, \dots, v_k)). We say that P is a (v_1, v_k) -path and that P visits all vertices in $V(P)$.

Temporal graphs and temporally disjoint paths. A temporal graph $\mathcal{G} = (V, (E_i)_{i \in [T]})$ consists of a set V of vertices and lifetime T many edge sets E_1, E_2, \dots, E_T over V . The pair (e, i) is a *time edge* of \mathcal{G} if $e \in E_i$. The graph (V, E_i) is called the *i -th layer* of \mathcal{G} . The *underlying graph* of \mathcal{G} is the static graph $(V, \bigcup_{i=1}^T E_i)$. A *temporal (s, z) -walk* (or *temporal walk* from s to z) of length k from vertex $s = v_0$ to vertex $z = v_k$ in \mathcal{G} is a

sequence $P = ((v_{i-1}, v_i, t_i))_{i=1}^k$ of triples such that for all $i \in [k]$ we have that $\{v_{i-1}, v_i\} \in E_{t_i}$ and for all $i \in [k-1]$ we have that $t_i \leq t_{i+1}$. The *arrival time* of P is t_k . We say that P *visits* the vertices $V(P) := \{v_i \mid i \in [0, k]\}$. In particular, P visits vertex v_i during the time interval $[t_i, t_{i+1}]$, for all $i \in [k-1]$. Furthermore, we say that P visits v_0 during time interval $[t_1, t_1]$ and P visits v_k during time interval $[t_k, t_k]$. A temporal (s, z) -walk $P = ((v_{i-1}, v_i, t_i))_{i=1}^k$ of length k is called a *temporal (s, z) -path* (or *temporal path* from s to z) if $v_i \neq v_j$ whenever $i \neq j$. Given two temporal walks P_1, P_2 we say that P_1 and P_2 *temporally intersect* if there exists a vertex v and two time intervals $[a_1, b_1], [a_2, b_2]$, where $[a_1, b_1] \cap [a_2, b_2] \neq \emptyset$, such that v is visited by P_1 during $[a_1, b_1]$ and by P_2 during $[a_2, b_2]$. Now, we can formally define our problem.

TEMPORALLY DISJOINT PATHS

Input: A temporal graph $\mathcal{G} = (V, (E_i)_{i \in [T]})$ and a multiset S of source-sink pairs containing elements from $V \times V$.

Question: Are there pairwise temporally non-intersecting temporal (s_i, z_i) -paths for all $(s_i, z_i) \in S$?

Analogously, TEMPORALLY DISJOINT WALKS gets the same input but asks whether there are pairwise temporally non-intersecting temporal (s_i, z_i) -walks for all $(s_i, z_i) \in S$. From the NP-hardness of DISJOINT PATHS [2], we immediately get the following.

Observation 1 TEMPORALLY DISJOINT (PATHS/WALKS) is NP-hard even if $T = 1$.

For an instance $(\mathcal{G} = (V, (E_i)_{i \in [T]}), S)$ of TEMPORALLY DISJOINT (PATHS/WALKS) we assume throughout this paper that in the input, \mathcal{G} is given by the vertex set V followed by the ordered (by time label) subsequence of E_1, \dots, E_T , only containing the non-empty edge sets. To make the presentation simpler, we apply a linear-time preprocessing step to the input by renumbering these non-empty edge sets E_i (still keeping their relative order the same) such that all non-empty sets are consecutive. Note that this preprocessing step creates an instance that is equivalent to the original input instance. Therefore, without loss of generality we assume in the remainder of our work that all edge sets $E_i, i = 1, 2, \dots, T$, are non-empty (where the new lifetime T is now the number of non-empty edge sets in the original input). Hence, the *size* of \mathcal{G} is $|\mathcal{G}| := |V| + \sum_{t=1}^T |E_t|$.

Parameterized complexity. Let Σ denote a finite alphabet. A parameterized problem $L \subseteq \{(x, k) \in \Sigma^* \times \mathbb{N}_0\}$ is a subset of all instances (x, k) from $\Sigma^* \times \mathbb{N}_0$, where k denotes the *parameter*. A parameterized problem L is (i) FPT (*fixed-parameter tractable*) if there is an algorithm that decides every instance (x, k) for L in $f(k) \cdot |x|^{O(1)}$ time, and (ii) contained in the class XP if there is an algorithm that decides every instance (x, k) for L in $|x|^{f(k)}$ time, where f is any computable function only depending on the parameter. If a parameterized

problem L is W[1]-hard, then it is presumably not fixed-parameter tractable. We refer to Downey and Fellows [5] for more details.

3 The case of few source-sink pairs

In this section, we study the computational complexity of TEMPORALLY DISJOINT (PATHS/WALKS) for the case that the size of the multiset S of source-sink pairs is small. We start by showing that TEMPORALLY DISJOINT PATHS is NP-hard even for two sink-source pairs. This is a similar situation as for finding vertex-disjoint paths in directed static graphs, which is also NP-hard for two paths [15]. However, in the temporal setting there is a surprising difference between finding walks and paths that does not have an analogue in the static setting. We will show that TEMPORALLY DISJOINT WALKS is W[1]-hard for the number $|S|$ of source-sink pairs and is contained in XP for the same parameter. First, however, we consider the computationally seemingly even more intractable path variant.

Theorem 2 TEMPORALLY DISJOINT PATHS is NP-hard even if $|S| = 2$ and $T = 3$.

Proof We show that TEMPORALLY DISJOINT PATHS is NP-hard even if $|S| = 2$ and $T = 3$ by a polynomial-time reduction from the NP-complete EXACT (3,4)-SAT problem [40]. EXACT (3,4)-SAT asks whether a Boolean formula ϕ is satisfiable, assuming that it is in conjunctive normal form, each clause has exactly three literals, and each variable appears in exactly four clauses.

Construction. Let ϕ be an instance of EXACT (3,4)-SAT with n variables x_1, x_2, \dots, x_n and m clauses. We construct an instance $I = (\mathcal{G} = (V, (E_1, E_2, E_3)), S = \{(s_1, z_1), (s_2, z_2)\})$ in the following way. Intuitively, the first two layers contain the assignment gadget for the variables. The idea is that in time step one the temporal (s_1, z_1) -path P_1 goes from s_1 to s' and hereby sets all variables to true or false depending on which route was taken. At time step two, the unique temporal (s_2, z_2) -path departs and arrives. Moreover, this unique temporal (s_2, z_2) -path visits all vertices of the assignment gadget except s' . This ensures that P_1 does not “wait” at any other vertex than s' . In the third layer, P_1 must go from s' to z_1 through all clause gadgets. Since P_1 cannot visit a vertex twice, this validates whether the assignment satisfies ϕ . Fig. 2 depicts the resulting temporal graph.

