

Disentangling the Computational Complexity of Network Untangling

Vincent Froese, Pascal Kunz and Philipp Zschoche

Technische Universität Berlin, Faculty IV, Institute of Software Engineering and Theoretical Computer Science, Algorithmics and Computational Complexity
 {vincent.froese, p.kunz.1, zschoche}@tu-berlin.de

Abstract

We study the recently introduced network untangling problem, a variant of VERTEX COVER on temporal graphs—graphs whose edge set changes over discrete time steps. There are two versions of this problem. The goal is to select at most k time intervals for each vertex such that all time-edges are covered and (depending on the problem variant) either the maximum interval length or the total sum of interval lengths is minimized. This problem has data mining applications in finding activity timelines that explain the interactions of entities in complex networks.

Both variants of the problem are NP-hard. In this paper, we initiate a multivariate complexity analysis involving the following parameters: number of vertices, lifetime of the temporal graph, number of intervals per vertex, and the interval length bound. For both problem versions, we (almost) completely settle the parameterized complexity for all combinations of those four parameters, thereby delineating the border of fixed-parameter tractability.

1 Introduction

The classical VERTEX COVER problem is among the most well-studied NP-hard problems on static graphs especially in the context of parameterized algorithmics. In fact, it is often referred to as the *Drosophila* of parameterized complexity [Downey and Fellows, 2013; Fellows *et al.*, 2018]. Real-world graphs, however, are often dynamic and change over time. In modern applications, this temporal information is readily available and allows for more realistic models. This leads algorithmic research to focus on temporal graphs [Boehmer *et al.*, 2021; Bumpus and Meeks, 2021; Deligkas and Potapov, 2020; Erlebach *et al.*, 2021; Fan *et al.*, 2021; Holme and Saramäki, 2012; Michail and Spirakis, 2016; Singer *et al.*, 2019]. In a temporal graph, edges can appear and disappear over time, which is formalized by a sequence E_1, \dots, E_τ of edge sets over a fixed set V of vertices, where $\tau \geq 1$ denotes the lifetime of the temporal graph. Several classical graph problems have been studied on temporal graphs [Akrida *et al.*, 2020; Casteigts *et al.*, 2021; Klobas

et al., 2021; Mertzios *et al.*, 2020; Mertzios *et al.*, 2021; Zschoche *et al.*, 2020].

In this paper, we study the network untangling problem introduced by Rozenshtein *et al.* [2017; 2021], which is a temporal variant of VERTEX COVER and motivated by data mining applications such as discovering event timelines and summarizing temporal networks. Here, edges in the temporal graph model certain interactions between entities (vertices). The goal is to explain the observed interactions by selecting few (and short) activity intervals for each entity such that at the time of an interaction at least one of the two entities is active. The formal definition is as follows: Let $\mathcal{G} = (V, (E_i)_{i \in [\tau]})$ be a temporal graph. A k -activity timeline is a set \mathcal{T} containing at most k time intervals for each vertex, that is, $\mathcal{T} \subseteq \{(v, a, b) \in V \times [\tau] \times [\tau] \mid a \leq b\}$ such that $|\{(v, a, b) \in \mathcal{T}\}| \leq k$ for each $v \in V$. We say that \mathcal{T} covers \mathcal{G} if for each $t \in [\tau]$ and each $\{u, v\} \in E_t$, \mathcal{T} contains some (u, a, b) with $t \in [a, b]$ or some (v, a, b) with $t \in [a, b]$. That is, for each time step t , the set $\{v \in V \mid (v, a, b) \in \mathcal{T}, t \in [a, b]\}$ must be a vertex cover for the graph $G_t = (V, E_t)$. The task is to find a k -activity timeline that minimizes some objective regarding the interval lengths. Rozenshtein *et al.* [2021] introduced the following two problems.

MINTIMELINE_∞

Input: A temporal graph $\mathcal{G} = (V, E_1, \dots, E_\tau)$ and $k, \ell \in \mathbb{N}_0$.

Question: Is there a k -activity timeline \mathcal{T} covering \mathcal{G} such that $\max_{(v,a,b) \in \mathcal{T}}(b-a) \leq \ell$?

MINTIMELINE₊

Input: A temporal graph $\mathcal{G} = (V, E_1, \dots, E_\tau)$ and $k, \ell \in \mathbb{N}_0$.

Question: Is there a k -activity timeline \mathcal{T} covering \mathcal{G} such that $\sum_{(v,a,b) \in \mathcal{T}}(b-a) \leq \ell$?

An example instance and solutions for each of these two problems are pictured in Figure 1. Note that both problems are equivalent if $\ell = 0$.

Rozenshtein *et al.* [2021] showed that both problems are NP-hard and hard to approximate in polynomial time. They focus on developing heuristics and show in a case study that activity timelines can successfully be recovered for real-world data of Twitter users. In this work, we provide a deeper




Figure 1: An example of a temporal graph: The areas between the dotted lines represent the individual layers of this graph. The intervals highlighted in gray form a solution for MINTIMELINE_+ (with $\ell = 4, k = 2$). The intervals indicated by dashed boxes form a solution for $\text{MINTIMELINE}_\infty$ (with $\ell = 1, k = 2$).

complexity-theoretic understanding of both problems through the lens of parameterized complexity theory. To this end, we study the influence of several natural problem parameters on the algorithmic complexity, namely, the number n of vertices, the lifetime τ , the maximum number k of intervals per vertex, and the interval length bound ℓ . In our multivariate analysis, we identify for both problems which parameter combinations yield fixed-parameter tractability and which combinations are intractable. In doing so, we reveal interesting connections to seemingly unrelated problems like graph coloring and bin packing. Our results yield an (almost) tight characterization and pave the way for a comprehensive picture of the computational complexity landscape of network untangling.

Related Work. As the literature on edge covering is extremely rich, we only consider studies closely related to our setting. For a broader overview on the topic of temporal network mining, we refer to a recent tutorial [Rozenstein and Gionis, 2019]. Our main reference is the work by Rozenstein et al. [2021] who introduced both problems and showed that $\text{MINTIMELINE}_\infty$ is polynomial-time solvable for $k = 1$, whereas MINTIMELINE_+ is NP-hard for $k = 1$. Moreover, they showed that both problems are NP-hard for $\ell = 0$ with unbounded k and thus not approximable in polynomial time. They develop efficient heuristics to solve the problems and provide experiments to evaluate the performance of their approaches.

