# Temporal Reachability Dominating Sets: contagion in temporal graphs 

*Corresponding author(s). E-mail(s): david.c.kutner@durham.ac.uk;<br>laura.larios-jones@glasgow.ac.uk;<br>${ }^{\dagger}$ Authors contributed equally to this work.


#### Abstract

Given a population with dynamic pairwise connections, we ask if the entire population could be (indirectly) infected by a small group of $\boldsymbol{k}$ initially infected individuals. We formalise this problem as the Temporal Reachability Dominating Set (TaRDiS) problem on temporal graphs. We provide positive and negative parameterized complexity results in four different parameters: the number $\boldsymbol{k}$ of initially infected, the lifetime $\boldsymbol{\tau}$ of the graph, the number of locally earliest edges in the graph, and the treewidth of the footprint graph $\mathcal{G}_{\downarrow}$. We additionally introduce and study the MAxMinTARDIS problem, where the aim is to schedule connections between individuals so that at least $\boldsymbol{k}$ individuals must be infected for the entire population to become fully infected. We classify three variants of the problem: Strict, Nonstrict, and Happy. We show these to be coNP-complete, NP-hard, and $\boldsymbol{\Sigma}_{\mathbf{2}}^{P}$-complete, respectively. Interestingly, we obtain hardness of the Nonstrict variant by showing that a natural restriction is exactly the well-studied Distance-3 Independent Set problem on static graphs.


Keywords: Temporal Graphs, Temporal Reachability, Treewidth, Computational Complexity, Polynomial Hierarchy

## 1 Introduction

A natural problem in the study of networks is that of finding a small set of individuals which together can affect the entire network. This problem is practically relevant in identifying influential entities in social networks, in gauging the risk of viral spread in biological networks, or in choosing sources to broadcast from in wireless networks. We study this problem through the lens of temporal graphs. That is, graphs which change over time, capturing the dynamic nature intrinsic to many real-world networks.

We formalize the notion of a set of sources which together reach the whole temporal graph as a Temporal Reachability Dominating Set, or TaRDiS. We study the complexity of the problem of finding a TaRDiS of a given size in a given temporal graph. Later, we ask: if we can choose when connections between individuals exist, can we maximize the size of the minimum TaRDiS? That is, for viral spread, can we guarantee that the entire population will not be contaminated by at most $k$ initial infections. We refer to this problem as MaxMinTaRDiS.

Our problems sit at the intersection of reachability and covering, which are two rich classes of (temporal) graph problems. Reachability is a fundamental concept in temporal graph theory. Temporal reachability and related problems have been the subject of extensive study since the seminal work in the field by Kempe, Kleinberg and Kumar [1]. Several works focus on the complexity of choosing or modifying times to optimize reachability [1-5], while some consider as input temporal graphs in which times are immutably fixed [6-10]. A comparison between our problems and some of those in the literature is shown in Table 1.

MaxMinTaRDiS is closely related to the MinMaxReach problem studied by Enright, Meeks and Skerman [3], in which the objective is to minimize the maximum number of individuals reached by any single vertex. TARDIS, on the other hand, vastly generalizes Casteigts's [11] notion of $\mathcal{J}^{1 \forall}$ connectivity, in which the question is whether any single vertex reaches the entire network. This framing of reachability asks what the worst-case spread is from a single source in the temporal graph. In reality, studied populations are often infected by several individuals. For example, SARS-CoV-2 had been independently introduced to the UK at least 1300 times by June 2020 [12]. This, a dynamic population infected by many sources, is precisely the setting motivating our TARDIS problem.

| Problem | Choose | such that | reach | Aim/motivation |
| :--- | :---: | :---: | :---: | :---: |
| MAXMinTARDIS | times for edges | no coalition of $k$ nodes | all nodes | Minimize connectivity |
| TARDiS | nothing | a coalition of $k$ nodes | all nodes | Assess connectivity |
| MinMAXREACH [3] | times for edges | no single node | more than $k$ nodes | Minimize connectivity |
| $\mathcal{J}^{1 \forall}$ connectedness [11] | nothing | a single node | all nodes | Assess connectivity |
| REACHABILITY <br> InFERENCE [1] | times for edges | a designated root node | all "good" nodes and no <br> "bad" nodes | Network design |
| REACHFAST [13], | edges to delay | each source in the <br> designated set | all nodes | Maximize connectivity |
| MinREACHDELETE [4, 14] | edges to delete | the specified coalition | at most $r$ vertices | Minimize connectivity |
| MinREAChDELAY [4] | edges to delay | the specified coalition | at most $r$ vertices | Minimize connectivity |

Table 1: Comparison of our problems TARDiS and MAxMinTARDiS to problems in the literature.

Our work is focused on the computational (parameterized) complexity of TARDIS and MaxMinTaRDiS, which depends heavily on which formalisation of the problems is considered. In particular, instantaneous spread through a large swath of the population, while realistic in some computer networks, is inconsistent with viral spread in a biological system. Further, should multiple interactions between the same pair of individuals be allowed? Lastly, should it be possible for a single individual to simultaneously interact with several others?

We consider all combinations of answers to these questions, and in all cases for both problems: show that the problem is computationally hard in general; identify the maximum number of discrete times at which edges can appear such that the problem remains tractable; and provide (parameterized) algorithms. Interestingly, we also show through a nontrivial proof that MaxMinTARDIS generalizes the well-studied Distance-3 Independent Set problem.

### 1.1 Problem Setting

We begin with some standard definitions. We denote by $[i, j]$ the set $\{i, i+1, \ldots, j\}$ and $[j]$ the set $[1, j]$. Let $G=(V, E)$ be an undirected graph with $(u, v) \in E$. We say that $u$ and $v$ are adjacent (also neighbours) and that the edge $(u, v)$ is incident to both $u$ and $v$. If $S \subseteq V$ is a set of vertices, we say an edge $(u, v)$ is incident to $S$ if one of its endpoints $u$ or $v$ is in $S$. For any vertex $v \in V$, the closed neighbourhood of $v$ is denoted $N[V]:=\{v\} \cup\{u: u$ is adjacent to $v\}$. We say $G$ is planar if it can be drawn in a plane without any edges intersecting.

A temporal graph $\mathcal{G}=(V, E, \lambda)$ consists of a set of vertices $V$, a set of edges $E$ and a function $\lambda: E \rightarrow$ $2^{\tau} \backslash\{\emptyset\}$ which maps each edge to a discrete set of times where the lifetime $\tau \in \mathbb{N}$ of a temporal graph is the value of the latest timestep. We refer to $\lambda$ as the temporal assignment of $\mathcal{G}$. If $t \in \lambda(e)$ then we call the pair $(e, t)$ a time-edge, and say $e$ is active at time $t$. The set of all time-edges is denoted $\mathcal{E}$. We abuse notation and write $\lambda(u, v)$ to mean $\lambda((u, v))$.

For a static graph $G=(V, E)$, we denote the temporal graph $(V, E, \lambda)$ by $(G, \lambda)$. We also use $V(\mathcal{G}), E(\mathcal{G})$ to refer to the vertex and edge sets of $\mathcal{G}$, respectively, and use $E_{t}(\mathcal{G})$ to refer to the set of edges active at time $t$, and call $G_{t}(\mathcal{G})=\left(V(\mathcal{G}), E_{t}(\mathcal{G})\right)$ the snapshot at time $t$. When $\mathcal{G}$ is clear from the context we may omit it. Also, we use the convention that no snapshot is empty, which guarantees $\tau \leq|\mathcal{E}|$. The static graph $\mathcal{G}_{\downarrow}=(V, E)$ is referred to as the footprint of $\mathcal{G}$.

A strict (respectively nonstrict) temporal path from a vertex $u$ to a vertex $v$ is a sequence of time-edges $\left(e_{1}, t_{1}\right), \ldots,\left(e_{l}, t_{l}\right)$ such that $e_{1} \ldots e_{l}$ is a static path from $u$ to $v$ and $t_{i}<t_{i+1}\left(\right.$ resp. $\left.t_{i} \leq t_{i+1}\right)$ for $i \in[1, l-1]$. A vertex $u$ temporally reaches (or just "reaches") a vertex $v$ if there is a temporal path from $u$ to $v$. The


Fig. 1: Spread in a temporal graph from source $s$ through snapshots. Vertices are shaded (half-shaded) when reached from $s$ by a strict (nonstrict) temporal path.
reachability set $R_{u}(\mathcal{G})$ of a vertex $u$ is the set of vertices reachable from $u$. When the graph is clear from context, we may simply use $R_{u}$ to refer to the reachability set of a vertex $u$. We say a vertex $u$ is reachable from a set $S$ if for some $v \in S, u \in R_{v}(\mathcal{G})$. A set of vertices $T$ is temporal reachability dominated by another set of vertices $S$ if every vertex in $T$ is reachable from $S$. Domination of and by single vertices is analogously defined. Strict and nonstrict reachability are illustrated in Figure 1. We differentiate between strict and nonstrict reachability by introducing a superscript $<$ or $\leq$ to the appropriate operators. For example, in Figure $1, u$ is in $R_{\bar{s}}^{\leq}$, but is not in $R_{s}^{<}$. Note that any strict temporal path is also a nonstrict temporal path, but the converse does not necessarily hold.

Casteigts, Corsini, and Sarkar [10] define three useful properties a temporal graph may exhibit. A temporal graph is called simple if each edge is active exactly once, proper if each snapshot has maximum degree one, and happy if it is both simple and proper. Figure 2 provides examples of the different types of temporal graph. For simple graphs, we define the temporal assignment as $\lambda: E \rightarrow[\tau]$ for convenience. We also use these three terms to describe the temporal assignment of a graph with the corresponding property. Happy temporal graphs have the property that any nonstrict temporal path is also a strict temporal path. Part of the utility of happy temporal graphs is that hardness results on them generalize to the strict and nonstrict settings. We can now introduce our protagonist.


Fig. 2: Four small examples of (non)simple and (non)proper temporal graphs.

Definition 1 (TaRDiS). In a temporal graph $\mathcal{G}$, a (strict/nonstrict) Temporal Reachability Dominating Set (TaRDiS) is a set of vertices $S$ such that every vertex $v \in V(\mathcal{G})$ is temporally reachable from a vertex in $S$ by a (strict/nonstrict) temporal path.

A minimum TaRDiS is a TaRDiS of fewest vertices in $\mathcal{G}$. We emphasize that, just as every strict temporal path is a nonstrict temporal path, every strict TaRDiS is a nonstrict TaRDiS. It is possible that the smallest nonstrict TaRDiS is strictly smaller than the smallest strict TaRDiS; for example, in Figure 2b, there is a nonstrict TaRDiS of size 1 (namely $\{u\}$ ) but every strict TaRDiS has size at least 2 . We now formally define our problems.

## (Strict/Nonstrict) TaRDiS

Input: A temporal graph $\mathcal{G}=(V, E, \lambda)$ and an integer $k$.
Question: Does $\mathcal{G}$ admit a (strict/nonstrict) TaRDiS of size at most $k$ ?
The restriction to happy inputs $\mathcal{G}$ is a subproblem of both Strict TaRDiS and Nonstrict TaRDiS.

## Happy TaRDiS

Input: A happy temporal graph $\mathcal{G}=(V, E, \lambda)$ and an integer $k$.
Question: Does $\mathcal{G}$ admit a TaRDiS of size at most $k$ ?
MAXMinTARDiS is an extension of our problem in which we look to find a temporal assignment such that no TaRDiS of cardinality less than $k$ exists. As seen in Figure 3, scheduling social events is natural combinatorial problem. In a similar manner to how Edge Colouring corresponds straightforwardly to scheduling meetings between one pair of people at a time to avoid a scheduling conflict, the MaxMinTaRDiS problem can be thought of as scheduling interactions (potentially simultaneously) so that the risk of widespread contagion is limited.


Fig. 3: xkcd $\# 2450$ [15] depicts the scheduling of social events to minimize the risk of contagion.

```
(Strict/Nonstrict) MaxMinTaRDiS
Input: A static graph \(H=(V, E)\) and integers \(k\) and \(\tau\).
Question: Does there exists a temporal assignment \(\lambda: E \rightarrow 2^{\tau} \backslash\{\emptyset\}\) such that every strict/nonstrict TaRDiS admitted by \((H, \lambda)\) is of size at least \(k\) ?
```

Likewise, the variant of this problem in which the temporal assignment $\lambda$ is required to be happy is referred to as Happy MaxMinTaRDiS. Note this is not a subproblem of Strict MaxMinTaRDiS or Nonstrict MaxMinTaRDis. We will later show that Happy TaRDiS generalizes Edge Colouring (and hence conflict-free scheduling in our earlier analogy).

## Happy MaxMinTaRDiS

Input: A static graph $H=(V, E)$ and integers $k$ and $\tau$.
Question: Does there exists a happy temporal assignment $\lambda: E \rightarrow[\tau]$ such that every TaRDiS admitted by $(H, \lambda)$ is of size at least $k$ ?

In this work, by "each variant" we refer to the Strict, Nonstrict and Happy variants of the problems.

### 1.2 Our Contribution

Our work highlights the complexity intrinsic to the dynamic behavior of spreading processes as soon as timevarying elements are incorporated into natural models. At a high level, we identify the minimum lifetime $\tau$ for which each problem is computationally hard. The existence of hardness results even with bounded lifetime justifies the need for parameters other than $\tau$ to obtain tractability results. We provide a fixed-parameter tractable (fpt) ${ }^{1}$ algorithm for TARDIS with parameters $\tau$ and the treewidth ${ }^{2}$ of the footprint graph $\operatorname{tw}\left(\mathcal{G}_{\downarrow}\right)$, and show existence of such an fpt algorithm for MaxMinTARDiS with parameters $\tau, \operatorname{tw}\left(\mathcal{G}_{\downarrow}\right)$, and $k$.

Our results relating lifetime and computational complexity are highlighted in Table 2, and our parameterized complexity results are summarized in Table 3. For the case of happy temporal graphs, we exactly characterize the complexity of both TARDIS and MaxMinTaRDIS with lifetime $\tau \leq 3$. Both problems

[^0]| $\tau$ | TARDIS |  |  | MaxMinTaRDiS |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Strict | Nonstrict | Happy | Strict | Nonstrict | Happy |
| 1 | NP-c <br> (Cor. 2) | Linear <br> (Lem. 4) | Linear <br> (Lem. 4) | $\begin{aligned} & \text { coNP-c } \\ & (\text { Cor. 6) } \end{aligned}$ | Linear <br> (Lem. 18) | Linear <br> (Lem. 18) |
| 2 |  | NP-c <br> (Thm. 2) |  |  | NP-c <br> (Cor. 7) |  |
| 3 |  |  | NP-c |  | $\in \Sigma_{2}^{P}$ | $\begin{aligned} & \Sigma_{2}^{P}-c \\ & (\text { Thm. } 4) \end{aligned}$ |
| $\geq 4$ |  |  | (Thm. 1) |  | (Lem. 15) | $\begin{aligned} & \hline \in \text { coNP-h } \cap \Sigma_{2}^{P} \\ & (\text { Cor. } 5, \text { Lem. } 15) \end{aligned}$ |

Table 2: Computational complexity of our problems and its dependence on $\tau$.

|  | TARDIS |  |  | MaxMinTARDIS |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Parameter | Strict | Nonstrict | Happy | Strict | Nonstrict | Happy |
| $\Delta+\tau$ | para-NP-h <br> (Cor. 2) | para-NP-h <br> (Thm. 2) | para-NP-h <br> (Thm. 1) | para-NP-h <br> (Cor. 6) | para-NP-h <br> (Cor. 7) | para-NP-h <br> (Thm. 4) |
| $\#$ LEE | n/a | FPT | (Lem. 32) |  | n/a |  |

Table 3: Summary of our parameterized complexity results. Parameters are: maximum degree of the footprint graph $\Delta$, lifetime $\tau$, number of (weakly) locally earliest edges \#LEE, input $k$, and treewidth of the footprint graph $\operatorname{tw} \mathcal{G}_{\downarrow}$. Problems which are $\mathrm{W}[1]-$, (co-)W[2]-, or para-NP-hard are ones for which there presumably exists no fpt algorithm.
are trivially solvable in linear time for $\tau \leq 2$. We show NP-completeness of HAPPY TARDIS and $\Sigma_{2}^{P-}$ completeness of Happy MaxMinTaRDiS when $\tau=3$ - even when restricted to planar inputs of bounded degree.

For MaxMinTaRDiS, membership of NP is nontrivial since the existence of a polynomial-time verifiable certificate is uncertain. Indeed, the $\Sigma_{2}^{P}$-completeness of Happy MaxMinTaRDiS indicates no such certificate exists in general unless the Polynomial Hierarchy collapses. Interestingly, we show equivalence ${ }^{3}$ of Nonstrict MaxMinTaRDiS restricted to inputs where $\tau=2$ and Distance-3 Independent Set, which is NP-complete [16], in Section 3.3. Given the uncertain membership of NP for the general problem, NP-completeness of this restriction of Nonstrict MaxMinTaRDiS indicates it is possibly easier than the unrestricted problem.