The construction is done as follows. For each variable x_j , $1 \leq j \leq n$, we construct the *variable gadget* $\mathcal{G}^{x_j} = (\{a_{x_j}, a_{x_{j+1}}\} \cup \{x_j^i, \overline{x_j^i} \mid i \in [4]\}, (E_i^{x_j})_{i \in [2]})$, where

$$\begin{aligned} E_1^{x_j} &= E_T^{x_j} \cup E_F^{x_j} \text{ and} \\ E_2^{x_j} &= E_T^{x_j} \cup \{\{\overline{x_j^i}, \overline{x_j^{i+1}}\} \mid i \in [3]\} \cup \{\{\overline{x_j^4}, \overline{x_{j+1}^1}\} \mid j < n\} \text{ with} \\ E_T^{x_j} &= \{\{a_{x_j}, x_j^1\}, \{x_j^4, a_{x_{j+1}}\}\} \cup \{\{x_j^i, x_j^{i+1}\} \mid i \in [3]\} \text{ and} \\ E_F^{x_j} &= \{\{a_{x_j}, \overline{x_j^1}\}, \{\overline{x_j^4}, a_{x_{j+1}}\}\} \cup \{\{\overline{x_j^i}, \overline{x_j^{i+1}}\} \mid i \in [3]\}. \end{aligned}$$

Let $C_i = (\ell_j^f \vee \ell_p^g \vee \ell_q^h)$ be a clause, where ℓ_α^β is a literal of variable x_α and its β -th appearance when iterating the clauses in the order of the indices. We now




Fig. 2: An excerpt of layers $G_j = (V, E_j), j \in [3]$ of the temporal graph \mathcal{G} used in the proof of Theorem 2. Vertex z_1 and the vertices corresponding to the clauses are isolated in G_1 and G_2 and omitted in the illustration. In G_3 exemplary the edges for clause $C_i = (x_1 \vee x_2 \vee \bar{x}_3)$ are drawn, where C_i contains the fourth appearance of x_1 and the third appearance of x_2 and x_3 .

abuse our notation and say $\ell_\alpha^\beta \equiv \bar{x}_\alpha^\beta$ if ℓ_α^β is a negation of x_α and otherwise we say $\ell_\alpha^\beta \equiv x_\alpha^\beta$. We construct the *clause gadget* $\mathcal{G}^{C_i} := (\{C_i, C_{i+1}, \ell_j^f, \ell_p^g, \ell_q^h\}, (E_t^{C_i})_{t \in [3]})$ where $E_1^{C_i} = E_2^{C_i} = \emptyset$ and

$$E_3^{C_i} = \{\{C_r, \ell_j^f\}, \{C_r, \ell_p^g\}, \{C_r, \ell_q^h\} \mid r \in \{i, i+1\}\}.$$

Now we set $\mathcal{G} = (V, (E_t)_{t \in [3]})$, where

$$\begin{aligned} V &= \bigcup_{j=1}^n V(\mathcal{G}^{x_j}) \cup \bigcup_{i=1}^m V(\mathcal{G}^{C_i}) \cup \{s_1, s_2, z_1, z_2, s'\}, \\ E_1 &= \bigcup_{j=1}^n E_1^{x_j} \cup \{\{s_1, a_{x_1}\}, \{a_{x_{n+1}}, s'\}\}, \\ E_2 &= \bigcup_{j=1}^n E_2^{x_j} \cup \{\{s_2, x_1^1\}, \{x_n^4, a_{x_1}\}, \{a_{x_{n+1}}, z_2\}\}, \text{ and} \\ E_3 &= \bigcup_{i=1}^m E_3^{C_i} \cup \{\{s', C_1\}, \{z_1, C_{m+1}\}\}. \end{aligned}$$

Observe that I can be constructed in polynomial time.

Correctness. (\Leftarrow): Assume I is a *yes*-instance. Hence, there are non-intersecting P_1, P_2 such that P_i is a temporal (s_i, z_i) -path, for $i \in [2]$. Observe that there is only one unique temporal (s_2, z_2) -path in \mathcal{G} . Hence, P_1 visits s' , otherwise P_1 intersects P_2 . Moreover, on its way to s' the temporal path P_1 visits for each variable x_i either $x_i^1, x_i^2, x_i^3, x_i^4$ or $\bar{x}_i^1, \bar{x}_i^2, \bar{x}_i^3, \bar{x}_i^4$. We set variable x_i to true if and only if P_1 visits \bar{x}_i^1 . Now observe that P_1 visits all vertices C_1, C_2, \dots, C_{m+1} (in that order) at time 3.

8 *Interference-free Walks in Time: Temporally Disjoint Paths*

In particular, the only way from C_i to C_{i+1} leads via one of literals of C_i . Since P_1 cannot visit a vertex twice, one of the literals is set to true by our assignment of the variables. Thus, ϕ is satisfiable.

(\Rightarrow): Assume ϕ is satisfiable. We fix one assignment for the variables which satisfy ϕ . It is straightforward to verify that by construction \mathcal{G} contains one unique temporal (s_2, z_2) -path. Let P_2 be the unique temporal (s_2, z_2) -path in \mathcal{G} . Our P_1 starts with the time edge $(\{s_1, a_{x_1}\}, 1)$ and to get from a_{x_1} to $a_{x_{n+1}}$ at time one, P_1 takes for each $j \in [n]$ the induced path $(V, E_1)[\{a_{x_j}, a_{x_{j+1}}\} \cup \{x_j^i \mid i \in [4]\}]$ if x_j is set to false, otherwise P_1 takes the induced path $(V, E_1)[\{a_{x_j}, a_{x_{j+1}}\} \cup \{\bar{x}_j^i \mid i \in [4]\}]$. Afterwards, P_1 takes the time edges $(\{a_{x_{n+1}}, s'\}, 1), (\{s', C_1\}, 1)$. Let $i \in [m]$ and $C_i = (\ell_j^f \vee \ell_p^g \vee \ell_q^h)$ be a clause, where ℓ_α^β is a literal of variable x_α and its β -th appearance. To get from C_i to C_{i+1} at time three, P_1 takes the induced path $(V, E_3)[\{C_i, C_{i+1}, \ell_\alpha^\beta\}]$, where ℓ_α^β is a satisfied literal in C_i . Finally, P_1 takes the time edge $(\{C_{m+1}, z_1\}, 3)$. Note that P_1 does not visit a vertex twice, and hence it is a temporal (s_1, z_1) -path in \mathcal{G} . Moreover, it visits s' at times $\{1, 2, 3\}$ and all other vertices either at time one or time three. Note that P_2 is a temporal path that exclusively visits vertices at time two and does not visit vertex s' . Thus, P_1 is temporally non-intersecting with P_2 . It follows that I is a *yes*-instance. \square

The reduction behind Theorem 2 heavily relies on the fact that we are dealing with paths. Indeed, for temporally disjoint walks we presumably cannot obtain NP-hardness for a constant number of sink-source pairs since, as we will show at the end of this section, TEMPORALLY DISJOINT WALKS can be solved in polynomial time if the number of source-sink pairs is constant. However, before that we show W[1]-hardness for TEMPORALLY DISJOINT WALKS parameterized by the number $|S|$ of source-sink pairs, presumably excluding the existence of an FPT-algorithm for this parameter.

Theorem 3 TEMPORALLY DISJOINT WALKS is W[1]-hard when parameterized by $|S|$, even if all edges have exactly one time label.

Proof We present a parameterized reduction from DISJOINT PATHS ON DAGS. In this problem, we are given a directed acyclic graph (DAG) $D = (U, A)$ and ℓ source-sink pairs $\{(s_1, t_1), \dots, (s_\ell, t_\ell)\} \subseteq U \times U$, and we are asked whether there exist ℓ paths P_i from s_i to t_i that are pairwise vertex-disjoint. This problem parameterized by the number ℓ of paths is W[1]-hard [16].²

Construction. Let $(D = (U, A), \{(s_1, t_1), \dots, (s_\ell, t_\ell)\} \subseteq U \times U)$ be an instance of DISJOINT PATHS ON DAGS. We first compute a topological ordering σ for the vertices in U [41, Theorem 4.2.1]. Recall that σ is a linear ordering on the vertices in U with the property that $(u, v) \in A \Rightarrow u <_\sigma v$. We denote with $\sigma(u)$ the position of vertex u in the ordering σ .

Using this ordering, we construct a temporal graph \mathcal{G} as follows. We add the following vertices to V :

²Slivkins [16] shows W[1]-hardness for the arc-disjoint version of this problem. However, there are straightforward (parameterized) reductions between DISJOINT PATHS ON DAGS and ARC-DISJOINT PATHS ON DAGS.