Akrida et al. [2020] studied a different variant of VERTEX COVER on temporal graphs. Their model expects an edge to be covered at least once over every time window of some given size Δ . That is, they define a temporal vertex cover as a set $S \subseteq V \times [\tau]$ such that, for every time window of size Δ and for each edge $e = \{v, w\}$ appearing in a time step contained in the time window, it holds that $(v, t) \in S$ or $(w, t) \in S$ for some t in the time window with time-edge (e, t) . Among other results, they provide NP-hardness results in very restricted settings.

Fluschnik et al. [2019] and Heeger et al. [2021] studied the parameterized complexity of a “multistage” variant of VERTEX COVER. Here, one is given a temporal graph and seeks a sequence of vertex covers of size at most k (one for each layer of the temporal graph) such that consecutive vertex covers are in some sense similar, e.g., the symmetric difference of two consecutive vertex covers is upper-bounded by some value [Fluschnik et al., 2019] or the sum of symmetric differences of all consecutive vertex covers is upper-bounded by some value [Heeger et al., 2021].

Our Contributions. We almost completely settle the parameterized complexity of $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ for all combinations of the parameters n, τ, k , and ℓ . Table 1 gives an overview of the central results. It turns out that both problems are already NP-hard if the temporal graph contains three identical layers and $k = 2$ and $\ell = 0$ (Corollary 3). This strengthens the hardness result by Rozenstein et al. [2021] for $\ell = 0$ and implies that fixed-parameter tractability can only be achieved in combination with the parameter n . Interestingly, the two problems behave somewhat differently here. Our two main results state that $\text{MINTIMELINE}_\infty$ parameterized by n is W[1]-hard even for $\ell = 1$ (Theorem 11), while MINTIMELINE_+ parameterized by $n + \ell$ is in FPT (Theorem 8). We further show that both problems are in XP when parameterized by n (Theorem 6) and in FPT if $\ell = 0$ (Lemma 9). Theorem 6 shows that MINTIMELINE_+ is also in FPT for $n + k$ and Theorem 7 shows the same for $\text{MINTIMELINE}_\infty$. As regards the case $k = 1$, we strengthen the NP-hardness of MINTIMELINE_+ showing that it already holds for two layers (Theorem 5). Note that $n + \tau$ trivially yields fixed-parameter tractability since the instance size is bounded in these numbers. Thus, for $\text{MINTIMELINE}_\infty$, we obtain a full complexity dichotomy. For MINTIMELINE_+ , the parameterized complexity regarding n remains open. Notably, our hardness results also imply running time lower bounds based on the Exponential Time Hypothesis¹ (ETH) (Corollaries 4 and 14).

Due to space constraints, several proofs (marked by \star) are deferred to the full version of this paper.²

2 Preliminaries

We denote the positive integers by \mathbb{N} and let $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$.

If not specified otherwise, a graph is static and undirected. We refer to Diestel [2016] for details on notation in graph theory. If $G = (V, E)$ is a graph and $S \subseteq V$, let $\text{inc}(S) := \{e \in E \mid e \cap S \neq \emptyset\}$ denote the set of edges incident to at least one vertex in S . A *temporal graph* $\mathcal{G} := (V, (E_i)_{i=1}^\tau)$ consists of a finite vertex set V and a sequence of edge sets $E_1, \dots, E_\tau \subseteq \binom{V}{2}$. 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 *size* of $\mathcal{G} := (V, (E_i)_{i=1}^\tau)$ is $|\mathcal{G}| := |V| + \sum_{t=1}^\tau \max\{1, |E_t|\}$.

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) in $\Sigma^* \times \mathbb{N}_0$, where k denotes the *parameter*. A parameterized problem L is (i) *fixed-parameter tractable* (or contained in the class FPT) if there is an algorithm that decides every instance (x, k) for L in $f(k) \cdot |x|^{O(1)}$ time, (ii) contained in the class XP if there is an algorithm that decides every instance (x, k) for L in $|x|^{f(k)}$ time, and (iii) *para-NP-hard* if L is NP-hard for some constant value of the parameter, where f is any computable function that only depends on the parameter. Note that $\text{FPT} \subseteq \text{XP}$. If a parameterized problem is

¹The Exponential Time Hypothesis [Impagliazzo and Paturi, 2001] states that 3-SAT cannot be solved in $O(2^{cn})$ time for some constant $c > 0$, where n is the number of variables of the input formula.

²On the arXiv: <https://arxiv.org/abs/2204.02668>.

Parameter	$\text{MINTIMELINE}_\infty$	MINTIMELINE_+
$n + k$	$\text{FPT}^{\text{[Thm. 7]}}$	$\text{FPT}^{\text{[Thm. 6]}}$
$n + \ell$	$\text{W}[1]\text{-h.}^{\text{[Thm. 11]}}$, $\text{XP}^{\text{[Thm. 6]}}$	$\text{FPT}^{\text{[Thm. 8]}}$
n	$\text{W}[1]\text{-h.}^{\text{[Thm. 11]}}$ ($\ell = 1$), $\text{XP}^{\text{[Thm. 6]}}$	$\text{XP}^{\text{[Thm. 6]}}$
$\tau = 3, k = 2, \ell = 0$	$\text{NP-h.}^{\text{[Cor. 3]}}$	$\text{NP-h.}^{\text{[Cor. 3]}}$

 Table 1: Overview of results. Parameters: n number of vertices, τ lifetime, k number of intervals per vertex, ℓ interval length bound.

$W[1]$ -hard, then it is presumably not in FPT, and if it is para-NP-hard, then it is not in XP (unless $P = NP$). Further details can be found in [Downey and Fellows, 2013].

3 First Results

This section sets the basis of our parameterized complexity analysis containing results for most of the parameter combinations (except for $n + \ell$).