Having shown $\tau$ and planarity of the footprint graph alone are insufficient for tractability, we give an algorithm which solves TARDIS on trees in time polynomial in the number of time-edges $|\mathcal{E}|$ in the input graph. In addition, we give an algorithm for TARDIS on nice tree decompositions [17]. This gives us tractability with respect to lifetime and treewidth of the footprint of the input graph combined. We also show existence of an algorithm for MaxMinTaRDiS which is fixed-parameter tractable with respect to $\tau$, $k$, and treewidth. This is achieved by applying Courcelle's theorem [18].

### 1.3 Related Work

Reachability and connectivity problems on temporal graphs have drawn significant interest in recent years. These have been studied in the context of network design [6, 19, 20] and transport logistics [21] (where

[^1]maximizing connectivity and reachability at minimum cost is desired), and the study of epidemics $[3,14$, $22,23]$ and malware spread [24](where it is not).

In research on networks, broadcasting refers to transmission to every device. In a typical model, communication rather than computation is at a premium, and there is a single source in a graph which does not vary with time[25]. Broadcasting-based questions deviating from this standard have been studied as well. Namely, there is extensive study of the complexity of computing optimal broadcasting schedules for one or several sources [26, 27], broadcasting in ad-hoc networks or time-varying graphs [28], and the choice of multiple sources (originators) for broadcasting in minimum time in a static graph [29, 30]. Ours is the first work to focus on the hardness of choosing multiple sources in a temporal graph to minimize the number of sources, in an offline setting.

One metric closely related to a temporal graph's vulnerability to contagion is its maximum reachability; that is, the largest number $k$ such that some vertex reaches $k$ vertices in the graph. In Enright, Meeks and Skerman [3] and Enright, Meeks, Mertzios and Zamaraev's [14] works, the problems of deleting and reordering edges in order to minimize $k$ are shown to be NP-complete. Problems on temporal graphs are often more computationally complex than their static counterparts [31-34]. In such cases, efficient algorithms may still be obtained on restricted inputs. In work exploring disease spread through cattle, Enright and Meeks find that these real-world networks naturally have low treewidth [35]. Treewidth and other structural parameters have been used with varying degrees of success to give tractability on some of these temporal problems [31, 33, 36, 37].

A powerful tool in parameterized complexity is Courcelle's theorem, which gives tractability on graphs of bounded treewidth for problems that can be represented in monadic second order logic [38]. Unfortunately, there are many temporal problems which remain intractable even when the underlying graph is very strongly restricted, for example when it is a path, a star, or a tree [31-34], all of which have treewidth 1 . When this is the case, it is sometimes sufficient to additionally bound the lifetime of the temporal graph in addition to its underlying structure. This motivates our use of treewidth in combination with lifetime as parameters for the study of our problems.

The problems we study generalize two classical combinatorial graph problems, namely Dominating Set and Distance-3 Independent Set (D3IS). For a static graph $G=(V, E)$, a dominating set is a set of vertices $S$ such that every vertex is either adjacent to a vertex in $S$ or is in $S$ itself, and a D3IS is a set of vertices $S$ such that no pair of vertices $u, v \in S$ has a common neighbour $w$ (so all pairs are at distance at least three). The corresponding decision problems ask, for an integer $k$, if there exists a dominating set (respectively D3IS) of size at most (resp. at least) $k$, and are W[2]- (resp. W[1]-) complete [16, 39]; that is, even if $k$ is fixed it is unlikely there exists an algorithm solving the problem with running time $f(k) \cdot n^{O(1)}$. However, Dominating Set can be solved in polynomial time on graphs of bounded treewidth [17, 40].

TARDiS is exactly the problem of solving the directed variant of Dominating SEt on a Reachability Graph [41]. A Reachability Graph is also referred to as the transitive closure of journeys in a temporal graph and is shown to be efficiently computable by Whitbeck, Amorim, Conan, and Guillaume [9]. Temporal versions of dominating set and other classical covering problems have been well studied [11, 19, 42], however these interpretations do not allow a chosen vertex to dominate beyond its neighbours. Furthermore, many other problems looking to optimally assign times to the edges of a static graph have been studied $[1,3,5]$. TaRDiS also generalises Temporal Source, which asks whether a single vertex can infect every other vertex in the graph, which is equivalent to the graph being a member of the class $\mathcal{J}^{1 \forall}$ [11] mentioned earlier. There has also been extensive research into modifying an input temporal graph (subject to some constraints) to achieve some desired reachability objective [4, 11, 13]. Deligkas, Eiben, and Skretas [13] have the objective of modifying $\lambda$ through a delaying operation so that each vertex in a designated set reaches all vertices in the graph; that is $\bigcap_{u \in S} R_{u}=V(\mathcal{G})$. Molter, Renken and Zschoche [4] consider both delay and deletion operations to modify $\lambda$ with the objective of minimizing the cardinality of the set $\bigcup_{u \in S} R_{u}$. Contrast these with our problem MaxMinTaRDiS, which has the objective of maximizing the minimum cardinality of any set $S$ satisfying $\bigcup_{u \in S} R_{u}=V(\mathcal{G})$.

### 1.4 Organization

This paper is organised as follows. We begin with classical complexity results for TaRDIS in Section 2. This consists of some preliminary observations pertaining to our problems in Section 2.1, which also convey some of the intuition for TARDIS and provide tools for later technical results. Here we also consider temporal graphs with very small lifetime, showing that Nonstrict TARDiS and Happy TARDiS are efficiently solvable in temporal graphs with lifetime $\tau=1$ and $\tau=2$ respectively. In Sections 2.2 and 2.3 we show that these are the largest lifetime for which the problem is efficiently solvable in general. Finally, in Section 2.4, we show that TARDIS is efficiently solvable when the footprint of the temporal graph is a tree.

Section 3 presents classical complexity results for our MaxMinTaRDiS problems. We again begin with some preliminary results in Section 3.1 which provide the scaffolding for our later proofs as well as some easy bounds against which we compare our more technical results. This section allows us to draw the comparison between our problems and the classical problems Edge Colouring and Dominating Set. The remainder of this section includes our proofs of the contrasting results that Nonstrict MaxMinTaRDiS generalizes D3IS (and is in NP when restricted to $\tau=2$ ), and of the $\Sigma_{2}^{P}$-completeness of Happy MaxMinTaRDiS.

Having established the hardness of both problems in the general case, and provided an algorithm for TARDIS when the footprint graph is a tree, we turn to parameterized complexity in Sections 4 and 5. The majority of these sections focus on the structural parameter of treewidth of the footprint graph. We give an algorithm that solves TARDIS on a nice tree decomposition of a graph in Section 4.3. In Section 5.1 we show existence of such an algorithm for MaxMinTaRDiS by leveraging Courcelle's theorem. Finally, we give some concluding remarks and future research directions in Section 6.

## 2 Classical complexity results for TARDIS

In this section, we establish NP-completeness of each variant of TARDIS, and characterize the maximum lifetime $\tau$ for which the problem tractable. We also give an algorithm solving the problem in polynomial time when the footprint graph $\mathcal{G}_{\downarrow}$ is a tree. A more general algorithm for footprint graphs of bounded treewidth is later given in Section 4.

### 2.1 Containment in NP, useful tools, and small lifetime

We begin this section by showing containment of Strict, Nonstrict and Happy TaRDiS in NP and introducing the notions of a (weakly) locally earliest edge and (weakly) canonical TaRDiS. We then show that there always exists a canonical minimum TaRDiS in happy temporal graphs. This allows us to reduce the number of cases we must consider when solving Happy MaxMinTaRDiS in Section 3.2. Finally, we consider the restrictions of the lifetime to $\tau=1$ and $\tau \leq 2$, where Nonstrict TaRDiS and Happy TaRDiS respectively are easily shown to be efficiently solvable. We later show that these are the largest values of $\tau$ for which these problems are tractable.
Lemma 1. Each variant of TARDIS is in NP.
Proof. The reachability set of a vertex can be computed in polynomial time by a modification of breadthfirst search. Therefore, we can verify whether a set temporal reachability dominates a temporal graph in polynomial time.

We now introduce the notion of a canonical TaRDiS. The following results will then allow us to make assumptions about the composition of a minimum TaRDiS when working with proper or happy temporal graphs, or on Nonstrict TaRDiS.
Definition 2 ((Weakly) locally earliest, (weakly) canonical TaRDiS). In a temporal graph $\mathcal{G}$, a time-edge $((u, v), t)$ is locally earliest if every other time-edge incident to either $u$ or $v$ is at a time $t^{\prime}>t$. If the weaker constraint $t^{\prime} \geq t$ holds, then we call the time-edge weakly locally earliest. We say an edge ( $u, v$ ) is (weakly) locally earliest if, for some $t$, the time-edge $((u, v), t)$ is (weakly) locally earliest. A (weakly) canonical TaRDiS consists exclusively of vertices which are incident to a (weakly) locally earliest edge.

In Figure 4 b , each of $\{b, d\},\{b, c, d\}$ and $\{a, c, e\}$ is a TaRDiS, but only the former two are canonical.

(a) $R_{a}^{\leq} \neq R_{a}^{<}$and $R_{e}^{\leq} \neq R_{e}^{<}$.

(b) $\{a, e\}$ is not canonical.

(c) Every TaRDiS is canonical.

Fig. 4: Three temporal graphs, all admitting $\{a, e\}$ as a minimum TaRDiS.

Lemma 2. In a temporal graph, there always exists a minimum weakly canonical nonstrict TaRDiS. In a proper temporal graph, there always exists a canonical TaRDiS.

Proof. We may assume without loss of generality that $\mathcal{G}_{\downarrow}$ is a connected graph. Note that a weakly locally earliest edge in a proper temporal graph is necessarily a locally earliest edge. Consequently, for the remainder of the proof, we focus on weakly locally earliest edges.

Given a TaRDiS $S$, we can find a weakly canonical TaRDiS $S^{\prime}$ by replacing vertices not incident to weakly locally earliest edges by vertices that are. To see this, suppose $x \in S$ is not incident to a weakly locally earliest edge. Then there exist vertices $u, v \in V(\mathcal{G})$ and time-edges $((x, u), t),\left((u, v), t^{\prime}\right)$ such that $t \geq t^{\prime}$. Suppose that $t$ is the earliest time that this is true for. Then, since all vertices reachable from $x$ must be temporally reachable from $v$ by a path appending the time-edges $\left((v, u), t^{\prime}\right),((u, x), t), R_{x} \subseteq R_{v}$. Therefore, we can swap $x$ for $v$ without losing domination of the whole graph. We perform this replacement repeatedly for every such $x$ until we obtain a set of vertices each incident to at least one weakly locally earliest edge. This is precisely a weakly canonical TaRDiS.

This gives us a result that mirrors Lemma 54 from [43].
Corollary 1. The number of weakly locally earliest edges upper-bounds the size of a minimum nonstrict TaRDiS.

We now show that the classical problem Dominating Set is a special case of Strict TaRDiS. Dominating Set is known to be NP-complete, even on planar graphs with maximum degree $\Delta=3$ (hence para-NP-hard with respect to $\Delta$ ) [44], and W[2]-hard with respect to the size of the dominating set [39].
Lemma 3. For all positive integers $\tau$ and instances $(G, k)$ of Dominating Set, a temporal graph $\mathcal{G}$ of lifetime $\tau$ can be found in linear time, where $(\mathcal{G}, k)$ is a yes-instance of Strict TARDIS if and only if $(G, k)$ is a yes-instance of Dominating Set.
Proof. The construction of $\mathcal{G}$ can be seen in Figure 5. We append a path $P$ of length $\tau-1$ to an arbitrary vertex $v$ in $G$. The temporal assignment of $\mathcal{G}$ is defined as follows: all edges in $E(G)$ are assigned time 1 , and the edges in $P$ are assigned times in ascending order from $v$. It is clear that there exists a dominating set of cardinality $k$ in $G$ if and only if there exists a strict TaRDiS of cardinality $k$ in $\mathcal{G}$.

We note that, in our construction, $k$ is unchanged and the maximum degree of $\mathcal{G}_{\downarrow}$ is bounded by $\Delta(G)+1$; giving us the following corollary.
Corollary 2. Strict TaRDIS is NP-complete, W[2]-hard with respect to $k$, and para-NP-hard with respect to $\Delta+\tau$, where $\Delta$ is the maximum degree of the footprint graph and $\tau$ is the lifetime of the temporal graph.


Fig. 5: Illustration of the construction used to reduce from Dominating Set to Strict TaRDiS with arbitrary lifetime $\tau$.

We now consider the problems Nonstrict TaRDiS and Happy TaRDiS with very restricted lifetimes.
Lemma 4. Nonstrict TaRDiS can be computed in linear time when $\tau=1$. When $\tau \leq 2$, Happy TaRDiS is solvable in linear time.

Proof. If every edge is active at the same time, every vertex can be temporally reached by a nonstrict path from any vertex in the same connected component. Therefore, if $\tau=1$, we find a minimum TaRDiS by choosing exactly one vertex in every connected component of the graph.

If a temporal graph is happy and the lifetime of the graph is at most 2 , the footprint of the graph must be 2-edge colourable, and so consist of only paths or even cycles. Any TaRDiS must include every vertex which
has no neighbours. Then, for any paths, we choose a leaf and add its neighbour to the TaRDiS. Following this, we work along the path adding an endpoint of an edge if it is active before the edge previously considered. We use a similar method on even cycles by simply picking an arbitrary vertex $v$ incident to an edge in $E_{1}(\mathcal{G})$ to be in the TaRDiS. From here, we find the subgraph induced by removing the vertices reachable from $v$, giving us a path to which we can apply the previous procedure.

### 2.2 NP-completeness of Happy TaRDiS with lifetime 3

Above, we establish that Happy TARDiS can trivially be solved in linear time when the input has lifetime $\tau \leq 2$ by Lemma 4. Here we show that the problem immediately becomes NP-complete for inputs where $\tau=3$, even when $\mathcal{G}_{\downarrow}$ is planar.

We give a reduction from the problem Planar exactly 3-bounded 3-SAT, which asks whether the input Boolean formula $\phi$ admits a satisfying assignment. We are guaranteed that $\phi$ is a planar formula in $3-$ CNF (a disjunction of clauses each containing at most 3 literals) with each variable appearing exactly thrice and each literal at most twice. This problem is shown to be NP-hard by Tippenhauer and Muzler [45].

### 2.2.1 Construction

Let $\phi$ be an instance Planar exactly 3-bounded 3-SAT consisting of clauses $\left\{c_{1}, \ldots, c_{m}\right\}$ over variables $X=\left\{x_{1}, \ldots, x_{n}\right\}$. Let $l_{2 i}$ (resp. $l_{2 i-1}$ ) denote the positive (resp. negative) literal for variable $x_{i}$, and $l_{2 n+1}$ denote the special literal $\perp$, which is always False. We rewrite every 2-clause of $\phi$ to include the special literal $\perp$ (e.g. the clause $\left(x_{i} \vee \neg x_{j}\right)$ becomes $\left(x_{i} \vee \neg x_{j} \vee \perp\right)$ then $\left(l_{2 i} \vee l_{2 j-1} \vee l_{2 n+1}\right)$ ). Let $\phi^{\prime}$ denote the new formula obtained by doing this. Note that $\phi^{\prime}$ admits a satisfying assignment (in which $\perp$ evaluates to False) if and only if $\phi$ admits a satisfying assignment.


Fig. 6: Gadgets and adjacent literal vertices in the reduction from Planar exactly 3-bounded 3-SAT to HAPPY TARDiS. Left: the variable gadget for $x_{i}$, which appears twice negatively. Right: the clause gadget for the clause $c_{j}$.

In this construction $E(\mathcal{G})=E_{1} \cup E_{2} \cup E_{3} ; \lambda$ is implicitly defined on this basis. We introduce literal vertices as follows. Iterate over $\phi^{\prime}$ : at the $a$ th appearance of each literal $l_{i}$, create two new vertices $l_{i}^{a}$ and $\overline{l_{i}^{a}}$, and let $\left(l_{i}^{a}, \overline{l_{i}^{a}}\right) \in E_{3}$. We refer to the set of all literal vertices as $L$. We say the vertex $l_{i}^{a}$ belongs to the clause $c_{j}$ in which $l_{i}$ appears for the $a$ th time.

For each variable $x_{i}$, we introduce vertices $V_{i}=\left\{T_{i}^{1}, T_{i}^{2}, F_{i}^{1}, F_{i}^{2}, v_{i}^{1}, v_{i}^{2}, a_{i}, b_{i}\right\}$, connected in a cycle labeled as shown in Figure 6. Specifically, we have $\left(T_{i}^{1}, T_{i}^{2}\right),\left(F_{i}^{1}, F_{i}^{2}\right),\left(v_{i}^{1}, v_{i}^{2}\right) \in E_{1},\left(a_{i}, v_{i}^{1}\right),\left(b_{i}, v_{i}^{2}\right) \in E_{2}$, and $\left(a_{i}, T_{i}^{1}\right),\left(b_{i}, F_{i}^{1}\right) \in E_{3}$. Then, for each positive (resp. negative) literal vertex $\overline{l_{2 i}^{a}}$ (resp. $\overline{l_{2 i-1}^{a}}$ ), we add $\left(\overline{l_{2 i}^{a}}, T_{i}^{a}\right)$ to $E_{2}$.