- For every vertex $u \in U$ we add a “vertex-vertex” x_u to V .
- For every arc $(u, v) \in A$ we add an “arc-vertex” $y_{(u,v)}$ to V .

Next, we add the following time edges: For every arc $(u, v) \in A$ we add an edge from x_u to $y_{(u,v)}$ with time label $2\sigma(u) - 1$ and we add an edge from $y_{(u,v)}$ to x_v with label $2(\sigma(v) - 1)$.

Finally, for every source-sink pairs (s_i, t_i) of the DISJOINT PATHS ON DAGS instance, we add the source-sink pair (x_{s_i}, x_{t_i}) to S . The reduction can clearly be performed in polynomial time and the number of source-sink pairs in the constructed instance is the same as the number of source-sink pairs in the given DISJOINT PATHS ON DAGS instance.

Correctness. We now show that the given DISJOINT PATHS ON DAGS instance is a *yes*-instance if and only if our constructed instance is a *yes*-instance of TEMPORALLY DISJOINT WALKS.

(\Rightarrow): Assume $(D = (U, A), \{(s_1, t_1), \dots, (s_\ell, t_\ell)\}) \subseteq U \times U$ is a *yes*-instance of DISJOINT PATHS ON DAGS. Then there is a set \mathcal{P} of pairwise vertex-disjoint paths for all source-sink pairs $\{(s_1, t_1), \dots, (s_\ell, t_\ell)\}$. Let $P_i = ((v_{j-1}, v_j))_{j=1}^k \in \mathcal{P}$ with $v_0 = s_i$ and $v_k = t_i$ be the path from s_i to t_i . Then, by construction of \mathcal{G} , we have that $Q_i = ((x_{s_i}, y_{(s_i, v_1)}, 2\sigma(s_i) - 1), (y_{(s_i, v_1)}, x_{v_1}, 2(\sigma(v_1) - 1)), (x_{v_1}, y_{(v_1, v_2)}, 2\sigma(v_1) - 1), \dots, (y_{(v_{k-1}, t_i)}, x_{t_i}, 2(\sigma(t_i) - 1)))$ is a strict temporal path from x_{s_i} to x_{t_i} in \mathcal{G} that alternately visits the vertex-vertices and arc-vertices that correspond to the vertices and arcs visited by P_i in D . Since the paths $P_i \in \mathcal{P}$ are pairwise vertex disjoint (and hence also arc-disjoint), it is clear that the temporal paths $\{Q_i\}_{i=1}^\ell$ are vertex-disjoint, and hence also temporally disjoint.

(\Leftarrow): Note that a temporal walk can visit vertices multiple times. But there are only two ways in which a temporal walk W visits a vertex x_{v_j} multiple times: either W contains subwalk $((x_{v_j}, y_{(v_j, v_{j'})}, 2\sigma(v_j) - 1), (y_{(v_j, v_{j'})}, x_{v_{j'}}, 2(\sigma(v_{j'}) - 1)))$ for some j' , or W contains the subwalk $((x_{v_j}, y_{(v_j, v_{j'})}, 2\sigma(v_{j'}) - 1), (y_{(v_j, v_{j'})}, x_{v_j}, 2(\sigma(v_{j'}) - 1)))$ for some j' . In both cases we can remove the subwalk from W and still obtain a temporal walk from the same starting vertex to the same end vertex as W . We can remove subwalks of this kind repeatedly until we obtain a temporal path. The observations above allow us to assume that if we face a *yes*-instance of TEMPORALLY DISJOINT WALKS, then there is a set \mathcal{P} of pairwise temporally disjoint *paths* for all source-sink pairs in S .

Now observe that all vertices in \mathcal{G} are x_{v_j} -incident with time edges of exactly two time labels: $2\sigma(v_j) - 1$ and $2(\sigma(v_j) - 1)$. Hence, our paths “enter” a vertex-vertex x_{v_j} with a time edge that has time label $2\sigma(v_j) - 1$ and “leave” x_{v_j} with a time edge that has time label $2(\sigma(v_j) - 1)$. It follows that no two of the temporally disjoint paths in \mathcal{P} visit the same vertex-vertex x_{v_j} . The same holds for arc-vertices $y_{(v_j, v_{j'})}$ since they are only incident with two time edges. It follows that the temporal paths in \mathcal{P} are pairwise vertex-disjoint. Furthermore, by the construction of \mathcal{G} , they alternately visit vertex-vertices and arc-vertices, where the latter correspond to arcs in D that connect the vertices corresponding to the vertex-vertices visited directly before and after. It follows that we can translate temporal paths in \mathcal{P} directly to a set of pairwise vertex-disjoint paths in D that connect the corresponding source-sink pairs of the DISJOINT PATHS ON DAGS instance. \square

We contrast Theorem 3 by showing that TEMPORALLY DISJOINT WALKS is contained in XP for the parameter number $|S|$ of source-sink pairs.

Theorem 4 TEMPORALLY DISJOINT WALKS is in the class XP when parameterized by $|S|$, as it can be solved in $O(|V|^{2|S|+2} \cdot T)$ time.

Proof Consider an instance $I = (\mathcal{G}, S = \{(s_1, z_1), (s_2, z_2), \dots, (s_k, z_k)\})$ of TEMPORALLY DISJOINT WALKS. We use the following dynamic programming table D with Boolean entries. Intuitively, we want that for all $t \in \{1, \dots, T\}$ and $v_1, \dots, v_k \in V$ we have that $D[t, v_1, \dots, v_k] = \top$ if and only if there are temporally non-intersecting temporal (s_i, v_i) -walks P_1, \dots, P_k with arrival time $t_i \leq t$. However, for technical reasons, we have slightly stronger requirements for D . First of all, we have a “dummy” time label zero that we use to encode the sources in the dynamic program. Formally, we initialize D as follows:

For all $v_1, \dots, v_k \in V$ we have that

$$D[0, v_1, \dots, v_k] := \begin{cases} \top, & \text{if } \forall i \in [k]: v_i = s_i \\ \perp, & \text{otherwise.} \end{cases}$$

Furthermore, we have to model that the temporal walks we are looking for do not have to start immediately at their respective sources. Hence, if a temporal walk is still “waiting” at its source, the source vertex is not “blocked” for other temporal walks. We have a symmetric situation if temporal walks already arrived at their respective sink. In other words, if we have an entry $D[t, v_1, \dots, v_k]$ with $v_i = v_j$ for some $i \neq j$, then it is a necessary condition for $D[t, v_1, \dots, v_k] = \top$ that at least one of the two temporal walks i, j is either still waiting at its source or already arrived at its sink. In the latter case, we additionally need that the temporal walk arrived at the sink in a previous time step, otherwise the sink would still be blocked. We now look up in D whether there all these conditions are met for a set of temporally disjoint walks that arrive at some vertices at time $t - 1$ such that they can be extended in time step t to reach the vertices v_1, \dots, v_k .

Formally, for all $t \in [T]$ we have that $D[t, v_1, \dots, v_k] = \top$ if for all $i, j \in [k]$ with $v_i = v_j$ such that there exists $p \in \{i, j\}$ with $v_p \in \{s_p, z_p\}$, and there exist $u_1, \dots, u_k \in V$ the following holds:

- $\forall i, j \in [k], i \neq j, v_i = z_i: u_i = z_i \vee v_i \neq v_j$, and
- $D[t - 1, u_1, \dots, u_k] = \top$, and
- $((V, E_t), \{(u_i, v_i) \mid i \in [k], \{u_i, v_i\} \neq \{z_i\}, v_i \neq s_i\})$ is a *yes*-instance of DISJOINT PATHS.

Otherwise we have $D[t, v_1, \dots, v_k] = \perp$.