We start with the general observation (which we will use later) that, for $\ell = 0$, the temporal order of the layers is irrelevant since every element in a k -activity timeline covers only edges of one layer.

Observation 1. Let $\pi: [\tau] \rightarrow [\tau]$ be a permutation. Then, the instance $((V, (E_i)_{i \in [\tau]}), k, 0)$ is a yes-instance of $\text{MINTIMELINE}_\infty$ (MINTIMELINE_+) if and only if $((V, (E_{\pi(i)})_{i \in [\tau]}), k, 0)$ is a yes-instance.

Moreover, if all layers of the temporal graph are equal, then, for $\ell = 0$, both problems are equivalent to an extension of k -coloring. For a graph $G = (V, E)$ and $a, b \in \mathbb{N}$ with $a \geq b$, an $(a:b)$ -coloring of G is a function $c: V \rightarrow \binom{[a]}{b}$ such that $c(u) \cap c(v) = \emptyset$ for all $\{u, v\} \in E$. Note that a $(k:1)$ -coloring of a graph is simply a k -coloring. Given a graph G , the problem $(a:b)$ -COLORING asks whether there is an $(a:b)$ -coloring for G . We get the following.

Theorem 2. If $E_1 = E_2 = \dots = E_\tau$, $\ell = 0$, and $\tau \geq k$, then $\text{MINTIMELINE}_\infty$ (MINTIMELINE_+) is equivalent to $(\tau:\tau - k)$ -COLORING.

Proof. Suppose that \mathcal{T} is a k -activity timeline that covers $\mathcal{G} = (V, E_1, \dots, E_\tau)$ with $E_1 = \dots = E_\tau =: E$ and $\ell = 0$. We can assume that \mathcal{T} contains exactly k intervals for each $v \in V$. Let $G = (V, E)$. Then, $c: V \rightarrow \binom{[\tau]}{\tau - k}$ with $c(v) := \{t \in [\tau] \mid (v, t, t) \notin \mathcal{T}\}$ is a $(\tau:\tau - k)$ -coloring of G .

Conversely, if $c: V \rightarrow \binom{[\tau]}{\tau - k}$ is a $(\tau:\tau - k)$ -coloring of $G = (V, E)$, then $\mathcal{T} := \{(v, t, t) \in V \times [\tau] \times [\tau] \mid v \in V, t \notin c(v)\}$ is a k -activity timeline that covers \mathcal{G} . \square

The NP-hardness of 3-COLORING, or (3:1)-COLORING, implies the following.

Corollary 3. $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ are NP-hard, even if $k = 2$, $\ell = 0$, $\tau = 3$, and all layers are identical.

Bonamy et al. [2019] mention a result by Nederlof [2008] showing that $(a:b)$ -COLORING can be solved in $(b + 1)^n \cdot n^{O(1)}$ time. They show that the ETH implies that there is no $2^{o(\log b)n}$ -time algorithm for $(a:b)$ -COLORING [Bonamy et al., 2019, Theorem 1.1]. Thus, we have the following.

Corollary 4. $\text{MINTIMELINE}_\infty$ (MINTIMELINE_+) with $E_1 = E_2 = \dots = E_\tau$ and $\ell = 0$ cannot be solved in $f(\tau - k) \cdot 2^{o(\log(\tau - k))n}$ time for any function f unless the ETH fails.

Note that, if $\tau = 2$ then $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ are both trivial unless $k = 1$. Rozenstein et al. [2021] showed that $\text{MINTIMELINE}_\infty$ is polynomial-time solvable in the case of $k = 1$. By contrast, we show that MINTIMELINE_+ is NP-hard in the case of $k = 1$ and $\tau = 2$ (strengthening the NP-hardness for $k = 1$ with unbounded τ by Rozenstein et al. [2021, Proposition 3]). Moreover, $\text{MINTIMELINE}_\infty$ is FPT when parameterized by ℓ with this restriction. This gives us the following.

Theorem 5. MINTIMELINE_+ with two layers is NP-hard and FPT when parameterized by ℓ .

Proof. We show NP-hardness with a reduction from ODD CYCLE TRANSVERSAL. In this problem, the input consists of a graph $G = (V, E)$ and an integer s and the task is to decide whether there is a vertex set $X \subseteq V$ with $|X| \leq s$ such that $G - X$ is bipartite. This problem is NP-hard [Lewis and Yannakakis, 1980].

Given an instance $(G = (V, E), s)$, the reduction outputs the instance $(\mathcal{G} := (V, E_1 := E, E_2 := E), k := 1, \ell := s)$. If $X \subseteq V$ with $|X| \leq s$ is an odd cycle transversal in G such that V_1 and V_2 are the two color classes in the bipartite graph $G - X$, then $\mathcal{T} := \{(v, 1, 2) \mid v \in X\} \cup \{(v, i, i) \mid i \in \{1, 2\}, v \in V_i\}$ is a 1-activity timeline that covers \mathcal{G} with $\sum_{(v,a,b) \in \mathcal{T}} (b - a) \leq s$. Conversely, if \mathcal{T} is a 1-activity timeline that covers \mathcal{G} with $\sum_{(v,a,b) \in \mathcal{T}} (b - a) \leq s$, then $X := \{v \in V \mid (v, 1, 2) \in \mathcal{T}\}$ is an odd cycle transversal of size at most s in G .

For fixed-parameter tractability, we show that MINTIMELINE_+ with $\tau = 2$ can be reduced to ALMOST 2-SAT. In this problem, the input consists of a Boolean formula φ in 2-CNF and an integer s and the task is to decide whether φ can be made satisfiable by deleting at most s clauses. The problem is known to be in FPT when parameterized by s [Razgon and O’Sullivan, 2009].

Let $(\mathcal{G} = (V, E_1, E_2), k, \ell)$ be an instance for MINTIMELINE_+ . If $k \neq 1$, the instance is trivial. Otherwise, we introduce two Boolean variables x_1^v and x_2^v for each $v \in V$. Intuitively, x_i^v is true if v is used to cover edges in E_i . The reduction outputs the instance $(\varphi, s := \ell)$, where φ is constructed as follows: For every edge $\{u, v\} \in E_i$, $i \in \{1, 2\}$, φ contains $s + 1$ copies of the clause $(x_i^u \vee x_i^v)$ expressing that the edge must be covered. For every vertex $v \in V$, the formula contains the clause $(\neg x_1^v \vee \neg x_2^v)$, stating that v can only be used once unless this clause is deleted.