Every three literals vertices $u, v, w$ belonging to the same clause $c_{j}$ are connected with the clause gadget as shown in Figure 6. Namely, we introduce the vertices $Q_{j}=\left\{q_{j}^{1}, \ldots, q_{j}^{6}\right\}$, with $\left(q_{j}^{1}, q_{j}^{2}\right),\left(q_{j}^{3}, q_{j}^{4}\right),\left(q_{j}^{5}, q_{j}^{6}\right) \in E_{3}$, $\left(q_{j}^{2}, q_{j}^{3}\right),\left(q_{j}^{4}, q_{j}^{5}\right),\left(q_{j}^{6}, q_{j}^{1}\right) \in E_{1}$, and $\left(q_{j}^{1}, u\right),\left(q_{j}^{3}, v\right),\left(q_{j}^{5}, w\right) \in E_{2}$. If, for example, clause $c_{j}=\left(l_{5} \vee l_{17} \vee l_{20}\right)$ is the first appearance of $l_{5}$ and $l_{20}$ and the second appearance of $l_{17}$ in $\Phi$, then $u=l_{5}^{1}, v=l_{17}^{2}$ and $w=l_{20}^{1}$.

Lastly, we let $k=2 m+2 n$. This concludes our construction of the Happy TaRDiS instance $(\mathcal{G}, k)$.

### 2.2.2 Properties of the construction

We will show that $\mathcal{G}$ admits a TaRDiS $S$ of size $k$ if $\phi$ is satisfiable, and that any TaRDiS is of size at least $k+1$ otherwise. It is worth noting that our construction preserves planarity of $\phi$. That is, $\mathcal{G}_{\downarrow}$ is a planar graph.

We now show that the existence of a TaRDiS of size at most $k$ implies that of a satisfying assignment.
Lemma 5. Any canonical TaRDiS of $\mathcal{G}$ includes at least 2 vertices from each variable gadget and at least 2 vertices from each clause gadget, and so has size at least $k$.

Proof. Recall that a canonical TaRDiS $S$ by definition is incident only to locally earliest edges. In the case of $\mathcal{G}$, none of the literal vertices can belong to $S$, i.e. $S \subseteq V \backslash L$ where $L$ is the set of literal vertices. For each variable $x_{i}$ and clause $c_{j}$, no vertex from $V_{i}$ reaches any vertex from $Q_{j}$, and vice versa. Note that, for all $j \in[m]$, every vertex in $Q_{j}$ reaches exactly 4 other vertices in $Q_{j}$, for example $R_{q_{j}^{3}}=\left\{q_{j}^{1}, q_{j}^{2}, q_{j}^{3}, q_{j}^{4}\right\}$. Thus, $\left|S \cap Q_{j}\right| \geq 2$, else $S$ would not reach every vertex in $Q_{j}$. Similarly, for all $i \in[n]$, every vertex in $V_{i}$ reaches at most 6 other variable gadget vertices, so $\left|S \cap V_{i}\right| \geq 2$.

Lemma 6. If $\mathcal{G}$ admits a TaRDiS $S$ of size $k$ then $\phi$ is satisfiable.
Proof. Suppose $\mathcal{G}$ admits a TaRDiS $S$ of size at most $k$. We first apply Lemma 2, which allows us to assume that $S$ is canonical. Then, by Lemma 5 , we have that $S$ is of size exactly $k$, and that for all $i$ and $j,\left|S \cap V_{i}\right|=2$ and $\left|S \cap Q_{j}\right|=2$.

We show that a satisfying assignment $X$ to the variables of $\phi$ can be obtained from $S$. For each variable $x_{i}$, assign $x_{i}=$ True if $T_{i}^{1}$ or $T_{i}^{2} \in S$, and $x_{i}=$ False otherwise. Note that, since $S$ is a canonical TaRDiS, $a_{i}$ and $b_{i}$ are not in $S$. Hence $F_{i}^{1} \in S$ or $F_{i}^{2} \in S$ if and only if $x_{i}=$ False in our assignment. Suppose for contradiction there is some clause $c_{j}$ which is not satisfied. Then, under this assumption, the gadget $c_{j}$ is incident to three literal vertices $\{u, v, w\}$, none of which are reached from their respective variable gadgets. Since $S$ is a TaRDiS, $q_{j}^{1}$ or $q_{j}^{6}$ is in $S$ (as $u$ must be reached from within $Q_{j}=\left\{q_{j}^{1}, \ldots, q_{j}^{6}\right\}$ ), and likewise $q_{j}^{2}$ or $q_{j}^{3}$ and $q_{j}^{4}$ or $q_{j}^{5}$ are in $S$ (as $v$ and $w$ respectively must be reached from within $Q_{j}$ ). Hence $\left|Q_{j} \cap S\right| \geq 3$. By applying Lemma 5 we get that $|S| \geq 2 m+2 n+1=k+1$, a contradiction. Hence $X$ must satisfy $\phi$.

Lemma 7. If $\phi$ is satisfiable then $\mathcal{G}$ admits a TaRDiS $S$ of size $k$.
Proof. Given a satisfying assignment to $\phi$, we construct $S$ as follows. We start with $S=\cup_{i}\left\{v_{i}^{2}\right\}$. Then if $x_{i}=$ True (respectively $x_{i}=$ False) under the assignment, add $T_{i}^{1}$ to $S$ (resp. $F_{i}^{1} \in S$ ). Now $S$ has size $2 n$ and a vertex in $S$ reaches the literal vertices for every literal set to True in the assignment, but none of the clause gadget vertices.

For each clause $c_{j}$, there is some literal in $c_{j}$ which is assigned True, and hence at least one literal vertex incident to $Q_{j}$ is reached from the corresponding variable gadget. Denote this literal vertex $u$, and the other literal vertices incident to $Q_{j}$ (which may or may not be reached from their respective variable gadgets) $v$ and $w$, respectively. We add the neighbours of $v, w$ in $Q_{j}$ to $S$. That is, $S=S \cup\left((N[v] \cup N[w]) \cap Q_{j}\right)$. Now $S$ has size $2 n+2 m$ and reaches: all variable gadget vertices; all clause gadget vertices; and all literal vertices (since every literal vertex is incident to a clause gadget). Hence $S$ is a TaRDiS of size exactly $k$.

Theorem 1. Happy TaRDiS is NP-complete, even restricted to instances where the footprint of the temporal graph is planar and its lifetime is 3.

Proof. We have membership of NP from Lemma 1. The construction above produces an instance $(\mathcal{G}, k)$ of Happy TaRDiS from an instance $\phi$ of Planar Exactly 3-Bounded 3-SAT in polynomial time. Combining Lemmas 6 and 7 shows that $\mathcal{G}$ admits a TaRDiS of size at most $k$ if and only if $\phi$ is satisfiable.

This result generalises to both Strict and Nonstrict TaRDiS, giving us NP-completeness of all variants with bounded lifetime and a planar footprint.

### 2.3 NP-completeness of Nonstrict TaRDiS with lifetime 2

We show hardness of Nonstrict TaRDiS by reducing from Set Cover, which is known to be NP-complete [46] and W[2]-hard with respect to the parameter $k$ [17]. The SET Cover problem is defined as follows.

[^2]Given an instance of SET Cover, we build a temporal graph $\mathcal{G}$ with vertex set $V(\mathcal{G})=\mathcal{U} \cup \mathcal{S} \cup\left\{a_{i}^{j} \mid \exists x_{i} \in\right.$ $\left.\mathcal{U}, s_{j} \in \mathcal{S}: x_{i} \in s_{j}\right\}$. To this vertex set we add the edges:

- connecting $s_{j}$ in a path at time 1 to all $a_{i j}^{j}$ for all $i$ such that $a_{i}^{j}$ exists;
- connecting $x_{i}$ in a path at time 2 to all $a_{i}^{j}$ for all $j$ such that $a_{i}^{j}$ exists;
- $\left(s_{j}, s_{j+1}\right)$ at time 2 for all $j \in[1, m-1]$.

A sketch of $\mathcal{G}$ can be found in Figure 7. Note that the maximum degree of $\mathcal{G}$ is 4 . We keep the same value of $k$ in our construction; that is, we construct an instance $(\mathcal{G}, k)$ of Nonstrict TARDiS from the instance ( $\mathcal{U}, \mathcal{S}, k$ ) of Set Cover.
Lemma 8. Any TaRDiS in $\mathcal{G}$ can be reduced to a TaRDiS of the same or smaller size in $V(\mathcal{G}) \cap \mathcal{S}$.
Proof. Suppose there is a TaRDiS $S$ in $\mathcal{G}$ containing a vertex $v \in V(\mathcal{G}) \backslash \mathcal{S}$. We have two cases:

1. $v=x_{i}$ for some $1 \leq i \leq n$, or
2. $v=a_{i}^{j}$ for some $1 \leq i \leq n$ and $1 \leq j \leq m$.

In the first case, we can swap $x_{i}$ for any neighbour (or delete $x_{i}$ if $N\left[x_{i}\right] \subset S$ ). This must be a vertex $a_{i}^{j}$ for some $1 \leq j \leq m$. Since the reachability set of $a_{i}^{j}$ is a strict superset of that of $x_{i}, S$ remains a TaRDiS when $x_{i}$ is replaced by its neighbour. This leaves us with case 2 . The reachability sets of $a_{i}^{j}$ are equal to the reachability set of $s_{j}$ for all $a_{i}^{j} \in V(\mathcal{G}) \backslash(\mathcal{U} \cup \mathcal{S})$. Therefore, we can replace vertices in $S$ with those corresponding to sets in $\mathcal{S}$ without changing the reachability set of the TaRDiS.
Lemma 9. The temporal graph $\mathcal{G}$ admits a TaRDiS of cardinality $k$ if and only if $(\mathcal{U}, \mathcal{S}, k)$ is a yes-instance of Set Cover.
Proof. Suppose there is a set $\left\{s_{j \in I}\right\}$ of cardinality $k$ which is a set cover of $\mathcal{U}$. Then we claim the same set $S$ is a $\operatorname{TaRDiS}$ in $\mathcal{G}$. If an element $x_{i}$ is in the set $s_{j}$, then by construction of $\mathcal{G}$, there is a path from $s_{j}$ to $x_{i}$ via $a_{i}^{j}$. Furthermore, all vertices $s_{j^{\prime}}$ for $1 \leq j^{\prime} \leq m$ are temporally reachable from any $s_{j}$. Since all $x_{i}$ appear in at least one of $\left\{s_{j \in I}\right\}$, we know that all $x_{i}$ and all $s_{j}$ are temporal reachability dominated by $S=\left\{s_{j \in I}\right\}$. What remains to check is that all $a_{i}^{j}$ are also reached by a vertex in $S$. Each $a_{i}^{j}$ is connected to $x_{i}$ at time 2. Therefore, since all $x_{i}$ are reached at time 2 from $S$, each $a_{i}^{j}$ must also be temporally reachable from $S$.

Now suppose that there is a TaRDiS $S$ of cardinality $k$ of $\mathcal{G}$. By Lemma 8, we can assume that $S \subseteq \mathcal{S}$. Observe that there is a nonstrict temporal path from a vertex $s_{j}$ reaches a vertex $x_{i}$ if and only if $x_{i} \in s_{j}$. Hence, a vertex $x_{i}$ is temporally reachable from some vertex in $S$ if and only if a set $s_{j}$ containing $x_{i}$ is in $S$. Since every $x_{i}$ is reachable from some vertex in $S$, then, $S$ is a set cover.

We note that since we use the same input $k$ for both problems, Nonstrict TaRDiS is $W[2]$-hard with respect to $k$. This proves the following theorem.
Theorem 2. For every lifetime $\tau \geq 2$ Nonstrict TARDIS is NP-complete and W[2]-hard with respect to $k$. This holds even on graphs with maximum degree 4.

We now extend our construction to obtain W[2]-hardness of HAPPY TARDIS. Consider the following construction for a happy temporal graph $\mathcal{G}^{\prime}$. Begin with the temporal graph $\mathcal{G}$ as constructed earlier in Figure 7. Then, for each $i$, add edges such that the set $\left\{s_{i}\right\} \cup \bigcup_{j} a_{j}^{i}$ forms a clique. For each $j$, we also add edges such that the set $\left\{x_{j}\right\} \cup \bigcup_{i} a_{j}^{i}$ forms a clique. Then we order all time-edges arbitrarily such that all edges in each clique containing edges formerly in $E_{1}$ are active before the edges in each clique containing all time-edges formerly in $E_{2}$. This concludes the description of $\mathcal{G}^{\prime}$. Observe that $\mathcal{G}^{\prime}$ is a happy temporal graph, since each snapshot contains just one edge.
Lemma 10. The temporal graph $\mathcal{G}^{\prime}$ admits a TaRDiS of cardinality $k$ if and only if $(\mathcal{G}, k)$ is a yes-instance of HAPPY TARDIS.

Proof. We show that a vertex $u$ temporally reaches a vertex $v$ in $\mathcal{G}$ if and only if $u$ also reaches $v$ in $\mathcal{G}^{\prime}$. Thus, proving that $\mathcal{G}^{\prime}$ admits a TaRDiS of cardinality $k$ if and only if $\mathcal{G}$ admits a TaRDiS of cardinality $k$.

Suppose we have a path $p$ from a vertex $u$ to a vertex $v$ in $\mathcal{G}$. We have two cases; namely, the edges in the path are assigned the same time in $\mathcal{G}$ or they are not. In the former, there must still be a (one-edge) path from $u$ to $v$ in $\mathcal{G}^{\prime}$ since we have added edges such that all connected components in either snapshot of $\mathcal{G}$ form a clique. In the latter, we can break $p$ into a path at time one $p_{1}$ and a path $p_{2}$ at time two. These paths can each be replaced by an edge as in the case where $p$ consists of edges which occur at the same time. Therefore, for every path in $\mathcal{G}$, there exists a path with the same starting and terminal vertex in $\mathcal{G}^{\prime}$.

Now suppose we have a path $p^{\prime}$ in $\mathcal{G}^{\prime}$ from $u^{\prime}$ to $v^{\prime}$. Since we have cliques for each connected component in the snapshots of $\mathcal{G}$, the path consisting of fewest edges in $\mathcal{G}^{\prime}$ from $u^{\prime}$ to $v^{\prime}$ must consist of at most 2 edges.


Fig. 7: Sketch for the reduction from Set Cover to Nonstrict TaRDiS

Suppose without loss of generality that $p^{\prime}$ is the shortest path from $u^{\prime}$ to $v^{\prime}$. If $p^{\prime}$ consists of a single edge, then $u^{\prime}$ and $v^{\prime}$ must be in the same connected component of $\mathcal{G}$, thus there exists a path from $u^{\prime}$ to $v^{\prime}$ in $\mathcal{G}$. Now suppose that $p^{\prime}$ consists of two edges. We know that $v^{\prime}$ and $u^{\prime}$ are not adjacent in $\mathcal{G}^{\prime}$, else $p^{\prime}$ would not be the shortest path from $u^{\prime}$ to $v^{\prime}$. Therefore, $u^{\prime}$ and $v^{\prime}$ must be members of two different cliques in $\mathcal{G}^{\prime}$. Let $w^{\prime}$ be the vertex adjacent to both $u^{\prime}$ and $v^{\prime}$ which $p^{\prime}$ traverses. Then, since $u^{\prime}$ and $w^{\prime}$ share an edge which is earlier than the edge $\left(w^{\prime}, v^{\prime}\right)$, they must be in the same connected component of $\mathcal{G}_{1}$. Similarly, $w^{\prime}$ and $v^{\prime}$ must be in the same connected component of $\mathcal{G}_{2}$. Therefore, there is a path from $u^{\prime}$ to $v^{\prime}$ in $\mathcal{G}$. Hence, a vertex $u$ temporally reaches a vertex $v$ in $\mathcal{G}$ if and only if $u$ also reaches $v$ in $\mathcal{G}^{\prime}$.

Corollary 3. HAPPY TARDIS is NP-complete and W[2]-hard with respect to $k$.

### 2.4 Algorithm for TaRDiS on Trees

In this section, we show that each variant of TARDIS can be solved in polynomial time on inputs which are trees, even when $k$ and lifetime are unbounded. We introduce some preliminary notions needed for the algorithm.

In this section, we deal with non-simple temporal graphs; each edge $e \in E(\mathcal{G})$ may appear several times. We define the functions $\lambda_{\min }: E \rightarrow[\tau]$ and $\lambda_{\max }: E \rightarrow[\tau]$ which return the earliest and latest appearance of an edge, respectively.