Here, DISJOINT PATHS is the problem, where we are given an undirected graph G and a set of k terminal pairs $\{(s'_i, z'_i) \mid i \in [k]\}$, and we ask whether there are k vertex-disjoint paths P_1, \dots, P_k in G such that P_i is an (s'_i, z'_i) -path and P_i and P_j are vertex-disjoint for all $i \in [k]$ and $j \in [k] \setminus \{i\}$. Note that $s'_i = z'_i$ is a valid input and that in this case s'_i is the only vertex on an (s'_i, z'_i) -path. We report that I is a *yes*-instance if and only if $D[T, z_1, \dots, z_k] = \top$.

Correctness. We show by induction that for all $t \in \{0, 1, \dots, T\}$ and $v_1, \dots, v_k \in V$ we have that $D[t, v_1, \dots, v_k] = \top$ if and only if there are temporally non-intersecting temporal walks P_1, \dots, P_k such that

- (a) P_i is a temporal (s_i, v_i) -walk in \mathcal{G} with arrival time $t_i \leq t$, and
- (b) $\forall i, j \in [k], v_i \neq v_j: \exists p \in \{i, j\}: v_p = s_p \vee (v_p = z_p \wedge t_p < t)$ is true.

Here, we say that an empty edge list (\emptyset) is a temporal (v, v) -walk with arrival time 0 which does not visit any vertex, for all $v \in V$. Hence, for $t = 0$ this claim is true by the definition of D . Now let $t \in [T]$ and assume that this claim holds true for all $t' < t$.

(\Leftarrow): Assume that there are temporally non-intersecting temporal walks P_1, \dots, P_k such that (a) and (b) hold true. Let $K_s := \{i \in [k] \mid s_i = v_i\}$, $K_z := \{i \in [k] \mid z_i = v_i, t_i < t - 1\}$, and $K := [k] \setminus (K_s \cup K_z)$. We set $u_i = s_i$, for all $i \in K_s$. We set $u_i = z_i$, for all $i \in K_z$. For all $i \in K$, let u_i be the first vertex which is visited at time t (in the visiting order), or v_i if P_i does not visit any vertex at time t . Observe that $\forall i, j \in [k], v_i = v_j : \exists p \in \{i, j\} : v_p \in \{s_p, z_p\}$ and $\forall i, j \in [k], i \neq j, v_i = z_i : u_i = z_i \vee v_i \neq v_j$ are both true. Observe for all $i \in K$ that we can split P_i into two parts P_i^1 and P_i^2 , where P_i^1 is a (possibly empty) temporal (s_i, u_i) -walk in \mathcal{G} with arrival time at most $t - 1$ and P_i^2 is a (possibly empty) (u_i, v_i) -path in the graph (V, E_t) . Note that $\{P_i^2 \mid i \in K\}$ is a solution for the DISJOINT PATH instance $((V, E_t), \{(u_i, v_i) \mid i \in [k], \{u_i, v_i\} \neq \{z_i\}, v_i \neq s_i\})$. We set $P_i^1 = (\emptyset)$ for all $i \in K_s$, and $P_i^1 = P_i$ for all $i \in K_z$. Clearly, P_1^1, \dots, P_k^1 are temporally non-intersecting such that

- (a) P_i is a temporal (s_i, u_i) -walk in \mathcal{G} with arrival time $t_i \leq t - 1$, and
- (b) $\forall i, j \in [k], u_i \neq u_j : \exists p \in \{i, j\} : u_p = s_p \vee (u_p = z_p \wedge t_p < t - 1)$.

Hence, $D[t-1, u_1, \dots, u_k] = \top$ and thus by definition of D we have $D[t, v_1, \dots, v_k] = \top$.

(\Rightarrow): Assume that $D[t, v_1, \dots, v_k] = \top$. By definition of D , we know that there are vertices u_1, \dots, u_k such that

- (i) $\forall i, j \in [k], v_i = v_j : \exists p \in \{i, j\} : v_p \in \{s_p, z_p\}$
- (ii) $\forall i, j \in [k], i \neq j, v_i = z_i : u_i = z_i \vee v_i \neq v_j$ are true,
- (iii) $D[t - 1, u_1, \dots, u_k] = \top$, and
- (iv) $((V, E_t), \{(u_i, v_i) \mid i \in [k], \{v_i, u_i\} \neq \{z_i\}, v_i \neq s_i\})$ is a *yes*-instance of DISJOINT PATH.

By assumption and (iii), there are temporally non-intersecting temporal walks P_1, \dots, P_k such that

- (a) P_i is a temporal (s_i, u_i) -walk in \mathcal{G} with arrival time $t_i \leq t - 1$, and
- (b) $\forall i, j \in [k], u_i \neq u_j : \exists p \in \{i, j\} : u_p = s_p \vee (u_p = z_p \wedge t_p < t - 1)$.

Let $K_s := \{i \in [k] \mid s_i = v_i\}$, $K_z := \{i \in [k] \mid u_i = z_i = v_i\}$, and $K = [k] \setminus (K_s \cup K_z)$. For all $i \in K_s$ we set $P_i' = (\emptyset)$. For all $i \in K_z$ we set $P_i' = P_i$. Due to (iv), there are vertex-disjoint paths $P_1^2, \dots, P_{|K|}^2$ such that P_i^2 is a (u_i, v_i) -path in (V, E_t) , for all $i \in K$. For all $i \in K$ we construct a temporal (s_i, v_i) -walk P_i' by concatenating P_i and P_i^2 . Note that P_1', \dots, P_k' are temporally non-intersecting such that (a) P_i' is a temporal (s_i, v_i) -walk with arrival time $t'_i \leq t$, for all $i \in [k]$. Moreover, if we have for some $i, j \in [k]$ that $v_i = v_j$, then there is a $p \in \{i, j\}$ such that either $v_p = s_p$ or $v_p = z_p$, because of (i). If $v_p = z_p$, then $u_p = z_p$ and thus $t'_i < t$, because of (ii). Hence, (b) $\forall i, j \in [k], i \neq j : v_i \neq v_j \vee v_i = s_i \vee (v_i = z_i \wedge t'_i < t)$ is true.

So, I is a *yes*-instance if and only if $D[T, z_1, \dots, z_k] = \top$.

Running time. The table size of D is $O(|V|^k \cdot T)$. To compute one entry of D we look at up to $O(|V|^k)$ other entries of D and then solve an instance of DISJOINT

PATHS in $O(n^2)$ time³ [6]. This gives an overall running time of $O(|V|^{2k+2} \cdot T)$ if k is a fixed constant. \square

Finally, we point out that Theorem 2 implies that for TEMPORALLY DISJOINT PATHS we presumably cannot achieve a result similar to Theorem 4 while Theorem 3 implies that we presumably cannot improve this result on a classification level, that is, we cannot expect to find an FPT-algorithm for TEMPORALLY DISJOINT WALKS for parameter $|S|$.

4 Temporal lines and trees

In this section, we investigate the computational complexity of TEMPORALLY DISJOINT (PATHS/WALKS) for restricted classes of underlying graphs, in particular so-called *temporal lines* and *temporal trees*. The former are temporal graphs that have a path as underlying graph and the latter are temporal graphs that have a tree as underlying graph. In particular, we first show that, surprisingly, the problems remain NP-hard on temporal lines (and thus also on temporal trees). On the positive side, we show that, on temporal trees, TEMPORALLY DISJOINT PATHS is fixed-parameter tractable with respect to the number of source-sink pairs. This result stands in stark contrast to the general case, where the problem is NP-hard even when the number of source-sink pairs is two (Theorem 2). If we further restrict all source-sink pairs to consist of the two end-points of the temporal line, however, then we obtain a polynomial-time algorithm.