Now, suppose that X is a set of at most ℓ clauses of φ and that α is a satisfying assignment for all other clauses. Then, $\mathcal{T} := \{(v, 1, 2) \mid (\neg x_1^v \vee \neg x_2^v) \in X\} \cup \{(v, i, i) \mid \alpha(x_i^v) = \text{true}, (\neg x_1^v \vee \neg x_2^v) \notin X\}$ is a 1-activity timeline that covers \mathcal{G} with $\sum_{(v,a,b) \in \mathcal{T}} (b-a) \leq \ell$. Conversely, if \mathcal{T} is a 1-activity timeline that covers \mathcal{G} with $\sum_{(v,a,b) \in \mathcal{T}} (b-a) \leq \ell$, then removing the at most ℓ clauses in $\{(\neg x_1^v \vee \neg x_2^v) \mid (v, 1, 2) \in \mathcal{T}\}$ from φ makes the formula satisfiable by the following assignment: $\alpha(x_i^v) = \text{true}$ if and only if $(v, i, i) \in \mathcal{T}$ or $(v, 1, 2) \in \mathcal{T}$. \square

Since $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ are NP-hard, even if $\tau + k + \ell$ is constant, the interesting parameterizations remaining involve the number n of vertices. Clearly, combining n and τ trivially yields fixed-parameter tractability for both problems, as the overall instance size is bounded in these two parameters. If we use only the number n of vertices as a parameter, then we can use dynamic programming to show that both problems are in XP.

Theorem 6 (*). (i) $\text{MINTIMELINE}_\infty$ is solvable in $(k+1)^n (\ell+2)^n 2^n \tau n^{O(1)}$ time and (ii) MINTIMELINE_+ is solvable in $(k+1)^n 2^{O(n)} \tau (\ell+1)$ time.

Note that, Theorem 6 (i) also implies that $\text{MINTIMELINE}_\infty$ parameterized by $n+k+\ell$ is in FPT and that Theorem 6 (ii) implies that MINTIMELINE_+ parameterized by $n+k$ is in FPT. To show that $\text{MINTIMELINE}_\infty$ is also in FPT when parameterized by $n+k$, we observe that it is sufficient to consider only two possibilities to cover a time-edge. This leads to a search tree algorithm similar to the one for the classical VERTEX COVER.

Theorem 7. $\text{MINTIMELINE}_\infty$ is solvable in $O(2^{nk} n^{2\tau})$ time.

Proof. Let $(\mathcal{G} = (V, (E_i)_{i \in [\tau]}, k, \ell))$ be a $\text{MINTIMELINE}_\infty$ instance. We solve the instance with a search tree algorithm. To this end, we store counters k_v , $v \in V$, for the number of intervals chosen for each vertex. We initially set all counters to zero and start with an empty solution $\mathcal{T} = \emptyset$.

Let $i = \min\{t \in [\tau] \mid E_t \neq \emptyset\}$ and let $\{u, v\} \in E_i$ be an arbitrary edge (if all layers are empty, then we return “yes”). If $k_u = k$ and $k_v = k$, then we return “no” since we cannot cover this edge. Otherwise, we simply branch into the (at most two) options of either taking (u, i, i') (if $k_u < k$) or (v, i, i') (if $k_v < k$) into \mathcal{T} , where $i' = \min(i + \ell, \tau)$. In each branch, we increase the corresponding counter (k_u or k_v) by one, delete all edges incident to the corresponding vertex (u or v) from the layers $E_i, \dots, E_{i'}$, and recursively proceed on the remaining instance. Clearly, the recursion terminates after at most nk calls. Thus, the running time is in $O(2^{nk} n^{2\tau})$. \square

As a k -activity timeline \mathcal{T} for a temporal graph \mathcal{G} with n vertices is of size at most kn , Theorem 7 and Theorem 6 (ii) imply that $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ are in FPT when parameterized by the solution size $|\mathcal{T}|$.

We continue with disentangling the parameterized complexity of MINTIMELINE_+ and $\text{MINTIMELINE}_\infty$ regarding n and ℓ in Sections 4 and 5.

4 Parameterizing MINTIMELINE_+ by $n + \ell$

The main goal of this section is to show the following.

Theorem 8. MINTIMELINE_+ parameterized by $n + \ell$ is in FPT.

The formal proof is deferred to the end of this section. The general idea behind the algorithm is to split the k -activity timeline \mathcal{T} into two parts $\mathcal{T}_0 \uplus \mathcal{T}_{>0} = \mathcal{T}$. Here, \mathcal{T}_0 contains all elements $(v, a, b) \in \mathcal{T}$ with $a = b$ and $\mathcal{T}_{>0}$ contains all elements $(v, a, b) \in \mathcal{T}$ with $a < b$. Note that $\sum_{(v,a,b) \in \mathcal{T}_0} (b-a) = 0$ and that $\sum_{(v,a,b) \in \mathcal{T}_{>0}} (b-a) \leq \ell$. We observe that $\mathcal{T}_{>0}$ induces an interval graph and that the number of possible induced interval graphs is upper-bounded by a function only depending on $n + \ell$. The algorithm iterates over all of these interval graphs and identifies the part of our temporal graph which can be covered by a corresponding $\mathcal{T}_{>0}$. It remains to find a \mathcal{T}_0 that covers the rest of our temporal graph. To this end, we will show that MINTIMELINE_+ parameterized by n is in FPT if $\ell = 0$. In fact, as both problems are equivalent if $\ell = 0$, we show that $\text{MINTIMELINE}_\infty$ and MINTIMELINE_+ are both in FPT when parameterized by n in this case. In the proof of Theorem 8, we will need a slightly more general version of these problems, in which there is not a single value k , but a value k_v for each $v \in V$. We call this generalization $\text{NONUNIFORM MINTIMELINE}_\infty$ (MINTIMELINE_+).