In the following we refer to rooted trees. A rooted tree is a tree $T=(V, E)$ with a special vertex $r$; the root. A vertex $w$ is a descendant of a vertex $v$ if $v$ lies on the unique path from $w$ to $r$. We refer to the vertices neighbouring $v$ which are descendants of $v$ as its children. If a vertex $w$ is a child of $v$, then $v$ is its parent. Similarly, the parent of a parent vertex is referred to as a grandparent. The subgraph induced by a vertex $v$ and its descendants is referred to as the subtree rooted at $v$. The input of the algorithm is a rooted temporal tree $\mathcal{G}$. By this, we mean a temporal graph whose footprint $\mathcal{G}_{\downarrow}$ is a rooted tree. We define the function $C(p)$ to return the set of children for a parent vertex $p$.

### 2.4.1 Intuition

The algorithm operates as follows. We initialise a counter $k$ and two empty sets. The first set $S$ will become a TaRDiS if a TaRDiS of size $k$ exists. The second set $M$ is a set of vertices which are temporally reachable from those in $S$ or will be temporally reachable from a later vertex added to $S$. Roughly, these are the vertices we do not need to worry about. We refer to them as marked. Vertices not in $M$ will be referred to as unmarked. We refer to the distance of a vertex in the tree from the root as its depth.

We work from the leaves of $\mathcal{G}$ to the root $r$. At each iteration, we either mark vertices or reduce the number of appearances of an edge. We continue until the unmarked graph is a star or empty. The star graph is the complete bipartite graph $K_{1, k}$. Its central vertex is a vertex incident to every edge the graph if $k \geq 1$, and is the sole vertex in the graph otherwise.

If a vertex $v$ with parent $p$ and grandparent $g$ is reachable from its grandparent using the earliest time on the edge $(v, p)$ and the latest appearance of $(p, g)$, it will be reached by a vertex in $S$ if its parent is. Therefore, we can mark that vertex. If a vertex $v$ is not marked and also not temporally reachable from its grandparent, we must add a vertex in $N[v]$ to $S$. We solve Strict TaRDiS or Nonstrict TaRDiS using the variable $s$ which depends on the Boolean flag Strict.

## Algorithm: TaRDiS on Trees

Input: A rooted temporal tree $\mathcal{G}$ and a Boolean flag Strict.
Output: A TaRDiS $S$ of minimum size.

1. Initialise $S=\emptyset, M=\emptyset$, and $s=0$.
2. If Strict:
(a) Set $s=1$.
3. While $M \neq V(\mathcal{G})$ :
(a) If $\mathcal{G} \backslash M$ is a star:
(i) Add the central vertex to $S$ and return $S$.
(b) Denote $L$ the set of vertices in $V(\mathcal{G}) \backslash M$ at maximum depth.
(c) Let $p$ be the parent of an arbitrary vertex from $L$.
(d) Choose the vertex $l \in C(p) \backslash M$ which minimizes $\lambda_{\max }(p, l)$.
(e) Let $g$ be the parent of $p$ (and hence the grandparent of $l$ ).
(f) While $l$ is not in $M$ :
(i) If $\lambda_{\max }(l, p)<\lambda_{\min }(p, g)+s$ :
(A) Add $p$ to $S$.
(B) Find $R_{p}$.
(C) Add $R_{p}$ to $M: M=M \cup R_{p}$.
(ii) Else, if $\lambda_{\max }(l, p) \geq \lambda_{\max }(p, g)+s$ :
(A) Do $M:=M \backslash\{p\}$ (note this does nothing if $p \notin M$ already).
(B) Add all children of $p$ to $M$.
(iii) Else, update $\lambda$ : do $\lambda(p, g)=\lambda(p, g) \backslash\left\{\lambda_{\max }(p, g)\right\}$.
4. Return $S$.

Lemma 11. The algorithm TaRDiS on Trees always terminates in time $O\left(|\mathcal{E}|^{2}\right)$, where $\mathcal{E}$ is the set of time-edges in $\mathcal{G}$.
Proof. The algorithm operates by adding vertices to a set $S$ until the whole tree is marked. Observe that each iteration of the inner while-loop (step: 3f) results in either an increase in size of the set $M$ of marked vertices or removal of an appearance of the time-edge between parent and grandparent vertices of the vertex in question. Note that, if there is only one appearance of the edge between parent and grandparent vertices, a vertex must be added to $M$. Therefore, there are at most $|\mathcal{E}|$ iterations of the inner while loop. Computing $R_{u}$ is achievable in linear time when $\mathcal{G}_{\downarrow}$ is a tree. Thus, each iteration can be completed in time linear in the number of time-edges. Hence, the algorithm terminates in time $O\left(|\mathcal{E}|^{2}\right)$.

We let $\mathcal{T}_{v}$ denote the subtree rooted at a vertex $v$.
Lemma 12. For any vertex $v$ added to the set $S$ by the algorithm TaRDiS on Trees, $S \cap \mathcal{T}_{v}$ is a TaRDiS of the subtree $\mathcal{T}_{v}$ rooted at $v$.

Proof. We show that all descendants of $v$ are temporal reachability dominated by $S \cap \mathcal{T}_{v}$. Note that the children of $v$ are necessarily temporal reachability dominated by $S$ since they are adjacent to the vertex $v$. Before $v$ is added to $S$, it must have a child $l$ which is an unmarked vertex at maximum depth. Therefore, all vertices deeper in the subtree must be in $M$.

We will show that all vertices in $M \cap \mathcal{T}_{v}$ are temporally reachable from a vertex in $S$. We claim that, if a vertex $u$ is in $M \cap \mathcal{T}_{v}$ but not temporally reachable from $S \backslash\{v\}$, then there is a temporal path from $v$ to $u$. This is the because edge from $u$ to its parent $p_{u}$ is active strictly later (or possibly at the same time in the nonstrict case) than the last appearance of the edge from $p_{u}$ to the grandparent vertex $g_{u}$. Otherwise, $u$
would not be marked. Note that this implies that neither $g_{u}$ nor $p_{u}$ are temporally reachable from a vertex in $S \backslash\{v\}$; otherwise $u$ would also be reachable by appending the edges $\left(g_{u}, p_{u}\right)$ and ( $\left.p_{u}, u\right)$ to such a path.

This either gives us an ancestor of $u$ which is not in $M$, or an ancestor of $u$ in $M$ which is not temporally reachable from a vertex in $S \backslash\{v\}$. If the latter is true, we repeat this logic until we reach an ancestor $w$ which is not in $M$. This must be either $v$ or a child of $v$ since $l$ is an unmarked vertex of maximum depth. For all descendants $w^{\prime}$ of $w$ that are $u$ or ancestors of $u$, we have shown that if the parent of $w^{\prime}$ is temporally reached by a path from an ancestor, then $w^{\prime}$ must be too. Since $v$ is either $w$ or the parent of $w, u$ and all of its ancestors in $\mathcal{T}_{v}$ must be temporally reachable from $v$ and thus $S \cap \mathcal{T}_{v}$ is a TaRDiS for $\mathcal{T}_{v}$.

Lemma 13. For any vertex $v$ added to the set $S$ by the algorithm TaRDiS on Trees, $S \cap \mathcal{T}_{v}$ is a minimum TaRDiS of the subtree $\mathcal{T}_{v}$ rooted at $v$.

Proof. Suppose for contradiction that such a vertex $v \in S$ and set $S^{\prime}$ exist where $S^{\prime} \subseteq T_{v}$ is a TaRDiS of $\mathcal{T}_{v}$ and strictly smaller than $S \cap \mathcal{T}_{v}$. We assume without loss of generality that $S^{\prime}$ is a minimal TaRDiS.

Consider the set $S^{\prime} \cap S$. If this is non-empty, let $R$ be the union of reachability sets of vertices in $S^{\prime} \cap S$. Let $\mathcal{T}^{\prime}$ be the subgraph of $\mathcal{T}_{v}$ induced by removing all vertices in $R \backslash\left(S^{\prime} \cap S\right)$. Since we have assumed that $\left|S^{\prime}\right|<|S|$, there must be a connected component $\mathcal{T}^{\prime \prime}$ of $\mathcal{T}^{\prime}$ where $\left|\mathcal{T}^{\prime \prime} \cap S^{\prime}\right|<\left|\mathcal{T}^{\prime \prime} \cap S\right|$.

Let $u \in \mathcal{T}^{\prime \prime} \cap S$ be the vertex in $\left(\mathcal{T}^{\prime \prime} \cap S\right) \backslash S^{\prime}$ of maximum depth. Observe that $u$ is not a leaf in $\mathcal{T}^{\prime \prime}$, otherwise it would not have an unmarked child before its addition to $S$ by the algorithm. Denote by $c$ the child of $u$ which minimises $\lambda_{\max }(u, c)$. Since $u$ is added to $S$, we must have that $\lambda_{\max }(c, u)<\lambda_{\min }(u, p)+s$ where $p$ is the parent of $u$. Therefore, the path consisting of time-edges $\left((u, p), \lambda_{\min }(u, p)\right),\left((c, u), \lambda_{\max }(c, u)\right)$ is not a valid temporal path. We now have two cases to consider. The case where appearances of the edge $(c, u)$ have been deleted and the case where no appearances of the edge $c u$ have been deleted by the algorithm.

If no appearances of $(u, c)$ have been removed by the algorithm, then $c$ is not temporally reachable from $p$ or any ancestors of $p$. Since we know that $u$ is in $S$ and not $S^{\prime}, c$ must be temporally reachable from a descendant $c^{\prime}$ of $u$ in $S^{\prime} \backslash S$. We claim that we can swap $c^{\prime}$ with $u$ without reducing the reachability set of vertices in $S^{\prime}$. Suppose otherwise; that there is a vertex $w$ no longer reached by $S^{\prime}$. Since we assumed that $u$ was a vertex of maximum depth in $S^{\prime} \backslash S$, this contradicts that $S \cap \mathcal{T}_{u}$ is a TaRDiS of $\mathcal{T}_{u}$. Thus, this contradicts Lemma 12.

Now suppose that there are appearances of the edge $(u, c)$ which have been removed by the algorithm. Then, a child of $c$ must have been an unmarked vertex of maximum depth. Call this vertex $g$. This implies that $g \in \mathcal{T}^{\prime \prime}$. If the algorithm were to add $c$ to $S$ when $g$ is unmarked, then $c$ would not be an unmarked vertex of maximum depth in any later iterations. Before deletion of appearances of $(u, c)$, we must have that $\lambda_{\max }(g, c)<\lambda_{\max }(c, u)+s$ and $\lambda_{\max }(g, c) \geq \lambda_{\min }(c, u)+s$. Following deletion of the appearance $\lambda_{\max }(c, u)$ (potentially multiple times), $\lambda_{\max }(g, c) \geq \lambda_{\max }(c, u)+s$ and $\lambda_{\max }(c, u)<\lambda_{\min }(u, p)+s$. Therefore, there is no temporal path from $p$ or an ancestor of $p$ to $g$. Since we know that $u$ is in $S$ and not $S^{\prime}, g$ must be temporally reachable from a descendant $g^{\prime}$ of $u$ in $S^{\prime} \backslash S$. We claim that we can swap $g^{\prime}$ with $u$ without reducing the reachability set of vertices in $S^{\prime}$. Suppose otherwise; that there is a vertex $w$ no longer reached by $S^{\prime}$. Since we assumed that $u$ was a vertex of maximum depth in $S^{\prime} \backslash S$, this contradicts that $S \cap \mathcal{T}_{u}$ is a TaRDiS of $\mathcal{T}_{u}$. Thus, this also contradicts Lemma 12.

In both cases, we have shown that there is a vertex in $S^{\prime} \backslash S$ that can be replaced by $u$ without reducing the reachability set of vertices in $S^{\prime}$. We repeat this with the vertex at maximum depth in $\mathcal{T}^{\prime \prime} \cap S \backslash S^{\prime}$ until either we contradict the assumption that $S^{\prime}$ is a $\operatorname{TaRDiS}$ of $\mathcal{T}_{v}$, or obtain $S=S^{\prime}$; contradicting our assertion that $\left|S^{\prime}\right|<|S|$. Therefore $S$ is a minimum TaRDiS of $\mathcal{T}_{v}$.

Lemma 14. The set $S$ output by the algorithm TaRDiS on Trees is a minimum TaRDiS of $\mathcal{G}$.
Proof. If the root $r$ of $\mathcal{G}$ is in $S$ as output by the algorithm, then combining Lemmas 12 and 13 gives us the desired result. Suppose $r$ is not in $S$. Then, rooting the tree at the final vertex added to $S$ gives the result without changing the vertices in $S$.

This gives us the following theorem and corollary.
Theorem 3. When the footprint of the graph is a tree, TARDIS is solvable in $O\left(|\mathcal{E}|^{2}\right)$ time.
We obtain a running time dependent only on the number of vertices for simple temporal graphs, where $|\mathcal{E}|=\left|E\left(\mathcal{G}_{\downarrow}\right)\right|=V(\mathcal{G})-1$.
Corollary 4. When the input temporal graph is simple and the footprint of the graph is a tree, TARDIS is solvable in $O\left(n^{2}\right)$ time.

In Section 4, we give a more general algorithm which solves any variant of TARDIS on a nice tree decomposition of the underlying graph. When the footprint graph is a tree, we have treewidth 1. Our algorithm for TARDIS on trees runs faster than the tree decomposition algorithm (whose runtime also grows with $\tau$ ), so this does not consume Theorem 3 .

## 3 Classical complexity results for MaxMinTaRDiS

We now shift our focus to MaxMinTaRDiS, where the objective is to find a temporal assignment which precludes the existence of a small TaRDiS. For each variant, we characterize the minimum lifetime such that the problem becomes intractable. Intriguingly, the expected leap in complexity to $\Sigma_{2}^{P}$ is only seen with Happy MaxMinTaRDis.

We give a full definition of $\Sigma_{2}^{P}$ in Section 3.2. It is widely believed that $\Sigma_{2}^{P}$ is a strict superclass of NP $\cup$ coNP. In this sense, Happy MaxMinTaRDiS is a harder problem than both Strict MaxMinTaRDiS (which is coNP-complete), and all TARDiS variants (which are NP-complete as shown in Section 2). We show Nonstrict MaxMinTaRDiS is at least as hard as all TaRDiS variants and at most as hard as Happy MaxMinTaRDiS, but leave open its exact complexity. Interestingly, this is achieved by proving that the well-studied Distance-3 Independent Set problem is a subproblem of Nonstrict MaxMinTaRDiS.

### 3.1 Containment in $\Sigma_{2}^{P}$, useful tools, and small lifetime

Here we give some preliminary complexity results for each variant of MaxMinTaRDiS. We begin by showing that they are all contained in $\Sigma_{2}^{P}$. We then show that we need only consider simple temporal assignments when trying to solve MaxMinTaRDiS. Finally, we draw comparisons between Happy MaxMinTaRDiS and Strict MaxMinTaRDiS and the classical problems Edge Colouring and Dominating Set respectively, which allow us to show NP-hardness and coNP-hardness of the respective problems.
Lemma 15. Each variant of the problem MaxMinTaRDIS is contained in $\Sigma_{2}^{P}$.
Proof. (Strict/Nonstrict/Happy) MaxMinTaRDiS admits a triple ( $H, k, \tau$ ) as a yes-instance if and only if there exists a (happy) temporal assignment $\lambda: E(H) \rightarrow \tau$ such that for every (Strict/Nonstrict) TaRDiS $S$ of $(H, \lambda),|S| \geq k$. It follows that a formula $\Phi$ over variable sets $X$ and $Y$ can be produced such that $(H, k, \tau)$ is a yes-instance if and only if there exists an assignment to the variables in $X$ such that for every assignment to the variables of $Y, \Phi(X, Y)$ is True. We refer the reader to Section 3.2 for a more detailed construction of such a Boolean formula.

We begin with the observation that we need only consider simple temporal assignments for the input graph.
Lemma 16. Let $(H, k, \tau)$ be a yes-instance of MAXMINTARDIS. Then there exists a simple temporal assignment $\lambda: E \rightarrow[\tau]$ such that the cardinality of the minimum TaRDiS on $(H, \lambda)$ is at least $k$.
Proof. We begin by supposing, for a contradiction, that $(H, k, \tau)$ is an instance of MaxMinTaRDiS such that any optimal solution is non-simple. Denote such a solution by $\lambda^{*}$. By our assumption, there is at least one edge $e^{*} \in E(H)$ such that $\left|\lambda^{*}\left(e^{*}\right)\right|>1$. Let $\lambda$ be a simple temporal assignment such that, for all edges $e \in E(H), \lambda(e) \in \lambda^{*}(e)$. Then, under $\lambda$, the reachability set of any vertex $v \in V(H)$ is a subset of the reachability set of $v$ under $\lambda^{*}$. Therefore, a minimal TaRDiS of $(H, \lambda)$ must be at most the cardinality of a minimal TaRDiS of $\left(H, \lambda^{*}\right)$. Since $\lambda$ is a simple temporal assignment, we have contradicted our assumption. Thus, for every instance of MAxMinTARDIS, there exists an optimal solution which is simple.