Before we proceed with our results in this section, first we recall some useful background. Given a temporal graph \mathcal{G} and two specific vertices s, z of it, a *foremost* temporal path from s to z starting at time t is a temporal path which starts at s not earlier than at time t and arrives at z with the earliest possible arrival time. A foremost temporal path from s to z starting at time t can be computed in linear $O\left(|V| + \sum_{i=1}^T |E_i|\right)$ time [27].

Theorem 5 TEMPORALLY DISJOINT (PATHS/WALKS) is NP-hard even on a temporal line where all temporal paths are to the same direction.

Proof We present here a polynomial-time reduction for TEMPORALLY DISJOINT WALKS. Towards the end of this proof, we argue that the same reduction also works for TEMPORALLY DISJOINT PATHS. The reduction is done from MULTICOLORED INDEPENDENT SET ON UNIT INTERVAL GRAPHS, which is known to be NP-complete [42, Lemma 2]. In this problem, the input is a unit interval graph $G = (V, E)$ with n vertices, where V is partitioned into k subsets of independent vertices; we interpret each of these subsets as a vertex color. The goal is to compute an independent set of size k in G which contains exactly one vertex from each color. By possibly slightly shifting the endpoints of the intervals in the given unit interval representation of G , we

³This running time bound considers k to be a constant. There is a factor $f(k)$ for some function f hidden in the O -notation.

can assume without loss of generality that all endpoints of the intervals are distinct. Furthermore, we can assume without loss of generality that each interval endpoint is an integer between $k + 1$ and $k + n^2$ (while all intervals still have the same length).

Construction. From the given multi-colored unit intervals in G , we construct a temporal line \mathcal{P} using the following procedure. Let $\{c_1, \dots, c_k\}$ be the set of all colors of the intervals in G . First we fix an arbitrary linear ordering $c_1 < c_2 < \dots < c_k$ of the k colors, and we add to the underlying path P of \mathcal{P} two vertices v_i^1 and v_i^2 , for every color c_i . We add to P also three basis vertices v_ℓ, v^*, v_r . The vertices of P are ordered starting from $v_1^1, v_2^1, \dots, v_k^1$, followed by the basis vertices v_ℓ, v^*, v_r , and finishing with $v_1^2, v_2^2, \dots, v_k^2$. At the end we have $P = (v_1^1, v_2^1, \dots, v_k^1, v_\ell, v^*, v_r, v_1^2, v_2^2, \dots, v_k^2)$.

We construct the multiset S of source-sink pairs as follows. Let m_i be the number of intervals of color c_i . For every color c_i we add the pair (v_i^1, v_i^2) to S . We refer to this source-sink pair as “the verification source-sink pair for color c_i ”. Furthermore, we add $m_i - 1$ copies of the pair (v_i^1, v_ℓ) to S and we add $m_i - 1$ copies of the pair (v_r, v_i^2) to S . We call these $2m_i - 2$ source-sink pairs the “dummy source-sink pairs for color c_i ”.

To fully define the temporal line \mathcal{P} , we still need to add time labels to the edges of P . Denote by a_i^j and b_i^j the start and end points of the j th interval of color c_i . We set up the edge labels of path P from v_i^1 to v_i^2 as follows. To edge $\{v_s^1, v_{s+1}^1\}$ with $s \in [k - 1]$, we add the labels a_i^j with $i \leq s$. To edges $\{v_k^1, v_\ell\}$ and $\{v_\ell, v^*\}$, we add all labels a_i^j . To edge $\{v_s^2, v_{s+1}^2\}$ with $s \in [k - 1]$, we add the labels b_i^j with $i > s$. To edges $\{v^*, v_r\}$ and $\{v_r, v_1^2\}$, we add all labels b_i^j . See Fig. 3 for an example. The construction can clearly be performed in polynomial time.

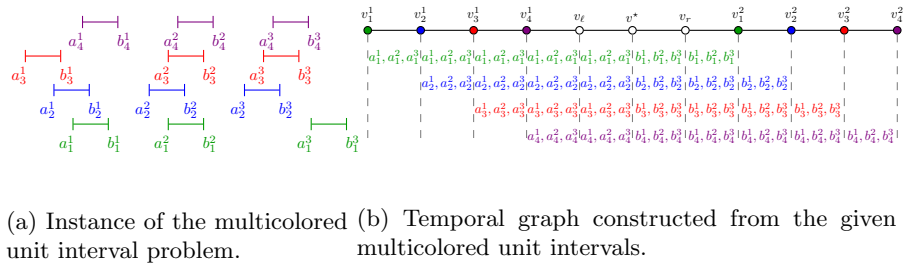


Fig. 3: Example of the reduction described in the proof of Theorem 5.

Correctness. (\Rightarrow): Assume there is a multicolored independent set $V' \subseteq V$ in G . Let $v_i \in V'$ be the vertex in the independent set with color c_i and let $[a_i^j, b_i^j]$ be the interval of v_i . Then for the verification source-sink pair of c_i we use the following temporal path: $((v_i^1, v_{i+1}^1, a_i^j), (v_{i+1}^1, v_{i+2}^1, a_i^j), \dots, (v_{k-1}^1, v_k^1, a_i^j), (v_k^1, v_\ell, a_i^j), (v_\ell, v^*, a_i^j), (v^*, v_r, b_i^j), (v_r, v_1^2, b_i^j), (v_1^2, v_2^2, b_i^j), \dots, (v_{i-1}^2, v_i^2, b_i^j))$. For the dummy source-sink pairs (v_i^1, v_ℓ) of c_i we use the temporal paths $((v_i^1, v_{i+1}^1, a_i^{j'}), \dots, (v_{k-1}^1, v_k^1, a_i^{j'}), (v_k^1, v_\ell, a_i^{j'}))$ with $j' \neq j$. Note that there are exactly $m_i - 1$ pairwise different paths of this kind. Analogously, for the dummy source-sink pairs (v_r, v_i^2)

of c_i we use the temporal paths $((v_r, v_1^2, b_i^{j'}), (v_1^2, v_2^2, b_i^{j'}), \dots, (v_{i-1}^2, v_i^2, b_i^{j'}))$ with $j' \neq j$. It is easy to check that the temporal paths for the dummy source-sink pairs of all colors do not temporally intersect. Now assume, for the sake of contradiction, that temporal paths of two verification source-sink pairs of colors c_i and $c_{i'}$ temporally intersect. Then they have to intersect in v^* , since this is the only vertex where the paths wait. By construction, the temporal path for the verification source-sink pairs of color c_i visits v^* during the interval $[a_i^j, b_i^j]$ and the verification source-sink pairs of color $c_{i'}$ visits v^* during the interval $[a_{i'}^{j'}, b_{i'}^{j'}]$. These two intervals correspond to the intervals of the vertices of colors c_i and $c_{i'}$ in the multicolored independent set V' . Hence, those intervals intersecting is a contradiction to the assumption that V' is in fact an independent set.

(\Leftarrow): Assume we have a set of pairwise temporally disjoint walks for all source-sink pairs in S . Note that all edges except $\{v_\ell, v^*\}$ and $\{v^*, v_r\}$ have as many time labels as temporal walks that need to go through them. Furthermore, note that $\{v_\ell, v^*\}$ has the same labels as $\{v_k^1, v_\ell\}$ and $\{v^*, v_r\}$ has the same labels as $\{v_r, v_1^2\}$. This in particular implies that all temporal walks are in fact paths since the only vertex that could be visited by a path for more than one time step is v^* . Therefore, for every pair $(s, z) \in S$, no temporal path from s to z can ever stop and wait at any vertex different from v^* . Furthermore, the only paths going through vertex v^* are the paths connecting vertices v_i^1 and v_i^2 (which correspond to color c_i); we will refer to this path as the color path of c_i . Consider color c_1 and its dummy source-sink pairs (v_1^1, v_ℓ) . By construction, the edge $\{v_1^1, v_1^2\}$ has time labels corresponding to the start points a_1^j of intervals from the m_1 vertices of G that have color c_1 . It follows that the temporal paths for these dummy source-sink pairs and the color path of c_1 use only time labels corresponding to the start points a_1^j of intervals from the m_1 vertices of G that have color c_1 until they are at v_ℓ or arrive at v^* , respectively, since they cannot wait at any vertex. Now by induction, this holds for all other colors c_i and by an analogous argument, this also holds for the “second half”. More specifically, we also have that temporal paths for the dummy source-sink pairs (v_r, v_i^2) as well as the “second part” of the color path of c_i use time labels corresponding to end points b_i^j of intervals from the vertices of G that have color c_i when going from v_r (respectively v^*) to their corresponding destinations.