Lemma 9. $\text{NONUNIFORM MINTIMELINE}_\infty$ (MINTIMELINE_+) parameterized by n is in FPT, if $\ell = 0$.

Proof. Given $\mathcal{I} = (\mathcal{G} = (V, (E_i)_{i \in [\tau]}, (k_v)_{v \in V}, 0))$, we create an integer linear program (ILP) with a number of variables bounded by some function of n . This ILP is feasible if and only if \mathcal{I} is a yes-instance. By Lenstra’s algorithm [1983], this implies that $\text{NONUNIFORM MINTIMELINE}_\infty$ (MINTIMELINE_+) parameterized by n is in FPT, if $\ell = 0$.

We use a variable X_E^S for every $E \subseteq \binom{V}{2}$ and $S \subseteq V$ where S is a vertex cover for (V, E) . Note that the number of these variables is at most $2^{\binom{n}{2} + n}$. Intuitively, the value of the variable X_E^S gives us the number of times the vertex cover S is used to cover a layer with edge set E .

For $E \subseteq \binom{V}{2}$, let $a(E) := |\{t \in [\tau] \mid E_t = E\}|$ denote the number of times the edge set E appears as a layer in \mathcal{G} . Let $\mathcal{C}(E) := \{S \subseteq V \mid S \text{ is a vertex cover of } (V, E)\}$. Then, the ILP constraints are as follows:

$$\sum_{S \in \mathcal{C}(E)} X_E^S = a(E), \text{ for all } E \subseteq \binom{V}{2}, \quad (1)$$

$$\sum_{E \subseteq \binom{V}{2}} \sum_{\substack{S \in \mathcal{C}(E) \\ \text{s.t. } v \in S}} X_E^S \leq k_v, \text{ for all } v \in V, \quad (2)$$

$$X_E^S \in \mathbb{N}, \text{ for all } E \subseteq \binom{V}{2}, S \in \mathcal{C}(E).$$

As we mentioned before, we are only interested in the feasibility of this ILP, so there is no objective function to optimize.

The above ILP is feasible if and only if \mathcal{I} is a yes-instance: Suppose that $(X_E^S)_{E \subseteq \binom{V}{2}, S \in \mathcal{C}(E)}$ is a solution. We define a

solution \mathcal{T} of \mathcal{I} as follows: For any $E \subseteq \binom{V}{2}$, let $\mathcal{C}(E) = \{S_1^E, \dots, S_{\tau}^E\}$. We use S_1^E to cover the first $X_E^{S_1}$ appearances of (V, E) in the layers of \mathcal{G} , that is, we add (v, t, t) to \mathcal{T} for every $v \in S_1^E$ and appearance (V, E_t) . Then we continue with S_2^E to cover the next $X_E^{S_2}$ appearances, and so on. Condition (1) guarantees that \mathcal{T} covers \mathcal{G} and Condition (2) ensures that at most k_v intervals are chosen for each vertex $v \in V$.

Conversely, suppose that \mathcal{I} admits a solution \mathcal{T} . For $t \in [\tau]$, let $X_t := \{v \in V \mid (v, t, t) \in \mathcal{T}\}$. Since $\ell = 0$, X_t must be a vertex cover for (V, E_t) . Then, set $X_E^S := |\{t \in [\tau] \mid E_t = E \text{ and } X_t = S\}|$. It is easy to check that this yields a solution of the ILP. \square

As a side result, Theorem 2 and Lemma 9 together imply the following.

Corollary 10. *(a:b)-COLORING parameterized by the number of vertices is in FPT.*

We are now set to show Theorem 8: MINTIMELINE_+ parameterized by $n + \ell$ is in FPT. The algorithm in the following proof is illustrated in Figure 2.

Proof of Theorem 8. For $(\mathcal{G} = (V, (E_i)_{i \in [\tau]}), k, \ell)$, we say that an ℓ' -pattern is a sequence $(F_0, \dots, F_{\ell'})$ of edge sets where $F_i \subseteq \binom{V}{2}$ for all $i \in [\ell']$. A solution pattern is a pair (P, \mathcal{I}) where P is an ℓ' -pattern with $1 \leq \ell' \leq \ell$ and $\mathcal{I} \subseteq V \times \{0, \dots, \ell'\} \times \{0, \dots, \ell'\}$ with $\mathcal{I} \neq \emptyset$ and $a < b$ for all $(v, a, b) \in \mathcal{I}$. Let \mathcal{P} denote the set of all solution patterns. Note that there are at most $2^{\binom{n}{2}(\ell+1)} \cdot \ell$ different ℓ' -patterns with $1 \leq \ell' \leq \ell$. This yields $|\mathcal{P}| \leq 2^{\binom{n}{2}(\ell+1) + n(\ell+1)^2} \cdot \ell \in 2^{O(n^2\ell^2)}$. The weight of a solution pattern (P, \mathcal{I}) is $w((P, \mathcal{I})) := \sum_{(v, a, b) \in \mathcal{I}} (b - a)$.

Our algorithm first checks all possibilities of how often a solution “matches” each solution pattern (P, \mathcal{I}) , that is, it uses intervals as specified by \mathcal{I} to cover a subsequence of layers of \mathcal{G} that is equal to P (a formal explanation follows below). Since \mathcal{I} contains an interval with positive length, we can choose at most ℓ such solution patterns. Hence, we iterate over all functions $f: \mathcal{P} \rightarrow \{0, \dots, \ell\}$ with $\sum_{x \in \mathcal{P}} f(x)w(x) \leq \ell$ and which have the property that $\sum_{(P, \mathcal{I}) \in \mathcal{P}} f((P, \mathcal{I})) |\{(v, a, b) \mid (v, a, b) \in \mathcal{I}\}| \leq k$ for all $v \in V$. There are at most $(\ell + 1)^{|\mathcal{P}|}$ such functions. Next, we try out every order in which the solution patterns are matched by a solution. The result is a sequence $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$ of $r \leq \ell$ solution patterns. The arguments above imply that we check at most $(\ell + 1)^{2^{O(n^2\ell^2)}} \cdot \ell!$ such sequences.