Lemma 17 ([10]). A static graph $H$ admits a happy temporal assignment $\lambda: E(G) \rightarrow[\tau]$ if and only if $H$ is $\tau$-edge colourable.
Proof. To see this, we assign each of the times a colour. Then, if the graph cannot be $\tau$-edge coloured, we cannot assign times to the edges such that no two adjacent edges share a time. Furthermore, if we can give $H$ a happy temporal assignment, then the corresponding assignment of colours to edges is a proper edge-colouring.

Happy MaxMinTaRDiS restricted to instances with $k=0$ asks only if there exists a happy assignment $\lambda$ with lifetime $\tau$ for the input graph $G$. This is equivalent to the Edge colouring problem with $\tau$ colours. Edge Colouring is NP-complete, even when the number of colours is 3 [47].
Corollary 5. Happy MAXMINTARDIS is NP-hard for any $\tau \geq 3$, even when $k=0$.
Lemma 18. Nonstrict MaxMinTaRDiS can be computed in linear time when $\tau=1$. When $\tau \leq 2$, Happy MaxMinTaRDiS is solvable in linear time.

Proof. If every edge is active at the same time, we have no choice in the temporal assignment. As shown in Lemma 4, we can find the size of the minimum nonstrict TaRDiS in linear time by counting the number of connected components.

As stated in Lemma 4, the footprint of a happy temporal graph with lifetime 2 necessarily consists of disconnected paths and even cycles. In paths and cycles of even length, there is only one happy temporal
assignment up to symmetry. In odd paths, the optimal ordering is that wherein both edges incident to leaves are assigned time 1 . This forces one of the endpoints of these edges to be in a TaRDiS, which is not the case if they are assigned time 2 .

Lemma 19. For any static graph $H$ and $k \in \mathbb{N}^{+},(H, k)$ is a yes-instance of Strict MaxMinTaRDiS if and only if $(H, k-1)$ is a no-instance of Dominating Set.

Proof. We first show that the constant function $\lambda$ is an optimal one (i.e. one which maximizes the size of the minimum TaRDiS for $(H, \lambda)$ ). Suppose otherwise, that we have an optimal temporal assignment $\lambda^{\prime}$ where there exist edges $e, e^{\prime}$ such that $\lambda^{\prime}(e) \neq \lambda^{\prime}\left(e^{\prime}\right)$. Then, under a temporal assignment $\lambda(e)=c$ for all edges $e \in E$ and $c \in \mathbb{N}$, all vertices have reachability sets whose cardinality are bounded above by the size of their reachability set under $\lambda^{\prime}$. Therefore a minimum TaRDiS under $\lambda$ must be at least the size of a minimum TaRDiS under $\lambda^{\prime}$. Applying Lemma 3, we see that $(H, k)$ is a yes-instance of Strict MaxMinTaRDiS if and only if $(H, k-1)$ is a no-instance of Dominating Set.

Corollary 6. Strict MaxMinTaRDiS is coNP-complete, co W[2]-complete with respect to $k$, and para$N P$-hard with respect to $\Delta+\tau$.

## $3.2 \Sigma_{2}^{P}$-completeness of Happy MaxMinTaRDiS with lifetime 3

As stated in Corollary 5, Happy MaxMinTaRDiS is trivially NP-hard even for lifetime $\tau=3$. Lemma 15 gives us that the problem is in $\Sigma_{2}^{P}$. We characterize the problem's complexity exactly, showing it is $\Sigma_{2}^{P}$ complete.
Definition 3 ( $\Sigma_{2}^{P}$ - adapted from [48], Definition 5.1). $\Sigma_{2}^{P}$ is the set of all languages $L$ for which there exists a polynomial-time Turing Machine $M$ and a polynomial $q$ such that

$$
x \in L \Longleftrightarrow \exists u \in\{0,1\}^{q(|x|)} \forall v \in\{0,1\}^{q(|x|)} M(x, u, v)=1
$$

for every $x \in\{0,1\}^{*}$.
Arora and Barak [48] note that NP $\cup \mathrm{co-NP} \subseteq \Sigma_{2}^{P}$ and that $\Sigma_{1}^{P}=\mathrm{NP}$. The complement of a $\Pi_{i}^{P}$-complete problem is necessarily $\Sigma_{i}^{P}$-complete [49].

We begin by presenting its problem Restricted Planar Satisfiability, which is $\Pi_{2}^{P}$ complete [50].

## Restricted Planar Satisfiability (RPS)

Input: An expression of form $(\forall X)(\exists Y) \Phi(X, Y)$ with $\Phi$ a CNF formula over the set $X \cup Y$ of variables. Each clause contains exactly 3 distinct variables, each variable occurs in exactly three clauses, every literal appears at most twice, and the graph $G_{\Phi}$ is planar.
Question: Is the expression true?
We also introduce the complement problem co-RPS, which we will reduce from.

```
co-Restricted Planar Satisfiability (co-RPS)
```

Input: An expression of form $(\exists X)(\nexists Y) \Phi(X, Y)$ with $\Phi$ a CNF formula over the set $X \cup Y$ of variables. Each clause contains exactly 3 distinct variables, each variable occurs in exactly three clauses, every literal appears at most twice, and the graph $G_{\Phi}$ is planar.
Question: Is the expression true?

Lemma 20. The problem co-RPS is $\Sigma_{2}^{P}$ complete.
Proof. Gutner shows Restricted Planar Satisfiability (RPS) to be $\Pi_{2}^{P}$-complete in [50]. Note that they do not state the restriction that each literal appears at most twice in the lemma stating their result, but this can be seen from their construction.

### 3.2.1 Intuition

We can imagine problems in $\Sigma_{2}^{P}$ as games with two players. In co-RPS the first player chooses an assignment to variables in $X$, then the second player chooses an assignment to the variables in $Y$. If the resulting assignment to $X \cup Y$ satisfies $\Phi$ then the second player wins; otherwise the first player wins. Analogously, in Happy MaxMinTaRDiS the first player chooses a happy temporal assignment $\lambda$ with lifetime $\tau$ for the edges of $G$, then the second player chooses a set $S$ of vertices in $G$ of size at most $k-1$. If $S$ is a TaRDiS of $(G, \lambda)$ then the second player wins; otherwise the first player wins.

We perform this reduction by creating gadgets which force the first player to choose a function $\lambda$ which is nice. Essentially, $\lambda$ must assign specific times to certain edges. This means that if the first player does not do this, the second player can always win. We can then replicate some techniques from the proof of Theorem 1 to encode clauses being satisfied. The choice of TaRDiS by the second player then encodes a truth assignment to variables in $Y$. By allowing a small amount of freedom in the choice of $\lambda$, we allow the first player to encode a truth assignment to the variables of $X$. We now give the formal construction.

### 3.2.2 Construction

Given an instance ( $\Phi, X, Y$ ) of co-RPS, we will produce an instance ( $G, k, \tau=3$ ) of MaxMinTaRDiS. The number of clauses in $\Phi$ is denoted by $m$. We label the number of variables in $X$ (respectively $Y$ ) with $n_{X}$ (resp. $n_{Y}$ ). The total number of variables is $n$. Note that $3 m \geq n \geq m$, so the size of the co-RPS instance is linear in $n$. Further, we denote $X=\left\{x_{1}, \ldots, x_{n_{x}}\right\}$ and $Y=\left\{y_{n_{x}+1}, \ldots, y_{n_{x}+n_{y}}\right\}$. We say that the literal $l_{2 i}$ (respectively, $l_{2 i-1}$ ) is the positive (resp., negative) literal for variable $x_{i}$ if $i \leq n_{x}$, and the positive (resp. negative) literal for variable $y_{i}$ if $n_{x}<i \leq n_{y}$.

We now describe the construction of $G$ and $k$. The construction of our gadgets is intended to guarantee that any optimal temporal assignment $\lambda$ will have certain properties. First, we define the Uncovered 2Gadget, Uncovered 3-Gadget, and Covered 2-Gadget. Each of these is treated in our construction as a vertex of degree one. The construction for each of these makes use of large values $\alpha$ and $\beta$. Let $\beta=100 n$ and $\alpha=600 n \beta+600 n+2$. In particular we require that $\alpha \gg \beta \gg n, \beta$ is even, and $\alpha \equiv 2 \bmod 12$.

### 3.2.3 Uncovered 2-Gadget (U2G)

This construction is illustrated in Figure 8. We use a U2G by connecting it to some vertex $x$ elsewhere in the construction. Given some such $x$, we create vertices $y$ and $w$ and a ladder graph on $2 \alpha$ vertices $\left\{u_{1}, \ldots, u_{\alpha}, v_{1}, \ldots, v_{\alpha}\right\}$, with the edges $(x, y),\left(y, u_{1}\right),\left(y, v_{1}\right),\left(w, u_{\alpha}\right),\left(w, v_{\alpha}\right)$. The ladder graph $L_{n}$ on $2 n$ vertices has vertex set $V\left(L_{n}\right)=\left\{u_{1}, \ldots, u_{n}, v_{1}, \ldots, v_{n}\right\}$, and edge set $E\left(L_{n}\right)=\left\{\left(u_{i}, u_{i+1}\right) \mid i \in[n-1]\right\} \cup$ $\left\{\left(v_{i}, v_{i+1}\right) \mid i \in[n-1]\right\} \cup\left\{\left(u_{i}, v_{i}\right) \mid i \in[n]\right\}$.

### 3.2.4 Covered 2-Gadget (C2G)

The construction is very similar to that of the U2G - the only difference is the incrementation of the ladder length by 1 . Given $x$, we create vertices $y$ and $w$ and a ladder graph on $2 \alpha+2$ vertices $\left\{u_{1}, \ldots, u_{\alpha+1}, v_{1}, \ldots, v_{\alpha+1}\right\}$, and edges $(x, y),\left(y, u_{1}\right),\left(y, v_{1}\right),\left(w, u_{\alpha+1}\right),\left(w, v_{\alpha+1}\right)$.


With $l=\alpha$ :


With $l=\alpha+1$ :


Fig. 8: A vertex $x$ incident to the Uncovered 2 Gadget (U2G) or Covered 2 Gadget (C2G). Left: construction of the gadget where $l=\alpha$ for a U2G and $l=\alpha+1$ for a C2G. Right: usage in later constructions. Note that planarity is preserved.

### 3.2.5 Uncovered 3-Gadget (U3G)

The construction is illustrated in Figure 9. We start with the ladder graph on $2 \alpha$ vertices $\left\{u_{1}, \ldots, u_{\alpha}, v_{1}, \ldots, v_{\alpha}\right\}$, and doubly subdivide the edge ( $u_{i}, u_{i+1}$ ) (respectively ( $\left.v_{i}, v_{i+1}\right)$ ) whenever $i$ is even (resp. odd) with two new vertices $a_{i}, b_{i}$. We say doubly subdivide an edge ( $u, v$ ) to refer to the deletion of ( $u, v$ ) and introduction of two vertices $a, b$ and three edges $(u, a),(a, b),(b, v)$. We add the vertices $y_{1}, y_{2}, y_{3}, w$ and edges $\left(x, y_{1}\right),\left(y_{1}, y_{2}\right),\left(y_{2}, y_{3}\right),\left(y_{3}, u_{1}\right),\left(y_{3}, v_{1}\right),\left(w, u_{\alpha}\right)$. We denote the set of internal vertices of the U3G by $V_{\mathrm{U} 3 \mathrm{G}}=\left\{a_{1}, \ldots, a_{\beta-1}, b_{1}, \ldots b_{\beta-1}, y_{1}, y_{2}, y_{3}, u_{1}, \ldots, u_{\beta}, v_{1}, \ldots, v_{\beta}, w\right\}$. Internal vertices are highlighted by a red (solid) box in Figure 9. We make any vertex not already of degree 3 in $V_{\mathrm{U} 3 \mathrm{G}}$ incident to a U2G.


Fig. 9: A vertex $x$ incident to the Uncovered 3 Gadget (U3G). Left: construction of the gadget; vertices of $V_{\mathrm{U} 3 \mathrm{G}}$ are in the red (solid) box. Right: usage in later constructions. Note that planarity is preserved.

### 3.2.6 Uncovered 1-Gadget (U1G)

For a U1G, we simply create a vertex $y$ incident to both a C2G and a U3G as shown in Figure 10, then add an edge from $y$ to the target vertex $x$.


Fig. 10: A vertex $x$ incident to the Uncovered 1 Gadget (U1G). Left: implementation of the gadget. Right: usage in other constructions.

### 3.2.7 Construction: literal vertices and clause gadgets

Exactly $3 m$ literal vertices and $m$ clause gadgets are created. We iterate over $\Phi$, and for each clause $c_{j}$ we create a cycle on 6 new vertices $Q_{j}=\left\{q_{j}^{1}, \ldots, q_{j}^{6}\right\}$, and make $q_{j}^{2}, q_{j}^{4}, q_{j}^{6}$ each incident to a U2G. Where the clause $c_{j}$ contains the $a$ th appearance of some literal $l_{i}$ in $\Phi$, we create two new literal vertices $l_{i}^{a}$ and $\overline{l_{i}^{a}}$ connected by an edge, and make each of $l_{i}^{a}$ and $\overline{l_{i}^{a}}$ incident to a new U1G. We make $l_{i}^{a}$ incident to a vertex from $\left\{q_{j}^{1}, q_{j}^{3}, q_{j}^{5}\right\}$, such that each of these is adjacent to exactly one literal vertex. Figure 11 illustrates this construction. Vertices $u, v, w, \bar{u}, \bar{v}, \bar{w}$ are literal vertices for which the corresponding literals appear in $c_{j}$.

### 3.2.8 Construction: $X$-variable gadget

This construction is illustrated in Figure 12. For each variable $x_{i} \in X$, we create 13 vertices $\left\{x_{i}^{1}, \ldots, x_{i}^{13}\right\}$ connected in a path, with vertex $x_{i}^{p}$ incident to a U2G if $p$ is even, and incident to a C2G for $p \in\{1,7,13\}$. Lastly we add the edges $\left(x_{i}^{3}, \overline{l_{2 i}^{1}}\right),\left(x_{i}^{5}, \overline{l_{2 i-1}^{1}}\right),\left(x_{i}^{9}, \overline{l_{2 i}^{2}}\right)$, and $\left(x_{i}^{11}, \overline{l_{2 i-1}^{2}}\right)$ whenever the appropriate literal vertex exists. Recall that, by the restrictions on $\Phi$, at most one of the vertices $l_{2 i}^{2}$ and $l_{2 i-1}^{2}$ exist in our construction for any input.

### 3.2.9 Construction: $Y$-variable gadget

This construction is illustrated in Figure 13. For each variable $y_{i} \in Y$, we create 8 vertices $\left\{a_{i}, v_{i}^{1}, v_{i}^{2}, b_{i}, F_{i}^{1}, F_{i}^{2}, T_{i}^{2}, T_{i}^{1}\right\}$ connected in a cycle, with vertices $v_{i}^{1}$ and $v_{i}^{2}$ incident to U3Gs and vertices $a_{i}$ and $b_{i}$ incident to U1Gs. Then we add edges $\left(T_{i}^{1}, \overline{l_{2 i}^{1}}\right),\left(T_{i}^{2}, \overline{l_{2 i}^{2}}\right),\left(F_{i}^{2}, \overline{l_{2 i-1}^{2}}\right),\left(F_{i}^{1}, \overline{l_{2 i-1}^{1}}\right)$ whenever the appropriate literal vertex exists. If vertex $T_{i}^{2}$ is not incident to a literal vertex, we connect it to a U2G. Otherwise, $F_{i}^{2}$ is connected to a U2G.


Fig. 11: The clause gadget for clause $c_{j}$ together with an example nice assignment. If, for example, clause $c_{j}=\left(l_{5} \vee l_{17} \vee l_{20}\right)$ is the first appearance of $l_{5}$ and $l_{20}$ and the second appearance of $l_{17}$ in $\Phi$, then $u=l_{5}^{1}$, $v=l_{17}^{2}$ and $w=l_{20}^{1}$. The neighbourhoods of dashed vertices are shown in Figures 12 and 13 respectively depending on whether they correspond to variables in $X$ or $Y$.


Fig. 12: The $X$-variable gadget for variable $x_{i}$ together with an example gadget-respecting assignment; note that either $A=1$ and $B=3$ or $B=1$ and $A=3$. Only in the former case, which corresponds to setting $x_{i}$ to True, does $x_{i}^{4}$ have a temporal path to $l_{2 i}^{1}$. Here $x_{i}$ appears twice negatively and once positively, hence $x_{i}^{9}$ is not incident to a literal vertex. The neighbourhoods of dashed vertices are shown in Figure 11.