It follows that each color path can enter and leave vertex v^* only at the time corresponding to the start and end points of its color intervals. In any other case some of the other vertices are blocked, which prevents the completion of other temporal S -paths. Recall that intervals of the same color are non-overlapping. Hence, for every color path corresponding to a color c_i we can find one interval $[a_i^j, b_i^j]$ such that the color path visits v^* in an interval that includes $[a_i^j, b_i^j]$. Since the color paths are temporally non-intersecting, the vertices corresponding to the intervals form a multicolored independent set in G .

This completes the proof for the case of TEMPORALLY DISJOINT WALKS. As, in the constructed reduction, all walks are actually just paths, it follows that also TEMPORALLY DISJOINT PATHS is NP-hard. \square

NP-hardness even in the case of temporal lines motivates to study the potential for parameterized tractability results. Next, we show fixed-parameter tractability of TEMPORALLY DISJOINT PATHS parameterized by the number $|S|$ of source-sink pairs if the underlying graph is a tree.




Fig. 4: Temporally disjoint walks on a temporal line are not necessarily equal to temporally disjoint paths. Suppose that one wants to determine edge-labels of the following walks: P_1 from v_1 to v_7 , P_2 from v_2 to v_4 and P_3 from v_7 to v_5 on the depicted temporal graph. In the feasible solution P_2 and P_3 are temporal paths, but P_1 has to be a walk.

Theorem 6 TEMPORALLY DISJOINT PATHS on temporal trees is fixed-parameter tractable when parameterized by $|S|$, as it can be solved in $O(|S|^{|S|+2} \cdot |G|)$ time.

Proof Let $I = (G, S)$ be an instance of TEMPORALLY DISJOINT PATHS, the underlying graph G being a tree and S consisting of k source-sink pairs $(s_1, z_1), \dots, (s_k, z_k)$. We solve I using the following procedure.

First we can observe that, since G is a tree, every source-sink pair (s_i, z_i) in S corresponds to exactly one path P_i in G . Furthermore, if in a tree two paths P_i and P_j intersect, then their intersection $P_i \cap P_j$ is a path in G (potentially containing only one vertex). If P_i and P_j intersect, then there are two ways that their intersection can be traversed: either first by P_i and then P_j , or P_j and then P_i . We enumerate all possible permutations π of the k source-sink pairs (s_i, z_i) . For every permutation π and for every $i = 1, 2, \dots, k$, we compute the foremost path $P_i(\pi)$ from s_i to z_i (i.e., a temporal path with the earliest arrival time). Let v_x be an arbitrary internal vertex of $P_i(\pi)$, and suppose that v_x is visited by $P_i(\pi)$ within the time interval $[a_x, b_x]$. Then, for all the edges that are incident to v_x , we remove all labels $\ell \leq b_x$, as these temporal edges cannot be used by any further temporal path $P_j(\pi)$, where $j > i$. We proceed by computing the foremost path $P_{i+1}(\pi)$ from s_{i+1} to z_{i+1} which only uses unmarked temporal edges. The permutation π leads to a feasible routing of the paths between the k source-sink pairs if and only if we can compute all these k foremost paths $P_1(\pi), P_2(\pi), \dots, P_k(\pi)$ as described above.

During the above procedure we construct $O(k!) = O(k^{k+1})$ different permutations π . For every permutation we calculate $|S| = k$ foremost paths, each in $O(|V| + \sum_{i=1}^T |E_i|)$ time [27]. In total, all above computations can be done in $O(|S|^{|S|+2} \cdot (|V| + \sum_{i=1}^T |E_i|))$ time. \square

We remark that it remains open whether a similar result can be obtained for TEMPORALLY DISJOINT WALKS, since we cannot assume w.l.o.g. that the temporally disjoint walks are actually paths, even on temporal lines (for an example see Fig. 4). Presumably (and in contrast to the general case) the walk version is computationally more difficult than the path version of our problem on temporal paths and trees.

Finally, we show that we can solve TEMPORALLY DISJOINT (PATHS/WALKS) in polynomial time if the underlying graph is a path and all source-sink pairs consist of the endpoints of that path.

Theorem 7 Let \mathcal{G} be a temporal line having $P = (v_0, v_1, v_2, \dots, v_n)$ as its underlying path. If S contains k times the source-sink pair (v_0, v_n) and $\ell = |S| - k$ times the source-sink pair (v_n, v_0) , then TEMPORALLY DISJOINT (PATHS/WALKS) can be solved on \mathcal{G} in $O(k\ell(k + \ell) \cdot |\mathcal{G}|)$ time.

Proof We first consider the problem version TEMPORALLY DISJOINT PATHS. Let $I = (\mathcal{P}, S)$ be an instance of TEMPORALLY DISJOINT PATHS, where \mathcal{P} is a given temporal line with $P = (v_0, v_1, v_2, \dots, v_n)$ as its underlying path. Assume that there have to be k (resp. $\ell = |S| - k$) temporally disjoint (v_0, v_n) - (resp. (v_n, v_0) -) paths in the output, i. e., they must have the orientation from v_0 to v_n , (resp. from v_n to v_0).

We solve the instance I using dynamic programming. The main idea is that, since all temporal paths start and end in endpoints of P , in any optimal solution, once a temporal path starts, it proceeds in the fastest possible way (without interfering with previously started paths). Therefore, assuming we start with (v_0, v_n) -temporal paths, we only need to find out *how many* (v_0, v_n) -temporal paths follow the starting path, after that *how many* (v_n, v_0) -temporal paths follow, then after that *how many* (v_0, v_n) -temporal paths follow, etc.

Let $0 \leq i \leq k$, $0 \leq j \leq \ell$, and $1 \leq t \leq T$. Then $L(i, j, t)$ denotes the earliest arrival time of $(k - i) + (\ell - j)$ temporally non-intersecting temporal paths with $k - i$ being (v_0, v_n) -temporal paths and $\ell - j$ being (v_n, v_0) -temporal paths, assuming that the earliest-starting temporal path is a (v_0, v_n) -temporal path that starts at time t . If it is not possible to route such $(k - i) + (\ell - j)$ temporally non-intersecting temporal paths starting at time t , then let $L(i, j, t) = \infty$. Similarly we define $R(i, j, t)$, with the only difference that here the earliest-starting temporal path needs to start at time t from v_n and finishes at v_0 . For the sake of completeness, we let $L(i, j, \infty) = R(i, j, \infty) = \infty$ for every $i \leq k$ and every $j \leq \ell$. Furthermore, for every t , every $i \leq k - 1$, and every $j \leq \ell - 1$, we let $L(k, j, t) = R(i, \ell, t) = \infty$. Finally we let $L(k, \ell, t) = R(k, \ell, t) = t - 1$. Note that, the input instance I is a *yes*-instance if and only if $\min\{L(0, 0, 1), R(0, 0, 1)\} \neq \infty$. Furthermore, note that, for every triple i, j, t , the value $\min\{L(i, j, t), R(i, j, t)\}$ is the earliest arrival time of all temporal paths in the subproblem where, until time $t - 1$, exactly i and j temporally non-intersecting temporal (v_0, v_n) - and (v_n, v_0) -paths, respectively, have been routed.