Next, we check whether $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$ can be matched by a solution and, if it can, compute the temporal graph that remains to be covered once we implement the solution patterns. The above sequence can be matched by a solution if there are $i_1 < j_1 < i_2 < j_2 < \dots < i_r < j_r \in [\tau]$ such that $P_s = (E_{i_s}, \dots, E_{j_s})$ for all $s \in [r]$. Algorithmically, we simply determine i_s and j_s by finding the earliest occurrence of P_s in E_1, \dots, E_τ with $i_s \geq j_{s-1} + 1$.

Next, we compute the remaining temporal graph obtained by implementing $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$ on $i_1, \dots, i_r, j_1, \dots, j_r \in [\tau]$. To this end, for each

$s \in [r]$ and each $t \in [i_s, j_s]$, let $E'_t := E_t \setminus \text{inc}\{v \in V \mid (v, a, b) \in \mathcal{I}_s, i_s + a \leq t \leq i_s + b\}$. We call $\mathcal{G}' := (V, E'_1, \dots, E'_\tau)$ the residual temporal graph. For each $v \in V$, let $k_v := k - \sum_{s=1}^r |\{(v, a, b) \mid (v, a, b) \in \mathcal{I}_s\}|$.

We now build the instance $(\mathcal{G}', (k_v)_{v \in V}, 0)$ of $\text{NONUNIFORM MINTIMELINE}_+$. By Lemma 9, solving this instance is FPT when parameterized by n . Our algorithm returns `true` if $(\mathcal{G}', (k_v)_{v \in V}, 0)$ is a yes-instance for at least one sequence $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$.

It remains to show that the algorithm we described is correct. If the algorithm returns `true`, then this result is clearly correct. Assume that, for $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$, the instance $(\mathcal{G}', (k_v)_{v \in V}, 0)$ has a solution \mathcal{T}_0 . Let $\mathcal{T}_{>0} := \bigcup_{s=1}^r \{(v, i_s + a, i_s + b) \mid (v, a, b) \in \mathcal{I}_s\}$. Then, by construction, it is easy to verify that $\mathcal{T} := \mathcal{T}_0 \cup \mathcal{T}_{>0}$ is a solution for (\mathcal{G}, k, ℓ) .

Conversely, suppose that \mathcal{T} is a solution for (\mathcal{G}, k, ℓ) . Let $\mathcal{T}_{>0} := \{(v, a, b) \in \mathcal{T} \mid a < b\}$. Consider the interval multiset $\mathcal{I} := \{[a, b] \subseteq [\tau] \mid \exists v \in V: (v, a, b) \in \mathcal{T}_{>0}\}$ and the corresponding interval graph $G_{\mathcal{I}} := (\mathcal{I}, E_{\mathcal{I}})$, where $E_{\mathcal{I}} := \{ \{[a, b], [a', b']\} \in \binom{\mathcal{I}}{2} \mid [a, b] \cap [a', b'] \neq \emptyset \}$. Since $\sum_{[a, b] \in \mathcal{I}} (b - a) \leq \ell$, it follows that $G_{\mathcal{I}}$ contains at most ℓ connected components and each of these components covers at most ℓ time steps. We will describe a solution pattern corresponding to each connected component of $G_{\mathcal{I}}$. For a connected component C of $G_{\mathcal{I}}$, let $i := \min_{[a, b] \in C} a$ and $j := \max_{[a, b] \in C} b$. Note that $j - i + 1 \leq \ell$. Consider the $(j - i + 1)$ -pattern $P = (E_i, \dots, E_j)$. Let $\mathcal{J} := \{(v, a - i, b - i) \mid (v, a, b) \in \mathcal{T}_{>0}, [a, b] \in C\}$. Then, the solution pattern corresponding to C is (P, \mathcal{J}) . By listing the solution pattern for each connected component of $G_{\mathcal{I}}$ in the order in which the components appear in the interval graph, we get the solution patterns $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$ where r is the number of components of $G_{\mathcal{I}}$.

We claim that our algorithm returns `true` for this particular sequence of solution patterns. Note that our algorithm may choose different indices $i_1, \dots, i_r, j_1, \dots, j_r$ to implement $(P_1, \mathcal{I}_1), \dots, (P_r, \mathcal{I}_r)$. However, the resulting residual temporal graph \mathcal{G}' is the same as the residual temporal graph derived from the connected components of $G_{\mathcal{I}}$ up to a permutation of the layers. Hence, by Observation 1, \mathcal{G}' is also a yes-instance for $\text{NONUNIFORM MINTIMELINE}_+$ with $\ell = 0$. \square

5 Parameterizing $\text{MINTIMELINE}_\infty$ by $n + \ell$

In this section, we prove that fixed-parameter tractability of $\text{MINTIMELINE}_\infty$ parameterized by $n + \ell$ is (in contrast to MINTIMELINE_+) unlikely.

Theorem 11. *$\text{MINTIMELINE}_\infty$ parameterized by n is $W[1]$ -hard for $\ell = 1$.*

The key difference to the case $\ell = 0$ (which is in FPT (Lemma 9)) is that, for $\ell = 1$, intervals can overlap and hence, the temporal ordering of the layers is relevant. We prove Theorem 11 in three steps by first showing that two generalizations of $\text{MINTIMELINE}_\infty$ with $\ell = 1$ are $W[1]$ -hard when parameterized by n . The key step is to show $W[1]$ -hardness for the following “multicolored” version of $\text{MINTIMELINE}_\infty$.




Figure 2: Illustration of the algorithm for $\ell = 7$ and $k = 2$ in the proof of Theorem 8: There are two solution patterns with the ℓ' -patterns highlighted in light gray and the intervals are highlighted in dark gray. The solution is completed with intervals of length 0, which are represented by diamonds.

MULTICOLORED MINTIMELINE $_{\infty}$

Input: A temporal graph $\mathcal{G} = (V, E_1, \dots, E_{\tau})$ where $V = V_1 \uplus \dots \uplus V_r$, $\ell \in \mathbb{N}_0$, and $k_1, \dots, k_r \in \mathbb{N}$.