Fig. 13: The $Y$-variable gadget for variable $y_{i}$ together with an example nice assignment. Note the similarity to the variable gadget in Figure 6. Here $x_{i}$ appears twice negatively and once positively, hence $T_{i}^{2}$ is connected to a U2G and not a literal vertex. The neighbourhoods of dashed vertices are shown in Figure 11.

### 3.2.10 Construction: $\boldsymbol{k}$

In order to state $k$, we first define some auxiliary variables. As stated earlier, the number of $X$-gadgets, $Y$-gadgets and clause-gadgets is $n_{X}, n_{Y}$ and $m$ respectively. We then define:

$$
\begin{aligned}
\#_{\text {lit }}=2 \cdot 3 m & \text { (The number of literal nodes.) } \\
\#_{\mathrm{U} 1 \mathrm{G}}=\#_{\mathrm{lit}}+2 \cdot n_{Y} & \text { (The number of U1Gs.) } \\
\#_{\mathrm{U} 3 \mathrm{G}}=2 n_{Y}+\#_{\mathrm{U} 1 \mathrm{G}} & \text { (The number of U3Gs.) } \\
\#_{\mathrm{C} 2 \mathrm{G}}=3 n_{X}+\#_{\mathrm{U} 1 \mathrm{G}} & \text { (The number of C2Gs.) } \\
\#_{\mathrm{U} 2 \mathrm{G}}=\#_{\mathrm{U} 3 \mathrm{G}}(2 \beta+2)+3 m+6 n_{X}+n_{Y} & \text { (The number of U2Gs.) }
\end{aligned}
$$

We now define $k$ :

$$
k=\#_{\mathrm{U} 2 \mathrm{G}} \cdot\left(\frac{\alpha+1}{3}\right)+\#_{\mathrm{C} 2 \mathrm{G}} \cdot\left(\frac{\alpha+4}{3}\right)+\#_{\mathrm{U} 3 \mathrm{G}} \cdot(\beta)+2 m+3 n_{X}+3 n_{Y}+1
$$

This concludes the construction of the Happy MaxMinTaRDiS instance.

### 3.2.11 Properties of the construction

We first define the temporal assignments of interest to us.
Definition 4 (Gadget-respecting, Nice Temporal Assignment). A happy temporal assignment $\lambda$ for the graph $G$ described above is:

- 2-gadget-respecting if every edge ( $x, y$ ) incident to a U2G or C2G is assigned time 2 under $\lambda$.
- 3-gadget-respecting (resp. 1-gadget-respecting) if every edge ( $x, y$ ) incident to a U3G (resp. U1G) is assigned time 3 (resp. 1) under $\lambda$.

If $\lambda$ is 1-, 2- and 3-gadget-respecting, we say that $\lambda$ is nice.
We will show that a happy temporal assignment $\lambda$ for $G$ satisfies that every TaRDiS of $(G, \lambda)$ has cardinality at least $k-1$ if and only if $\lambda$ is nice. We then show that there exists a $\lambda$ such that every TaRDiS of $(G, \lambda)$ has cardinality at least $k$ if and only if $(\Phi, X, Y)$ is a yes-instance of co-RPS.
Lemma 21. For any happy temporal assignment $\lambda$ on any U2 $G$ gadget incident to $x$ consisting of vertices $V_{U 2 G}=\left\{y, w, u_{1}, \ldots, u_{\alpha}, v_{1}, \ldots, v_{\alpha}\right\}$, if $\lambda(x, y)=2$

- exactly $\frac{\alpha+1}{3}$ vertices from $V_{U 2 G}$ are needed to temporal reachability dominate $V_{U 2 G}$;
- no choice of $\frac{\alpha+1}{3}$ vertices temporally dominates $V_{U 2 G} \cup\{x\}$.

If $\lambda(x, y) \neq 2$, it is possible to temporally dominate $V_{U 2 G}$ with exactly $\frac{\alpha+2}{4}$ vertices.
Proof. It is easy to check that there is only one proper 3-colouring of the edges of the induced graph on $V_{\mathrm{U} 2 \mathrm{G}} \cup\{x\}$ up to isomorphism. Denote by $A$ the colour of $(x, y), B$ the colour of $\left(y, u_{1}\right)$ and $C$ the colour of $\left(y, v_{1}\right)$. Note that all edges $\left(u_{i}, v_{i}\right)$ are given colour $A$, and the edges on the path $u_{1}, \ldots, u_{\alpha}$ (resp. $v_{1}, \ldots, v_{\alpha}$ ) alternate between $B$ and $C$.

Any possible happy temporal assignment then corresponds exactly to one of 6 possible assignments of times $\{1,2,3\}$ to the colours $\{A, B, C\}$. From this point, we abuse notation slightly and denote $\lambda(x, y)$ by $\lambda(A), \lambda\left(y, u_{1}\right)$ by $\lambda(B)$ and $\lambda\left(y, v_{1}\right)$ by $\lambda(C)$. By symmetry of the gadget construction, we can always assume $\lambda(B)<\lambda(C)$.

If $\lambda(A)=2$, we have $\lambda(B)=1$ and $\lambda(C)=3$; then $\left|R_{v}\right| \leq 6$ for all $v \in V_{\mathrm{U} 2 \mathrm{G}}$. Note in particular that $R_{v_{1}}=\left\{y, v_{1}, v_{2}, v_{3}, u_{1}, u_{2}\right\}$, and $R_{x} \cap V_{\mathrm{U} 2 \mathrm{G}}=\left\{y, v_{1}\right\}$. Thus for any TaRDiS $S$, at least $\left|V_{\mathrm{U} 2 \mathrm{G}}\right|-2=2 \alpha$ vertices in $V_{\mathrm{U} 2 \mathrm{G}}$ are reached by vertices of $S \cap V_{\mathrm{U} 2 \mathrm{G}}$, each of which reaches at most 6 vertices including itself. Recall that $\alpha \equiv 2 \bmod 12$. Hence, $\left|S \cap V_{\mathrm{U} 2 \mathrm{G}}\right| \geq\left\lceil\frac{2 \alpha}{6}\right\rceil=\frac{\alpha+1}{3}$. Note that, by the same logic, any set of size at most $\frac{\alpha+1}{3}$ reaches at most $2 \alpha+2$ vertices and hence no such set reaches every vertex in $V_{\mathrm{U} 2 \mathrm{G}} \cup\{x\}$ since $\left|V_{\mathrm{U} 2 \mathrm{G}} \cup\{x\}\right|=2 \alpha+3$. Further, the set $\left\{v_{1}, u_{4}, v_{7}, u_{10}, \ldots v_{\alpha-1}\right\}$ is of size $\frac{\alpha+1}{3}$ exactly and reaches every vertex in $V_{\mathrm{U} 2 \mathrm{G}}$. Conversely, if $\lambda(A) \in\{1,3\}$ then $\left\{v_{1}, v_{5}, v_{9}, \ldots, v_{\alpha-1}\right\}$ is a set of size $\frac{\alpha+2}{4}$ which reaches all vertices of $V_{\mathrm{U} 2 \mathrm{G}}$.
Lemma 22. For any happy temporal assignment $\lambda$ on any C2G gadget incident to vertex $x$ consisting of vertices $V_{C 2 G}=\left\{y, w, u_{1}, \ldots, u_{\alpha+1}, v_{1}, \ldots, v_{\alpha+1}\right\}$, if $\lambda(x, y)=2$

- exactly $\frac{\alpha+4}{3}$ vertices from $V_{C 2 G}$ are needed to temporally dominate $V_{C 2 G}$;
- it is possible to temporally dominate $V_{C 2 G} \cup\{x\}$ with the same number of vertices.

If $\lambda(x, y) \neq 2$ it is possible to temporally dominate $V_{C 2 G}$ with exactly $\frac{\alpha+6}{4}$ vertices.
Proof. As in the case of the U2G, there is only one proper 3 -colouring of the edges of the induced graph on $V_{\mathrm{U} 3 \mathrm{G}} \cup\{x\}$ up to isomorphism. We again denote $A$ the colour of $(x, y), B$ the colour of $\left(y, u_{1}\right)$ and $C$ the colour of $\left(y, v_{1}\right)$. Since gadget symmetry is preserved, we also may assume without loss of generality that $\lambda(B)<\lambda(C)$. Then the assignment of the edges on vertices from $V_{\mathrm{C} 2 \mathrm{G}} \cup\{x\}$ depends only on the value of $\lambda(x, y)$.

The dependence of reachability on this choice remains; namely if $\lambda(A)=2$ then $\left|R_{v}\right| \leq 6$ for all $v \in$ $V_{\mathrm{C} 2 \mathrm{G}}$ and $R_{x} \cap V_{\mathrm{C} 2 \mathrm{G}}=\left\{y, v_{1}\right\}$. Applying the same logic as above, for any TaRDiS $S$ at least $\left|V_{\mathrm{C} 2 \mathrm{G}}\right|-$ $2=2 \alpha+2$ vertices in $V_{\mathrm{C} 2 \mathrm{G}}$ are reached by vertices of $S \cap V_{\mathrm{C} 2 \mathrm{G}}$. Hence $\left|S \cap V_{\mathrm{C} 2 \mathrm{G}}\right| \geq\left\lceil\frac{2 \alpha+2}{6}\right\rceil=\frac{\alpha+4}{3}$. Now $\left\{y, u_{2}, v_{5}, u_{8}, v_{11}, \ldots u_{\alpha}\right\}$ is of size $\frac{\alpha+4}{3}$ exactly and reaches every vertex in $V_{\mathrm{C} 2 \mathrm{G}}$ and additionally reaches $x$ at time 2. Conversely if $\lambda(A) \neq 2$ then covering $V_{\mathrm{C} 2 \mathrm{G}}$ requires at most $\frac{\alpha+6}{4}$ vertices, for example $\left\{y, u_{2}, u_{6}, u_{10}, \ldots, u_{\alpha}\right\}$.

Lemma 23. For any happy function $\lambda$ which is not 2-gadget-respecting, $(G, \lambda)$ admits a TaRDiS of size less than $k-1$.

Proof. Recall that $n=n_{Y}+n_{X}$ and $n>m$. Observe that the number of vertices in $G$ which are not inside a U2G or C2G is exactly:

$$
\begin{aligned}
& \#_{\mathrm{U} 3 \mathrm{G}} \cdot(4 \beta+2)+\#_{\mathrm{U} 1 \mathrm{G}}+\#_{\mathrm{lit}}+8 n_{Y}+13 n_{X}+6 m \\
= & \left(2 n_{Y}+6 m\right)(4 \beta+2)+6 m+6 m+12 n_{Y}+13 n_{X}+16 m \\
\leq & 32 n \beta+31 n .
\end{aligned}
$$

which is strictly smaller than $\frac{\alpha}{12} \approx 50 n \beta+50 n$. Even if there exists some $\lambda_{\perp}$ such that every vertex in $G$ not belonging to a U 2 G or C 2 G is necessarily included in every TaRDiS of $\left(G, \lambda_{\perp}\right)$, any such TaRDiS would still have size less than $\#_{\mathrm{U} 2 \mathrm{G}} \cdot \frac{\alpha+1}{3}+\#_{\mathrm{C} 2 \mathrm{G}} \cdot \frac{\alpha+4}{3}$ and hence less than $k-1$.

Lemma 24. For any 2-gadget respecting happy temporal assignment $\lambda$ on any U3G gadget incident to vertex $x$ :

- If $\lambda\left(x, y_{1}\right)=3$ then $\beta$ vertices from $V_{U 3 G}$ are needed to temporally dominate $V_{U 3 G} \backslash\left\{y_{1}\right\}$, and no choice of $\beta$ vertices temporally dominates $V_{U 3 G} \cup\{x\}$.
- If $\lambda\left(x, y_{1}\right) \neq 3$, it is possible to temporally dominate $V_{U S G} \backslash\left\{y_{1}\right\}$ with exactly $\frac{\beta}{2}+1$ vertices.

Recall $V_{U 3 G}=\left\{a_{1}, \ldots, a_{\beta-1}, b_{1}, \ldots b_{\beta-1}, y_{1}, y_{2}, y_{3}, u_{1}, \ldots, u_{\beta}, v_{1}, \ldots, v_{\beta}, w\right\}$; the set of vertices in the red (solid) box in Figure 9.
Proof. There are only two possible 2-gadget-respecting assignments for the edges of $V_{\mathrm{U} 3 \mathrm{G}}$. In any case, for all $i \in[\beta], \lambda\left(u_{i}, v_{i}\right)=\lambda\left(x, y_{1}\right)$. We denote by $\lambda_{3}$ the assignment where $\lambda(x, y)=3$ and $\lambda_{1}$ the temporal assignment where $\lambda(x, y)=1$.

It can be manually verified that, under $\lambda_{3}, \beta$ vertices are necessary to temporally dominate $V_{\mathrm{U} 3 \mathrm{G}} \backslash\left\{y_{1}\right\}$. This is also sufficient: the set $\left\{v_{1}, u_{2}, v_{3}, u_{4}, \ldots, u_{\beta}\right\}$ is one possibility. On the other hand, under $\lambda_{1}$ there exists a set of size $\frac{\beta}{2}+1$ which temporally dominates $V_{\mathrm{U} 3 \mathrm{G}} \backslash\left\{y_{1}\right\}$, namely $\left\{u_{1}, u_{2}, u_{4}, \ldots, u_{\beta}\right\}$.
Lemma 25. For any 2-gadget-respecting $\lambda$, if for any edge e incident to a U3G $\lambda(e) \neq 3$ then $(G, \lambda)$ admits a TaRDiS of size less than $k-1$.

Proof. Observe that the number of vertices in $G$ which are not inside a U2G, C2G or U3G is exactly:

$$
\begin{aligned}
& \#_{\mathrm{U} 1 \mathrm{G}}+\#_{\mathrm{lit}}+8 n_{Y}+13 n_{X}+6 m \\
= & 6 m+6 m+8 n_{Y}+13 n_{X}+6 m \\
= & 18 m+8 n_{Y}+13 n_{X} \\
\leq & 31 n
\end{aligned}
$$

which is strictly smaller than $\frac{\beta}{2}$. Even if there exists some 2-gadget-respecting $\lambda_{\perp}$ such that every vertex in $G$ not belonging to a U2G, C2G, or U3G is necessarily included in every TaRDiS of ( $G, \lambda_{\perp}$ ), any such TaRDiS would still have cardinality less than $\#_{\mathrm{U} 2 \mathrm{G}} \cdot \frac{\alpha+4}{3}+\#_{\mathrm{C} 2 \mathrm{G}} \cdot \frac{\alpha+4}{3}+\#_{\mathrm{U} 3 \mathrm{G}} \cdot(\beta)$ and hence less than $k-1$. Recall we assigned $k=\#_{\mathrm{U} 2 \mathrm{G}} \cdot\left(\frac{\alpha+1}{3}\right)+\#_{\mathrm{C} 2 \mathrm{G}} \cdot\left(\frac{\alpha+4}{3}\right)+\#_{\mathrm{U} 3 \mathrm{G}} \cdot(\beta)+2 m+3 n_{X}+3 n_{Y}+1$, and that by our choice of $\beta$ we have $\frac{\beta}{2}>2 m+3 n_{X}+3 n_{Y}+1$.
Lemma 26. Any happy temporal assignment $\lambda$ such that every $\operatorname{TaRDiS}$ on $(G, \lambda)$ has size at least $k-1$ is nice.
Proof. By Lemmas 23 and 25, any happy temporal assignment $\lambda$ such that every TaRDiS on $(G, \lambda)$ has size at least $k-1$ is 2 -gadget and 3 -gadget-respecting. Observe that if $\lambda$ is happy, 2- and 3 -gadget-respecting, it must also be 1-gadget-respecting by the edge coloring constraint. Hence $\lambda$ is nice.

We say that a nice function $\lambda$ encodes a truth assignment to the variables of $X$, and define this assignment as follows: let variable $x_{i} \in X$ be set to True if $\lambda\left(x_{i}^{1}, x_{i}^{2}\right)=1$ and False otherwise. In Figure 12, True corresponds to the case where $A=1$.
Lemma 27. A nice function $\lambda$ encodes an assignment to $X$ under which $\Phi(X, Y)$ is satisfiable if and only if $(G, \lambda)$ admits a TaRDiS of size $k-1$.

Proof. We first prove the forward direction. That is, if $\lambda$ encodes an assignment $\mathcal{X}$ to $X$ under which $\Phi(X, Y)$ is satisfiable then $(G, \lambda)$ admits a TaRDiS of size at most $k-1$.

In this case, there is an assignment $\mathcal{Y}$ to $Y$ such that $\Phi(\mathcal{X}, \mathcal{Y})$ evaluates to True. We have, by Lemmas 21,22 and 24 , that exactly $\#_{\mathrm{U} 2 \mathrm{G}} \cdot \frac{\alpha+1}{3}+\#_{\mathrm{C} 2 \mathrm{G}} \cdot \frac{\alpha+4}{3}+\#_{\mathrm{U} 3 \mathrm{G}} \cdot \beta$ vertices are necessary and sufficient to cover all the vertices in U2G, U3G, C2G and U1G gadgets, and these vertices also reach every vertex adjacent to a C 2 G at time 2 exactly. Let $S$ include exactly those vertices.