The value $L(i, j, t)$ can be recursively computed as follows. Suppose that, in the optimal solution, $1 \leq p \leq k - i$ temporally non-intersecting (v_0, v_n) -temporal paths are first routed (starting at time t) before the first (v_n, v_0) -temporal path (among the $\ell - j$ ones) is routed. Let t_p be the earliest arrival time of these p paths if they can all be routed; if not, then we set $t_p = \infty$. Then:

$$L(i, j, t) = \min\{R(i + p, j, t_p + 1) \mid 1 \leq p \leq k - i\}. \quad (1)$$

The value $R(i, j, t)$ can be computed similarly:

$$R(i, j, t) = \min\{L(i, j + p, t_p^* + 1) \mid 1 \leq p \leq \ell - j\}, \quad (2)$$

where (v_n, v_0) -temporal paths are routed.

The values $\{t_p \mid 1 \leq p \leq k - i\}$ can be computed as follows. If $p = 1$, then t_p is the arrival time of the (v_0, v_n) -foremost temporal path P_1 . To determine t_2 , we first compute P_1 and then, for every internal vertex v_x of \mathcal{P} , if v_x is visited by P_1 within the time interval $[a_x, b_x]$, then we remove from the edges $\{v_{x-1}, v_x\}, \{v_x, v_{x+1}\}$ of \mathcal{P} all labels $l \leq b_x$. In the resulting temporal line we then compute the foremost temporal path P_2 , which arrives at v_n at time t_2 . By applying this procedure iteratively, we

either compute p temporally non-intersecting temporal paths P_1, P_2, \dots, P_p , starting at time t and arriving at time t_p , or we conclude that $t_p = \infty$. The values $\{t_p^* \mid 1 \leq p \leq \ell - j\}$ (for the (v_n, v_0) -temporal paths) can be computed in a symmetric way. All these computations together can be done in linear time.

From the above it follows that we can decide TEMPORALLY DISJOINT PATHS by checking whether $\min\{L(0, 0, 1), R(0, 0, 1)\}$ is finite or not. In total, there are $2k\ell T$ values $L(i, j, t)$ and $R(i, j, t)$. Observe that, for every pair i, j , we only need to compute the value $L(i, j, t)$ (resp. $R(i, j, t')$) for one specific value of t (resp. t'). This observation ensures that the running time of the algorithm is polynomial. For this we need the following observation. Assume that, in the recursion tree originated at $L(0, 0, 1)$, we need to compute at two different places the values $L(i, j, t)$ and $L(i, j, t')$, where $t' > t$. Then, since obviously $L(i, j, t) \leq L(i, j, t')$, at the second place we can just replace $L(i, j, t')$ by ∞ , thus stopping the recursive calculations at that branch of the recursion tree. Similarly, if we need to compute at two different places the values $R(i, j, t)$ and $R(i, j, t')$, where $t' > t$, we replace $R(i, j, t')$ by ∞ at the second place of the recursion tree. That is, for every pair i, j , we just need to compute only one value $L(i, j, t)$ (resp. $R(i, j, t)$). Therefore, we can build two matrices M^L and M^R , each of size $(k+1) \times (\ell+1)$, such that $M^L(i, j)$ (resp. $M^R(i, j)$) stores the unique value of t for which we need to compute $L(i, j, t)$ (resp. $R(i, j, t)$). That is, in the recursion tree originate at $L(0, 0, 1)$, for every pair i, j we only need to compute the values $L(i, j, M^L(i, j))$ and $R(i, j, M^R(i, j))$.

Similarly, for the recursion tree originating at $R(0, 0, 1)$ we need to build two other matrices N^L and N^R (each of size $(k+1) \times (\ell+1)$) for the same purpose, as the recursion tree originated at $R(0, 0, 1)$ is different to the one originated at $L(0, 0, 1)$. That is, in the recursion tree originating at $R(0, 0, 1)$, for every pair i, j we only need to compute the values $L(i, j, N^L(i, j))$ and $R(i, j, N^R(i, j))$.

Each of these four $(k+1) \times (\ell+1)$ matrices can be computed by running $O(k\ell(k+\ell))$ times the foremost temporal path algorithm (in order to compute at each step in linear time $O(|V| + \sum_{i=1}^T |E_i|)$ the values $\{t_p \mid 1 \leq p \leq k - i\}$ and $\{t_p^* \mid 1 \leq p \leq \ell - j\}$, respectively). Once we have built these four matrices, we can iteratively compute the value $L(0, 0, 1)$ (resp. $R(0, 0, 1)$) in at most $k\ell$ computations, each of which takes at most $O(k+\ell)$ time (see equations (1)-(2)). Thus, all computations can be done in $O(k\ell(k+\ell) \cdot (|V| + \sum_{i=1}^T |E_i|))$ time.

This completes the proof for the case of TEMPORALLY DISJOINT PATHS. Finally, it is easy to see that in the problem TEMPORALLY DISJOINT WALKS, in any optimal solution every temporal walk is a temporal path, as every temporal walk is from v_0 to v_n or from v_n to v_0 . Hence, the above algorithm for TEMPORALLY DISJOINT PATHS also solves TEMPORALLY DISJOINT WALKS. \square

5 Conclusion

Formally introducing temporally disjoint paths and walks, we modeled the property that agents moving along these never meet, even though they might visit the same vertices. We identified an unexpected difference in their computational complexity: TEMPORALLY DISJOINT PATHS is NP-hard even for two paths, while TEMPORALLY DISJOINT WALKS can be solved in polynomial time for a constant number of walks (however it becomes W[1]-hard when parameterized by the number of walks). On the contrary, while TEMPORALLY

DISJOINT PATHS becomes fixed-parameter tractable for the number of paths if the underlying graph is a path, we leave open whether we can obtain a similar result for TEMPORALLY DISJOINT WALKS which seems to be much more complicated than the path version.

Table 1 surveys our main results and the central (concrete) open question of this work. Furthermore, we leave open whether TEMPORALLY DISJOINT (PATHS/WALKS) is in FPT or W[1]-hard for parameters of the underlying graph that are unbounded on paths, such as the vertex cover number or the treedepth. Employing *temporal* graph parameters such as the so-called timed feedback vertex number [25] might also be worthy to investigate in future research.

Lastly, we believe that this work can be a starting point for the investigation of many well-motivated variants or generalizations of our problem. One can e.g. consider the case where a certain set of vertices need to be visited or a certain “amount” of intersection is acceptable. It is also of interest to investigate our problem for restricted temporal path model such as so-called restless temporal paths or walks [23, 25, 43].

Supplementary information. Not applicable.

Acknowledgments. George B. Mertzios was supported by the EPSRC grant EP/P020372/1 and by DFG RTG 2434 while visiting TU Berlin. Hendrik Molter was supported by the German Research Foundation (DFG), project MATE (NI 369/17), and by the Israeli Science Foundation (ISF), grant No. 1070/20. The main part of this work was done while Hendrik Molter was affiliated with TU Berlin.

Declarations

Not applicable.