Question: Is there an activity timeline \mathcal{T} covering \mathcal{G} such that $\max_{(v,a,b) \in \mathcal{T}} (b-a) \leq \ell$ and $|\{(v, a, b) \in \mathcal{T} \mid v \in V_i\}| \leq k_i$ for all $i \in [r]$?

The proof is by reduction from UNARY BIN PACKING, where we are given m items with sizes $s_1, \dots, s_m \in \mathbb{N}$, a number $\beta \in \mathbb{N}$ of bins, and a bin size $B \in \mathbb{N}$ with all integers encoded in unary. We are asked to decide if there is an assignment $f: [m] \rightarrow [\beta]$ of items to bins such that $\sum_{i \in f^{-1}(b)} s_i \leq B$ for all $b \in [\beta]$. This problem is W[1]-hard when parameterized by β [Jansen *et al.*, 2013]. The general idea behind the reduction is that bins can be represented by vertices and the bin size by the color budgets k_i . The complete proof is deferred to the full version.

Lemma 12 (*). MULTICOLORED MINTIMELINE $_{\infty}$ parameterized by n is W[1]-hard for $\ell = 1$.

Recall the definition of NONUNIFORM MINTIMELINE $_{\infty}$ given in Section 4. Using Lemma 12, we can now show that it is W[1]-hard when parameterized by n for $\ell = 1$. The proof is deferred to the full version.

Lemma 13 (*). NONUNIFORM MINTIMELINE $_{\infty}$ parameterized by n is W[1]-hard for $\ell = 1$.

To prove Theorem 11, we now show how to reduce NONUNIFORM MINTIMELINE $_{\infty}$ to MINTIMELINE $_{\infty}$.

Proof of Theorem 11. We reduce from NONUNIFORM MINTIMELINE $_{\infty}$. Given an instance $(\mathcal{G} = (V = \{v_1, \dots, v_n\}, (E_t)_{t \in [\tau]}), \ell = 1, (k_v)_{v \in V})$, we construct an instance $(\mathcal{G}' = (V' = (V \cup \{u_1, u_2\}), (E'_t)_{t \in [\tau]}), \ell = 1, k)$ of MINTIMELINE $_{\infty}$. We let $V' := V \cup \{u_1, u_2\}$, $\tau' := \tau + 2k(|V| + 2)$, and $k := \max_{v \in V} k_v$. The layers of \mathcal{G}' are as follows: For $t \in [\tau]$, we let $E'_t := E_t$. The layers $E'_{\tau+1}, \dots, E'_{\tau+4k}$ only contain the edge $\{u_1, u_2\}$. Then, for $i \in [n]$, the layers $E'_{\tau+2k(i+1)+1}, \dots, E'_{\tau+2k(i+1)+2(k-k_{v_i})}$ contain only the edge $\{v_i, u_1\}$, while the layers $E'_{\tau+2k(i+1)+2(k-k_{v_i})+1}, \dots, E'_{\tau+2k(i+2)}$ are empty.

Suppose that \mathcal{T} is an activity timeline that covers \mathcal{G} and contains only k_v intervals that use v for each $v \in V$. Then, we construct a k -activity timeline \mathcal{T}' that covers \mathcal{G}' as follows. We include all intervals in \mathcal{T} . We add $(u_1, \tau + 2a - 1, \tau + 2a)$ and $(u_2, \tau + 2k + 2a - 1, \tau + 2k + 2a)$ for all $a \in [k]$ and $(v_i, \tau + 2k(i+1) + a - 1, \tau + 2k(i+1) + a)$ for all $a \in [k - k_{v_i}]$.

Now suppose that \mathcal{T}' is a k -activity timeline that covers \mathcal{G}' . First, \mathcal{T}' must contain k intervals that use u_1 and k inter-

vals that use u_2 in order to cover the appearances of the edge $\{u_1, u_2\}$ in $E'_{\tau+1}, \dots, E'_{\tau+4k}$. Hence, the edges $\{u_1, v_i\}$ can only be covered by intervals that use v_i . This requires $k - k_{v_i}$ intervals that use v_i . Hence, $\mathcal{T} := \{(v, a, b) \in \mathcal{T}' \mid a \in [\tau], v \in V\}$ is an activity timeline that covers \mathcal{G} and only contains k_v vertices that use each $v \in V$. \square

Unless the ETH fails, UNARY BIN PACKING cannot be solved in time $f(\beta)|I|^{o(\beta/\log \beta)}$ for any function f , where $|I|$ is the input size [Jansen *et al.*, 2013, Theorem 3]. That our reduction yields a temporal graph with $O(\beta)$ vertices implies:

Corollary 14. Unless the ETH fails, MINTIMELINE $_{\infty}$ cannot be solved in time $f(n)|\mathcal{G}|^{o(n/\log n)}$, for any function f , even if $\ell = 1$.

6 Conclusion

We completely settled the computational complexity of MINTIMELINE $_{\infty}$ regarding the considered parameters. For MINTIMELINE $_{+}$, the open question remaining is whether it is in FPT when parameterized by the number of vertices. Besides this question, there are many others which could be studied in future work:

- Are there (polynomial) kernelizations for the FPT cases? Developing data reduction rules might be especially interesting from a practical perspective.
- Which other parameters yield tractable special cases? For example, can the number of vertices be replaced by a smaller parameter (e.g. the vertex cover number or the treewidth of the underlying graph)?
- What changes if one only bounds the overall number of activity intervals instead of bounding for each vertex? It should be possible to modify all our algorithms to solve this problem. Maybe some of our hard cases become tractable?
- What about “temporalizing” other vertex selection problems (like DOMINATING SET) in an analogous way (that is, at each time step the set of active vertices must be a valid selection for the current graph)? In fact, our dynamic programs and the ILP should easily work here as well, since they do not specifically depend on vertex covers.

Acknowledgments

This work was initiated at the research retreat of the Algorithmics and Computational Complexity group, TU Berlin, held at Zinnowitz, September 2021. The second author is supported by the DFG Research Training Group 2434 “Facets of Complexity”.