For each $Y$-variable gadget we add the vertices $v_{i}^{1}$ and $T_{i}^{1}$ to $S$, if $y_{i}$ is True in $\mathcal{Y}$, and vertices $v_{i}^{1}$ and $F_{i}^{1}$ are added otherwise. Further, let $S$ include vertices $x_{i}^{4}$ and $x_{i}^{10}$ from each $X$-variable gadget. Note that irrespective of how that gadget is labeled under $\lambda$, this choice of vertices is sufficient to reach all vertices not already covered from some vertex in $S$ within a C2G. Additionally, note that now every literal vertex corresponding to a literal set to True under the combined assignment to $X \cup Y$ is reached from $S$. Since this is a satisfying assignment, every set $Q_{j}$ of clause gadget vertices as shown in Figure 11 must be incident to at least one literal vertex which is already reached from $S$. By symmetry we may assume this literal vertex is $u$ (otherwise, cycle the labels $q_{i}^{1}, \ldots, q_{i}^{6}$ such that it is). Then let $S$ include vertices $q_{j}^{3}$ and $q_{j}^{5}$. Observe that $S$ is a TaRDiS of size precisely $k-1$; all gadget vertices, literal vertices, clause vertices and variable vertices are reachable from $S$.

We now prove the other direction. Namely, if $(G, \lambda)$ admits a TaRDiS of size at most $k-1$ then $\lambda$ encodes an assignment to $X$ under which $\Phi(X, Y)$ is satisfiable. Recall that, by Lemma 26, $\lambda$ must be nice. Let $S$ be a minimum TaRDiS of size at most $k-1$ in $(G, \lambda)$. We may assume, by Lemma 2 , that $S$ is canonical.

Since $S$ is minimum, we have, by Lemmas 21, 22 and 24 , that exactly $\#_{\mathrm{U} 2 \mathrm{G}} \cdot \frac{\alpha+1}{3}+\#_{\mathrm{C} 2 \mathrm{G}} \cdot \frac{\alpha+4}{3}+\#_{\mathrm{U} 3 \mathrm{G}} \cdot \beta$ vertices in $S$ cover all the vertices in U2G, U3G, C2G and U1G gadgets. Then at most $2 m+3 n_{X}+3 n_{Y}$ vertices of $S$ are not among these, and must be sufficient to cover the remaining vertices of $G$.

We define from $S$ a truth assignment to $Y$ as follows: if vertex $T_{i}^{1}$ or $T_{i}^{2}$ is in $S$, then we assign $y_{i}$ to be True, and we assign it to be False otherwise. We argue that this assignment together with $\mathcal{X}$ necessarily satisfies $\Phi(X, Y)$. Note that, since $S$ is canonical, it must contain at least 2 vertices from every clause gadget, 2 vertices from every $Y$-variable gadget, and 2 vertices from every $X$-variable gadget. All these bounds must be tight, else $S$ has cardinality more than $k-1$.

Suppose for contradiction that $\Phi(X, Y)$ is not satisfied by the assignment. Then some clause $c_{j}$ contains only False literals under the assignment, and at least 3 vertices in $Q_{j}$ must be in $S$, contradicting that at most 2 vertices from any clause gadget may be in the TaRDiS. Hence there is an assignment to $Y$ satisfying $\Phi(X, Y)$.
Lemma 28. co-RPS is polynomial-time reducible to Happy MaxMinTaRDiS.
Proof. The construction above can be achieved in polynomial time. We have shown:

- The constructed MaxMinTaRDiS instance $(G, k, \tau=3)$ is a yes-instance if and only if there is some nice $\lambda$ such that every TaRDiS for $(G, \lambda)$ has size at least $k-1$.
- Every nice $\lambda$ encodes a truth assignment to $X$.
- A nice function $\lambda$ encodes an assignment to $X$ under which $\Phi(X, Y)$ is satisfiable if and only if $(G, \lambda)$ admits a TaRDiS of size $k-1$.

That is, $(G, k, \tau=3)$ is a yes-instance of MaxMinTARDIS if and only if $(X, Y, \Phi)$ is a yes-instance of co-RPS.

Theorem 4. HAPPY MAXMInTARDIS is $\Sigma_{2}^{P}$-complete even restricted to inputs where $\tau=3$ and the input graph $G$ is planar.

Proof. By Lemma 28 we have that co-RPS is polynomially reducible to Happy MaxMinTaRDiS. coRPS is $\Sigma_{2}^{P}$-complete by Lemma 20 . The reduction above preserves planarity of $\Phi$. We have containment of MaxMinTaRDis in $\Sigma_{2}^{P}$ from Lemma 15.

### 3.3 NP-completeness of Nonstrict MaxMinTaRDiS with lifetime 2

Here we consider the restriction of Nonstrict MaxMinTaRDiS to instances with lifetime 2. We show the problem to be equivalent to the Distance-3 Independent Set (D3IS) decision problem. We say that two problems $X$ and $Y$ are equivalent if they have the same language - that is, an instance $I$ is a yes-instance of $X$ if and only if the same instance $I$ is a yes-instance of $Y$. Where $X$ has a language consisting of triples $(G, k, \tau)$ and $Y$ has a language of tuples $(G, k)$, we may say that $Y$ is equivalent to $X$ with $\tau$ fixed to some value.

We define the distance $d(u, v)$ between two vertices $u, v$ in a temporal graph to be the number of edges in the shortest path between them in the footprint of the graph.
Definition 5. A distance-3 independent set (D3IS) of a static graph $H$ is a set $S \subseteq V(H)$ such that for all distinct $u, v \in S, d(u, v) \geq 3$.

The decision problem D3IS is defined as follows.
Distance-3 Independent Set (D3IS)
Input: A static graph $H=(V, E)$ and an integer $k$.
Question: Is there a set $S \subseteq V(H)$ of cardinality $k$ that is a distance- 3 independent set?
We aim to show that a static graph $H$ and integer $k$ are a yes-instance of Nonstrict MaxMinTaRDiS with lifetime 2 if and only if the same graph $H$ and integer $k$ are a yes-instance of D3IS.

We begin by showing that existence of a maximal D3IS of size $k$ in a graph $H$ implies that we can find a temporal assignment $\lambda: E(H) \rightarrow\{1,2\}$ such that a minimum TaRDiS in $(H, \lambda)$ is of cardinality $k$. Given such a D3IS $S$ of $H$, we assign $\lambda(u, v)=1$, when $u \in S$ or $v \in S$, and $\lambda(u, v)=2$, otherwise.
Lemma 29. Let $S$ be a maximal D3IS of a static graph $H$ and $\lambda$ be a temporal assignment of $H$ where $\lambda(u, v)=1$ when $u \in S$ or $v \in S$, and $\lambda(u, v)=2$ otherwise. Then $S$ is a minimum TaRDiS of $(H, \lambda)$.

Proof. We first show that $S$ is a TaRDiS. We assume without loss of generality that $H$ is a single connected component. Suppose for contradiction some vertex $u$ is not reachable from any vertex in $S$. Note that every
vertex in $S$ trivially reaches its neighbours. So, by construction of $\lambda, u$ is incident only to edges at time 2 . Since we have assumed that $H$ consists of a single connected component, there must be a static path in $H$ from each vertex in $S$ to $u$. Let $z$ be the closest vertex in $S$ to $u$. Then the shortest path from $z$ to $u$ must have length 2 and consist of an edge assigned time 1 followed by an edge assigned time 2 . Else, $S$ is not maximal. Hence, the shortest path from $z$ to $u$ is a nonstrict temporal path and $u \in R_{z}$, contradicting our assumption. Thus $S$ is a TaRDiS of the constructed instance.

We now show minimality of $S$. By construction of $\lambda$, every vertex $v \in S$ is temporally reachable only from its closed neighbourhood $N[v]$. No temporal path originating outside $N[v]$ can include any edge incident to $v$ since any such path must contain an edge assigned time 2 before the final edge which is assigned time 1 . Since $S$ forms a D3IS, $N[u] \cap N[v]=\emptyset$ for all $u, v \in S$. Therefore, for $(H, \lambda)$ to be temporal reachability dominated, there must at least be a vertex from the neighbourhood of each vertex in $S$. These are disjoint sets, so any TaRDiS must have cardinality at least $k$. Hence $S$ is a minimum TaRDiS of $(H, \lambda)$.
Definition 6 (Sole Reachability Set). We define the sole reachability set of a vertex $v$ in a TaRDiS $S$ as the set $S R(\mathcal{G}, S, v)=R_{v}(\mathcal{G}) \backslash\left(\cup_{u \in S \backslash\{v\}} R_{u}(\mathcal{G})\right)$. Equivalently, it is the set of vertices reachable from $v$ and not any other vertex in $S$.

When $\mathcal{G}$ is clear from context, we write $S R(S, v)$ for $S R(S, \mathcal{G}, v)$. Note that, in a minimum TaRDiS, every vertex has a non-empty sole reachability set.
Definition 7. We call a TaRDiS S on a temporal graph $\mathcal{G}$ independent if and only if every vertex in $S$ is in its own sole reachability set under $S$.
Lemma 30. If a temporal graph $\mathcal{G}$ admits an independent nonstrict TaRDiS $S$, then $S$ is a D3IS in the footprint graph $\mathcal{G}_{\downarrow}$.
Proof. Consider two vertices $u, v \in S$ at distance $d(u, v)$ from one another in $\mathcal{G}_{\downarrow}$. If $u$ and $v$ are adjacent, then $u$ and $v$ reach each other, which would contradict independence of $S$. If $d(u, v)=2$ then there is at least one vertex $w \in N[u] \cap N[v]$, and either $\lambda(u, w) \leq \lambda(w, v)$ or $\lambda(u, w)>\lambda(w, v)$. So one of $u$ and $v$ must reach the other. Hence any two vertices in $S$ must be distance at least 3 from one another, the definition of a D3IS.

Lemma 31. If a temporal graph $\mathcal{G}$ with lifetime 2 has a minimum nonstrict TaRDiS of cardinality $k$, then $\mathcal{G}_{\downarrow}$ admits a D3IS of size $k$.
Proof. Our proof is constructive; given a minimum TaRDiS $S$, we show existence of an independent minimum TaRDiS $S^{*}$ of equal cardinality and then apply Lemma 30.

First, we justify some simplifying assumptions about $\mathcal{G}$. Since TARDIS can be computed independently in disconnected components of $\mathcal{G}_{\downarrow}$, we will assume $\mathcal{G}_{\downarrow}$ is connected. Further, if $E_{1}(\mathcal{G})=\emptyset$ or $E_{2}(\mathcal{G})=\emptyset$ we may choose any single vertex from $\mathcal{G}$ to be a minimum TaRDiS which is also independent; hence we assume $E_{1}(\mathcal{G}) \neq \emptyset$ and $E_{2}(\mathcal{G}) \neq \emptyset$. We also recall that $S R(S, x) \neq \emptyset$ for all $x \in S$ by minimality of $S$.

We construct $S^{*}$ by replacing every vertex $x$ in $S$ such that $x \notin S R(S, x)$ with some vertex $x^{*}$ with the property that $R_{x^{*}}=R_{x}$ and $x^{*} \in S R(S, x)$, as follows. Let $y \neq x$ be a vertex in $S$ such that $y$ reaches $x$. The path from $y$ to $x$ cannot arrive at time 1, else $S$ is not minimal as $R_{x}=R_{y}$ and $S \backslash\{x\}$ is a TaRDiS. Choose $x^{*}$ to be the closest vertex to $x$ in $S R(S, x)$. We know such a vertex exists by minimality of $S$. We claim that the path from $x$ to $x^{*}$ arrives at time 1 and so $R_{x}=R_{x^{*}}$. To see this, suppose otherwise. The path must begin with at least one edge at time 1 , otherwise $x^{*}$ would be reachable from $y$. If the path arrives at time 2, then the last vertex on the path reached at time 1 is closer to $x$ than $x^{*}$, and is in $S R(S, x)$. Else, some other vertex in $S$ reaches $x^{*}$.

This concludes our construction of $S^{*}$ as an independent minimum TaRDiS. By Lemma 30, $S^{*}$ is also a D3IS.

Combining Lemmas 29 and 31 gives us the following theorem.
Theorem 5. Nonstrict MaxMinTaRDiS with lifetime $\tau=2$ is equivalent to Distance-3 Independent Set.

Interestingly, the same does not hold for $\tau \geq 3$. A counterexample is shown in Figure 4c, where the minimum TaRDiS is larger than the maximum D3IS of the footprint. The Petersen graph also gives us a 3 -regular counterexample.

Eto, Guo and Miyano [16] show that D3IS is NP-complete even on planar, bipartite graphs with maximum degree 3. They also show D3IS to be $W[1]$-hard on chordal graphs with respect to the size of the distance-3 independent set. D3IS is also shown to be APX-hard on $r$-regular graphs for all integers $r \geq 3$ and admit a PTAS on planar graphs by Eto, Ito, Liu, and Miyano [51]. Agnarsson, Damaschke and Halldórsson [52] show that D3IS is tractable on interval graphs, trapezoid graphs and circular arc graphs.

Corollary 7. Nonstrict MaxMinTaRDiS restricted to inputs with $\tau=2$ remains NP-complete even when restricted to planar, bipartite graphs with maximum degree 3. Furthermore, the problem is $W$ [1]-hard with respect to the parameter $k$, the size of a minimum TaRDiS, and is APX-hard on r-regular graphs for all integers $r \geq 3$.
Corollary 8. Nonstrict MaxMinTaRDiS restricted to inputs with $\tau=2$ is in NP, admits a PTAS on planar graphs and is tractable on interval graphs, trapezoid graphs and circular arc graphs.

## 4 Parameterized complexity results for TARDIS

In Section 2, we showed that the variants of TARDIS are NP-complete and W[2]-hard with respect to $k$. This rules out $k$ as a candidate parameter for an fpt algorithm. We begin by showing inclusion of TaRDiS in FPT with respect to locally earliest edges and tractability when the input is heavily restricted. We then give an algorithm which solves each variant of TARDIS on a nice tree decomposition of the footprint graph. This generalises the tree algorithm given in Section 2.

A problem $\Pi$ is said to be fixed-parameter tractable (fpt) with respect to some parameter $k$ if there is an algorithm solving $\Pi$ in time $f(k) \cdot \operatorname{poly}(n)$ (where $n$ is the size of the instance of $\Pi$ ). The complexity class FPT consists of all problem-parameter pairs admitting an fpt algorithm.

### 4.1 FPT results with a restricted temporal assignment

We begin with two FPT results that require the temporal assignment to be restricted in some way to recover tractability of TARDIS. In the first, we consider Happy TaRDiS and Nonstrict TaRDiS when parameterised by the number of locally earliest edges. Following that, we consider Strict TaRDiS when each component of the input is restricted in some way.
Lemma 32. Happy TaRDiS and Nonstrict TaRDiS are in FPT with respect to the number of locally earliest edges and weakly locally earliest edges, respectively.
Proof. It suffices to observe that any instance with $t$ (weakly) locally earliest edges trivially admits a TaRDiS of cardinality $t$ (Corollary 1). If $k \geq t$, we must have a yes-instance. Otherwise there are $\binom{t}{k}$ possibilities for a (weakly) canonical TaRDiS, which can be checked by brute force in polynomial time when $t$ is bounded.

Lemma 33. Strict TARDIS is in FPT parameterized by maximum degree in $\mathcal{G}_{\downarrow}$, lifetime and $k$, and Happy TaRDIS is in FPT with parameters lifetime and $k$.
Proof. Any decidable language consisting of words of length at most some constant $c$ can be decided in constant time, hence is in $P$. Any strict temporal path in a temporal graph $\mathcal{G}$ has length at most $\tau$. This entails that, for all $v,\left|R_{v}^{<}\right| \leq 2 \Delta^{\tau}$ where $\Delta$ is the maximum degree of $\mathcal{G}_{\downarrow}$. Hence any strict TaRDiS $S$ in $\mathcal{G}$ must satisfy $|S| \geq \frac{V(\mathcal{G})}{2 \Delta^{\tau}}$. Therefore, no instance $(\mathcal{G}, k)$ satisfying $|V(\mathcal{G})| \geq 2 k \Delta^{\tau}$ can be in the language Strict TaRDiS. Note that, in a happy temporal graph, we can can apply the property that $\tau \geq \Delta$. Therefore, both problems restricted as above have constantly many yes-instances, each of size bounded by a constant.