References

- [1] Klobas, N., Mertzios, G.B., Molter, H., Niedermeier, R., Zschoche, P.: Interference-free walks in time: Temporally disjoint paths. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21, pp. 4090–4096 (2021)
- [2] Karp, R.M.: On the computational complexity of combinatorial problems. *Networks* **5**(1), 45–68 (1975)
- [3] Neil, R., Seymour, P.D.: Disjoint paths—a survey. Society for Industrial and Applied Mathematics. *Journal on Algebraic and Discrete Methods* **6**(2), 300–305 (1985)
- [4] Neil, R., Seymour, P.D.: Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B* **63**(1), 65–110 (1995)

- [5] Downey, R.G., Fellows, M.R.: *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, London (2013)
- [6] Kawarabayashi, K.-i., Kobayashi, Y., Reed, B.: The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B* **102**(2), 424–435 (2012)
- [7] Stern, R.: Multi-agent path finding - an overview. In: *Artificial Intelligence - 5th RAAI Summer School*, pp. 96–115. Springer, Dolgoprudny, Russia (2019)
- [8] Stern, R., Sturtevant, N.R., Felner, A., Koenig, S., Ma, H., Walker, T.T., Li, J., Atzmon, D., Cohen, L., Kumar, T.K.S., Boyarski, E., Barták, R.: Multi-agent pathfinding: Definitions, variants, and benchmarks. In: *Proceedings of the Twelfth International Symposium on Combinatorial Search, SOCS 2019*, pp. 151–159 (2019)
- [9] Chuzhoy, J., Kim, D.H.K., Nimavat, R.: New hardness results for routing on disjoint paths. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 86–99 (2017)
- [10] Golovach, P.A., Thilikos, D.M.: Paths of bounded length and their cuts: parameterized complexity and algorithms. *Discrete Optimization* **8**(1), 72–86 (2011)
- [11] Fluschnik, T., Kratsch, S., Niedermeier, R., Sorge, M.: The parameterized complexity of the minimum shared edges problem. *Journal of Computer and System Sciences* **106**, 23–48 (2019)
- [12] Fluschnik, T., Morik, M., Sorge, M.: The complexity of routing with collision avoidance. *Journal of Computer and System Sciences* **102**, 69–86 (2019)
- [13] Guo, L., Deng, Y., Liao, K., He, Q., Sellis, T.K., Hu, Z.: A fast algorithm for optimally finding partially disjoint shortest paths. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1456–1462 (2018)
- [14] Tao, B., Xiao, M., Zhao, J.: Finding minimum-weight link-disjoint paths with a few common nodes. In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI '20)*, pp. 938–945 (2020)
- [15] Fortune, S., Hopcroft, J., Wyllie, J.: The directed subgraph homeomorphism problem. *Theoretical Computer Science* **10**(2), 111–121 (1980)
- [16] Slivkins, A.: Parameterized tractability of edge-disjoint paths on directed acyclic graphs. *SIAM Journal of Discrete Mathematics* **24**(1), 146–157

- (2010)
- [17] Almagor, S., Lahijanian, M.: Explainable multi agent path finding. In: Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, pp. 34–42 (2020)
 - [18] Atzmon, D., Stern, R., Felner, A., Wagner, G., Barták, R., Zhou, N.-F.: Robust multi-agent path finding and executing. *Journal of Artificial Intelligence Research* **67**, 549–579 (2020)
 - [19] Standley, T.S.: Finding optimal solutions to cooperative pathfinding problems. In: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, pp. 173–178 (2010)
 - [20] Dondi, R., Sikora, F.: Finding disjoint paths on edge-colored graphs: more tractability results. *Journal of Combinatorial Optimization* **36**, 1315–1332 (2017)
 - [21] Santos, R.F., Andrioni, A., Drummond, A.C., Xavier, E.C.: Multicolour paths in graphs: NP-hardness, algorithms, and applications on routing in WDM networks. *Journal of Combinatorial Optimization* **33**(2), 742–778 (2017)
 - [22] Wu, B.Y.: On the maximum disjoint paths problem on edge-colored graphs. *Discrete Optimization* **9**(1), 50–57 (2012)
 - [23] Bentert, M., Himmel, A., Nichterlein, A., Niedermeier, R.: Efficient computation of optimal temporal walks under waiting-time constraints. *Applied Network Science* **5**(1), 73 (2020)
 - [24] Enright, J., Meeks, K., Mertzios, G.B., Zamaraev, V.: Deleting edges to restrict the size of an epidemic in temporal networks. *Journal of Computer and System Sciences* **119**, 60–77 (2021)
 - [25] Casteigts, A., Himmel, A., Molter, H., Zschoche, P.: Finding temporal paths under waiting time constraints. *Algorithmica* **83**(9), 2754–2802 (2021)
 - [26] Mertzios, G.B., Michail, O., Spirakis, P.G.: Temporal network optimization subject to connectivity constraints. *Algorithmica* **81**(4), 1416–1449 (2019)
 - [27] Wu, H., Cheng, J., Ke, Y., Huang, S., Huang, Y., Wu, H.: Efficient algorithms for temporal path computation. *IEEE Transactions on Knowledge and Data Engineering* **28**(11), 2927–2942 (2016)
 - [28] Fuchsle, E., Molter, H., Niedermeier, R., Renken, M.: Delay-robust routes

- in temporal graphs. In: Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science (STACS) (2022). Accepted for publication.
- [29] Füchsle, E., Molter, H., Niedermeier, R., Renken, M.: Temporal connectivity: Coping with foreseen and unforeseen delays. In: Proceedings of the 1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND) (2022). Accepted for publication.
- [30] Enright, J.A., Meeks, K., Molter, H.: Counting temporal paths. *CoRR* **abs/2202.12055** (2022)
- [31] Kempe, D., Kleinberg, J., Kumar, A.: Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences* **64**(4), 820–842 (2002)
- [32] Fluschnik, T., Molter, H., Niedermeier, R., Renken, M., Zschoche, P.: Temporal graph classes: A view through temporal separators. *Theoretical Computer Science* **806**, 197–218 (2020)
- [33] Zschoche, P., Fluschnik, T., Molter, H., Niedermeier, R.: The complexity of finding separators in temporal graphs. *Journal of Computer and System Sciences* **107**, 72–92 (2020)
- [34] Mertzios, G.B., Molter, H., Zamaraev, V.: Sliding window temporal graph coloring. In: AAI, pp. 7667–7674 (2019)
- [35] Mertzios, G.B., Molter, H., Niedermeier, R., Zamaraev, V., Zschoche, P.: Computing maximum matchings in temporal graphs. In: Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science (STACS), pp. 27–12714 (2020)
- [36] Akrida, E.C., Mertzios, G.B., Spirakis, P.G., Zamaraev, V.: Temporal vertex cover with a sliding time window. *Journal of Computer and System Sciences* **107**, 108–123 (2020)
- [37] Hamm, T., Klobas, N., Mertzios, G.B., Spirakis, P.G.: The complexity of temporal vertex cover in small-degree graphs. In: Proceedings of the 36th AAI Conference on Artificial Intelligence (AAAI) (2022). To appear
- [38] Mertzios, G.B., Molter, H., Renken, M., Spirakis, P.G., Zschoche, P.: The complexity of transitively orienting temporal graphs. In: Proceedings of the 46th International Symposium on Mathematical Foundations of Computer Science (MFCS), vol. 202, pp. 75–17518
- [39] Buß, S., Molter, H., Niedermeier, R., Rymar, M.: Algorithmic aspects of

- temporal betweenness. In: Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), pp. 2084–2092 (2020)
- [40] Tovey, C.A.: A simplified NP-complete satisfiability problem. *Discrete Applied Mathematics. The Journal of Combinatorial Algorithms, Informatics and Computational Sciences* **8**(1), 85–89 (1984)
- [41] Bang-Jensen, J., Gutin, G.Z.: *Digraphs – Theory, Algorithms and Applications*. Springer, London (2009)
- [42] Bevern, R.v., Mnich, M., Niedermeier, R., Weller, M.: Interval scheduling and colorful independent sets. *Journal of Scheduling* **18**(5), 449–469 (2015)
- [43] Thejaswi, S., Lauri, J., Gionis, A.: Restless reachability problems in temporal graphs. arXiv preprint arXiv:2010.08423 (2020)