References

- [Akrida *et al.*, 2020] Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis, and Viktor Zamaraev. Temporal vertex cover with a sliding time window. *J. Comput. Syst. Sci.*, 107:108–123, 2020.
- [Boehmer *et al.*, 2021] Niclas Boehmer, Vincent Froese, Julia Henkel, Yvonne Lasars, Rolf Niedermeier, and Malte Renken. Two influence maximization games on graphs made temporal. In *Proc. of the 30th IJCAI*, pages 45–51. ijcai.org, 2021.
- [Bonamy *et al.*, 2019] Marthe Bonamy, Łukasz Kowalik, Michał Pilipczuk, Arkadiusz Socała, and Marcin Wrochna. Tight lower bounds for the complexity of multicoloring. *ACM Trans. Comput. Theory*, 11(3):1–19, 2019.
- [Bumpus and Meeks, 2021] Benjamin Merlin Bumpus and Kitty Meeks. Edge exploration of temporal graphs. In *Proc. of the 32nd IWOCA*, pages 107–121. Springer, 2021.
- [Casteigts *et al.*, 2021] Arnaud Casteigts, Anne-Sophie Himmel, Hendrik Molter, and Philipp Zschoche. Finding temporal paths under waiting time constraints. *Algorithmica*, 83(9):2754–2802, 2021.
- [Deligkas and Potapov, 2020] Argyrios Deligkas and Igor Potapov. Optimizing reachability sets in temporal graphs by delaying. In *Proc. of the 34th AAI*, pages 9810–9817. AAI Press, 2020.
- [Diestel, 2016] Reinhard Diestel. *Graph Theory*. Springer, 5th edition, 2016.
- [Downey and Fellows, 2013] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- [Erlebach *et al.*, 2021] Thomas Erlebach, Michael Hoffmann, and Frank Kammer. On temporal graph exploration. *J. Comput. Syst. Sci.*, 119:1–18, 2021.
- [Fan *et al.*, 2021] Yujie Fan, Mingxuan Ju, Shifu Hou, Yanfang Ye, Wenqiang Wan, Kui Wang, Yinming Mei, and Qi Xiong. Heterogeneous temporal graph transformer: An intelligent system for evolving android malware detection. In *Proc. of the 27th KDD*, pages 2831–2839. ACM, 2021.
- [Fellows *et al.*, 2018] Michael R. Fellows, Lars Jaffke, Aliz Izabella Király, Frances A. Rosamond, and Mathias Weller. What is known about vertex cover kernelization? In *Adventures Between Lower Bounds and Higher Altitudes*, pages 330–356. Springer, 2018.
- [Fluschnik *et al.*, 2019] Till Fluschnik, Rolf Niedermeier, Valentin Rohm, and Philipp Zschoche. Multistage vertex cover. In *Proc. of the 14th IPEC*, pages 14:1–14:14, 2019.
- [Heeger *et al.*, 2021] Klaus Heeger, Anne-Sophie Himmel, Frank Kammer, Rolf Niedermeier, Malte Renken, and Andrej Sajenko. Multistage graph problems on a global budget. *Theor. Comput. Sci.*, 868:46–64, 2021.
- [Holme and Saramäki, 2012] Petter Holme and Jari Saramäki. Temporal networks. *Phys. Rep.*, 519(3):97–125, 2012.
- [Impagliazzo and Paturi, 2001] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [Jansen *et al.*, 2013] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.*, 79(1):39–49, 2013.
- [Klobas *et al.*, 2021] Nina Klobas, George B. Mertzios, Hendrik Molter, Rolf Niedermeier, and Philipp Zschoche. Interference-free walks in time: Temporally disjoint paths. In *Proc. of the 30th IJCAI*, pages 4090–4096. ijcai.org, 2021.
- [Lenstra, 1983] Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.
- [Lewis and Yannakakis, 1980] John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is NP-complete. *J. Comput. Syst. Sci.*, 20(2):219–230, 1980.
- [Mertzios *et al.*, 2020] George B. Mertzios, Hendrik Molter, Rolf Niedermeier, Viktor Zamaraev, and Philipp Zschoche. Computing maximum matchings in temporal graphs. In *Proc. of the 37th STACS*, pages 27:1–27:14, 2020.
- [Mertzios *et al.*, 2021] George B. Mertzios, Hendrik Molter, and Viktor Zamaraev. Sliding window temporal graph coloring. *J. Comput. Syst. Sci.*, 120:97–115, 2021.
- [Michail and Spirakis, 2016] Othon Michail and Paul G. Spirakis. Traveling salesman problems in temporal graphs. *Theor. Comput. Sci.*, 634:1–23, 2016.
- [Nederlof, 2008] Jesper Nederlof. Inclusion exclusion for hard problems, 2008. Master thesis. Department of Information and Computer Science, Utrecht University.
- [Razgon and O’Sullivan, 2009] Igor Razgon and Barry O’Sullivan. Almost 2-SAT is fixed-parameter tractable. *J. Comput. Syst. Sci.*, 75(8):435–450, 2009.
- [Rozenshtein and Gionis, 2019] Polina Rozenshtein and Aristides Gionis. Mining temporal networks. In *Proc. of the 25th KDD*, pages 3225–3226. ACM, 2019.
- [Rozenshtein *et al.*, 2017] Polina Rozenshtein, Nikolaj Tatti, and Aristides Gionis. The network-untangling problem: From interactions to activity timelines. In *Proc. of the ECML/PKDD ’17*, pages 701–716. Springer, 2017.
- [Rozenshtein *et al.*, 2021] Polina Rozenshtein, Nikolaj Tatti, and Aristides Gionis. The network-untangling problem: from interactions to activity timelines. *Data Min. Knowl. Discov.*, 35(1):213–247, 2021.
- [Singer *et al.*, 2019] Uriel Singer, Ido Guy, and Kira Radinsky. Node embedding over temporal graphs. In *Proc. of the 28th IJCAI*, pages 4605–4612. ijcai.org, 2019.
- [Zschoche *et al.*, 2020] Philipp Zschoche, Till Fluschnik, Hendrik Molter, and Rolf Niedermeier. The complexity of finding small separators in temporal graphs. *J. Comput. Syst. Sci.*, 107:72–92, 2020.