### 4.2 Preliminaries: treewidth and tree decompositions

We refer the interested reader to Chapter 10 of Niedermeier's Invitation to Fixed-Parameter Algorithms [53] for a fuller introduction. Definitions and results in this subsection are taken or adapted from that work.
Definition 8 (Tree Decomposition, Treewidth). We say a pair $(T, B)$ is a tree decomposition of $G$ if $T$ is a tree and $B=\{B(s): s \in V(T)\}$ is a collection of subsets of $V(G)$, called bags, satisfying:

1. $V(G)=\cup_{s \in V(T)} B(s)$.
2. $\forall(u, v) \in E(G): \exists s \in V(T):\{u, v\} \in B(s)$. That is, for each edge in the graph, there is at least one bag containing both of its endpoints.
3. $\forall v \in V(G): T[\{s: v \in B(s)\}]$ is connected; for each vertex, the subgraph obtained by deleting every node not containing $v$ in its bag from $T$ is connected.

The width of a tree decomposition is defined to be $\max \{|B(s)|: s \in V(T)\}-1$. The treewidth of a graph $G$ is the minimum $\omega$ such that $G$ has a tree decomposition of with $\omega$.

For a given tree decomposition $(T, B)$ of graph $G$, we denote by $V_{s} \subseteq V(G)$ the set of vertices in $G$ that occur in bags of the subtree of $T$ rooted at $s$. It is a well-known result by Bodlaender that finding a tree decomposition of width at most $\omega$, if one exists, is in FPT with parameter $\omega$.


Fig. 14: Graph $P_{3}$ and four of its tree decompositions. From left to right: $P_{3}$; the trivial decomposition of width 2 ; a decomposition of width 1 ; two possible nice tree decompositions of width 1 . Note that every graph admits the trivial tree decomposition of width $V(G)-1$, where there is only one bag $B(s)=V(G)$.

Theorem 6 (Bodlaender [54]). For all fixed constants $\omega \in \mathbb{N}$, there exists a linear time algorithm that tests whether a given graph $G=(V, E)$ has treewidth at most $\omega$, and if so outputs a tree decomposition of $G$ with treewidth at most $\omega$.

For ease, we describe our TARDiS algorithm on a tree decomposition with additional structural properties.
Definition 9 (Nice Tree Decomposition, Join/Introduce/Forget/Leaf Node). A tree decomposition ( $T, B$ ) is called a nice tree decomposition if:

- every node of $T$ has at most two children;
- if a node s has two children $s_{l}$ and $s_{r}$, then $B(s)=B\left(s_{l}\right)=B\left(s_{r}\right)$, and we call $s$ a join node;
- if node s has one child $s^{\prime}$ then either:
- $|B(s)|=\left|B\left(s^{\prime}\right)\right|+1$ and $B(s) \supset B\left(s^{\prime}\right)$ (s is an introduce node), or
$-|B(s)|=\left|B\left(s^{\prime}\right)\right|-1$ and $B(s) \subset B\left(s^{\prime}\right)$ (s is a forget node);
- if node s has no children then $B(s)=\emptyset$, and we call $t$ a leaf node.

An example of a graph and different tree decompositions of it can be found in Figure 14. It is a known result that a tree decomposition $(T, B)$ of a graph $G$ of width $\omega$ on $n$ nodes can be efficiently processed to obtain a nice tree decomposition $G$ of width $\omega$. We use this result as given below by Cygan, Fomin, Kowalik, Lokshtanov, Marx, Pilipczuk, Pilipczuk and Saurabh [17].
Lemma 34 (Cygan [17]). If a static graph $H$ admits a tree decomposition of width at most $\omega$, then it also admits a nice tree decomposition of width at most $\omega$. Moreover, given a tree decomposition ( $T, B$ ) of width at most $\omega$, on can compute a nice tree decomposition of $H$ of with at most $\omega$ that has at most $O(\omega V(H))$ nodes in time $O\left(\omega^{2} \max (V(T), V(H))\right.$.

### 4.3 Algorithm for TARDIS parameterized by treewidth and lifetime

The following gives an algorithm for TARDIS given a nice tree decomposition $(T, B)$ of the footprint $\mathcal{G}_{\downarrow}$ of the input temporal graph $\mathcal{G}$. Specifically, it computes the cardinality of a minimum TaRDiS $S$. Note that we use the word "vertex" when referring to the original graph and "node" when discussing the decomposition graph. We use the symbols $\prec$ and $\succ$ as place-holders for strict/nonstrict inequalities. More specifically, in Strict TaRDiS we use $<$ and $>$ in the place of $\prec$ and $\succ$ respectively and for Nonstrict TaRDiS we use $\leq$ and $\geq$. By substituting the correct inequalities, it can be seen that the algorithm described is correct for each of Strict TaRDiS and Nonstrict TaRDiS.

Denote by $G_{s}$ the subgraph $\left(V_{s}, E_{s}\right)$, where $V_{s}$ (respectively $E_{s}$ ) is the set of all vertices (resp. edges) introduced in the subtree rooted at node $s$ of the decomposition tree. Intuitively, the algorithm works from
leaf nodes to the root node and finds, at each node $s$, a partial solution consisting of a minimal TaRDiS $S$ for the subgraph $G_{s}$.

We define the arrival time of a temporal path at a vertex $v$ to be the time of the final time-edge of the path. For the trivial path from a vertex to itself, we say that the time of arrival is 0 . Similarly, we define the departure time of a temporal path as the time of the first time-edge in the path. For example, the temporal path consisting of a single time-edge has the same departure and arrival time. We note that, for a path to be a strict temporal path, the arrival time at a vertex $v$ must be strictly before the departure time from $v$. We allow the arrival and departure times to be equal in nonstrict temporal paths of any length. We refer to a temporal path from a vertex $v$ to a vertex $u$ with earliest arrival time as a foremost temporal path. We call the arrival time of a foremost path the foremost arrival time. A foremost path from a set of vertices $S$ to a single vertex $v$ is a foremost path from a vertex $u$ in $S$ to $v$ such that the arrival time of a path from any other vertex in $S$ to $v$ is the same or later.

### 4.3.1 States

A state of a bag $B(s)$ is a mapping $\psi: B(s) \rightarrow([0, \tau] \cup \perp) \times([0, \tau] \cup \perp)$ where $\psi(v)=\left(t_{a}(v), t_{p}(v)\right)$, such that $t_{a}(v) \geq t_{p}(v)$ if both values are integers and $t_{p}(v)=\perp$ only if $t_{a}(v)=\perp$. Conceptually, we think of $t_{a}(v)$ as the arrival time of some path to $v$ from our partial TaRDiS and $t_{p}(v)$ as the time by which we "promise" arrival of such a path from the eventual full TaRDiS (necessary to reach forgotten vertices). Denote the set of all states of a node $s$ by $\Psi(s)$.

In our definition of a temporal graph, we assume that all edges are active at strictly positive times. Therefore, in both the strict and nonstrict variants of the problem, the earliest time any vertex can be temporally reachable from a vertex in the TaRDiS is 1 , unless it is in the TaRDiS itself. Our intention is that, if $t_{a}(v) \neq \perp$ then $t_{a}(v)$ is exactly the time that $v$ is reached from some TaRDiS vertex. If $t_{a}(v)=0$ for a vertex $v$ in the bag, this corresponds to $v$ being included in the partial TaRDiS $S$. We use the notation $t_{a}^{-1}$, $t_{p}^{-1}$ to denote the preimage of the functions $t_{a}$ and $t_{p}$ respectively. That is, $t_{a}^{-1}(x)$ is the set of all vertices $v$ which are assigned $\left(x, t_{p}(v)\right)$ under a state $\psi$.

We call a state $\psi$ of a bag $B(s)$ consistent if and only if there exists a set of vertices $S \subseteq V\left(G_{s}\right)$ such that

- for all $v \in B(s) \backslash t_{a}^{-1}(\perp)$ the foremost temporal path from some vertex in $S$ arrives at $t_{a}(v)$ exactly;
- for any vertex $w$ in $V\left(G_{s}\right) \backslash B(s)$ which is not reachable from some vertex in $S$, there is a temporal path to $w$ from some vertex $u$ in $B(s)$ with departure time $t$ such that $t \succ t_{p}(u)$.
We call such a set $S$ a set that supports $\psi$ of $s$. Note that, if $t_{a}(v)=\perp$, then it is possible to have a valid state where $v$ is reachable from some vertex in $S$.


### 4.3.2 Signature

For a state $\psi$ of a node $s$, we define the signature $c(s, \psi)$ to be the cardinality of the smallest set $S^{*}$ which supports $\psi$ of $s$. If there is no such $S^{*}$, then we say that $c(s, \psi)=\infty$. We say that such a set $S^{*}$ supports the signature $c(s, \psi)$ if $S^{*}$ supports $\psi$ of $s$ and $\left|S^{*}\right|=c(s, \psi)$. The signature is a data structure with size $O\left(\tau^{2(\omega+1)}\right)$ where $\omega$ is the width of the nice tree decomposition of $\mathcal{G}_{\downarrow}$. Note that two bags may contain the same (possibly empty) set of vertices. Therefore, the signature must be both a function of the state $\psi$ and of the node $s$ which $\psi$ is a state of. Examples of the signatures of different states can be seen in Figure 15.

We use the convention that the root node $r$ is empty, i.e. $B(r)=\emptyset$. Consequently there is only one possible state for the root node, namely the empty function, and hence $\Psi(r)=\{\emptyset\}$. Therefore, we have a yes-instance of TaRDiS if and only if $c(r, \emptyset) \leq k$. Note that $\emptyset$ is always a consistent state for $r$. It is supported by the trivial TaRDiS $S=V(G)$.

We now discuss how we iteratively calculate the signature of a state for each type of node in a nice tree decomposition.

## Leaf nodes

We assume that leaf nodes $l$ are empty, so there is only one trivial state, namely the empty function $\emptyset$. The signature of this state for leaves $l$ is $c(l, \emptyset)=0$.

## Introduce Nodes

Let $s$ be an introduce node with child $s^{\prime}$. Then we must have that $B(s)=B\left(s^{\prime}\right) \cup\{v\}$ for some $v \notin B\left(s^{\prime}\right)$. To describe how to find $c(s, \psi)$ for an introduce node $s$ and state $\psi$, we must define some new notation.


Fig. 15: Five example bags, labelled with states and signatures. In the bottom right example, the infinite signature reflects the fact there exists no TaRDiS reaching $v$ at time 3 exactly.


Fig. 16: A valid state of an introduce node and a state of its child which are supported by the same partial TaRDiS. Here $x$ is a forgotten vertex.

Let $\left.\psi\right|_{B\left(s^{\prime}\right)}$ be the restriction of some state $\psi$ of $s$ to the bag $B\left(s^{\prime}\right)$. For a state $\psi$ and functions $g, f$, let the state $\psi^{t_{a}(A) \rightarrow g, t_{p}(B) \rightarrow h}$ be defined

$$
\psi^{t_{a}(A) \rightarrow g, t_{p}(B) \rightarrow h}(x):= \begin{cases}(g(x), h(x)) & \text { if } x \in A \cap B \\ \left(g(x), t_{p}(x)\right) & \text { if } x \in A \backslash B \\ \left(t_{a}(x), h(x)\right) & \text { if } x \in B \backslash A \\ \left(t_{a}(x), t_{p}(x)\right), & \text { otherwise. }\end{cases}
$$

For an introduce node $s$ with introduced vertex $v$ and state $\psi$ of $s$, we define $a_{v}$ to be the time of the earliest time-edge in $\left\{((v, u), t) \mid u \in B(s)\right.$ and $t_{a}(u) \prec t$ under $\left.\psi\right\}$. If no such edge exists, let $a_{v}=\infty$. Notionally, this is a time before which $v$ cannot be reached from a vertex in $S$ unless $v$ is itself in $S$. We define $R_{v}^{t}$ to be the set of vertices that are temporally reachable from $v$ by temporal paths which depart at a time $t^{\prime} \succ t$. We use the convention that $R_{w}^{\infty}=R_{w}^{\perp}=\emptyset$ for all vertices $w$. For a vertex $u$ in $R_{v}^{t}$, define foremost ${ }_{v}^{t}(u)$ to be the arrival time of a foremost path from $v$ to $u$ which departs at a time $t^{\prime} \succ t$. An example of a valid state of an introduce node can be seen in Figure 16.

Lemma 35. Let $\psi$ be a state of an introduce node $s$ with child $s^{\prime}$ where $v$ is the vertex introduced. Define $Z$ to be the set $\left(R_{v}^{t_{a}(v)}(\mathcal{G}) \cap B(s)\right) \backslash\left\{u:\right.$ foremost $\left._{v}^{t_{a}(v)}(u)>t_{a}(u)\right\}$, and define the function $f(w):=\min \left\{t_{p}(w)\right.$, foremost $\left._{v}^{t_{p}(v)}(w)\right\}$. Then

1. if, for any $u \in R_{v}^{a_{v}}(\mathcal{G}) \cap B(s), t_{a}(u) \neq \perp$ and foremost ${ }_{v}^{a_{v}}(u)<t_{a}(u)$, then $c(s, \psi)=\infty$ (a foremost path arrives before it is prescribed to);
2. else, if $t_{a}(v)=0, c(s, \psi)=1+c\left(s^{\prime},\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}\right.$ ) (the introduced vertex is added to the partial TaRDiS);
3. else, if we allow nonstrict paths (i.e. if $1 \prec 1$ ), $a_{v}=t_{a}(v)$ and there exists a nonempty set $W$ of neighbours of $v$ where, for all $w \in W, t_{a}(w)=t_{a}(v)=\lambda(v w)$ then $c(s, \psi)=$ $\min _{w \in W}\left\{c\left(s^{\prime},\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}\left(Z^{\prime}\right) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}\right)\right\}$ for $Z^{\prime}=Z \backslash\{w\}$ (a pair of neighbours each rely on the other to be reached on time);
4. else, if $t_{a}(v) \in\left\{a_{v}, \perp\right\}$ then $c(s, \psi)=c\left(s^{\prime},\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}\right)$ (the introduced vertex could result in a foremost path with an earlier arrival time);
5. else, $c(s, \psi)=\infty$ (we have an inconsistent state).

Proof. We begin by noting that for any set $S$ supporting $\psi$ of $s$ if the foremost arrival time to $v$ from a vertex in $S$ is $t_{a}(v)$, then foremost ${ }_{v}^{t_{a}(v)}(u)<t_{a}(u)$ contradicts that $t_{a}(u)$ gives the earliest time of arrival from a vertex in $S$ to $u$. Otherwise, we could append the temporal path to $v$ from $S$ with the path from $v$ to $u$ to find an temporal path with an earlier arrival time. Hence, in case (1) of the Lemma, $\psi$ is an inconsistent state.

We have characterised the value of $c(s, \psi)$ based on the value of $t_{a}(v)$ under $\psi$. This gives us four further cases $((2)-(5))$ to consider, accounting for all possible values of $t_{a}(v)$. We show that our calculation of the signature $c(s, \psi)$ is correct for each possible value of $t_{a}(v)$ for the introduced vertex $v$. Let $S$ be a set that supports the signature of the state $\psi$ of node $s$. That is, it is a set of minimal cardinality such that $t_{a}^{-1}(0)=S \cap B(s)$; all vertices $w$ in $B(s) \backslash t_{a}^{-1}(\perp)$ are temporally reachable from $S$ by a foremost path arriving at $t_{a}(w)$; and each vertex in $G_{s} \backslash B(s)$ not reachable from $S$ is temporally reachable from a vertex $u$ in $B(s)$ departing at some time $t \succ t_{p}(u)$. Note that, for a consistent state, if there is a forgotten vertex reachable from a vertex $w$ in $B(s)$ by a path departing at time $t \succ t_{p}(w)$, the change in $t_{p}$ values ensures that this remains true in the child states over which we take the minimum value.

In case $(2), t_{a}(v)=0$. Additionally, for all $u \in R_{v}^{t}, t_{a}(u) \leq$ foremost $_{v}^{t_{a}(v)}(u)$ or $t_{a}(u)=\perp$, since otherwise we would have case (1). Hence for any set $S$ supporting $\psi, v \in S$ and all vertices in $R_{v}^{0}$ are reached by time foremost ${ }_{v}^{0}(u)$. We claim that a set $S$ supports $\psi$ if and only if $S \backslash\{v\}$ supports the state $\psi^{\prime}=\left(t_{a}^{\prime}, t_{p}^{\prime}\right):=$ $\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow+t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}$ of $s^{\prime}$. That is, every vertex $w$ in $B(s) \backslash R_{v}^{0}$ is temporally reached by a vertex in $S \backslash\{v\}$ at $t_{a}^{\prime}(w)$ and every vertex in $V\left(G_{s}\right) \backslash B(s)$ that is not temporally reachable from $S \backslash\{v\}$ is temporally reachable from a vertex $y$ in $B\left(s^{\prime}\right)$ by a path departing at time $t^{\prime} \succ t_{p}^{\prime}(y)$. It is clear that, if states $\psi$ and $\psi^{\prime}$ are valid, the former statement is true. This is because the value of $t_{a}$ for a vertex in $B(s) \backslash R_{v}^{0}$ is the same under both states.

To verify the latter statement, note that if a vertex $y^{\prime}$ in $V\left(G_{s}\right) \backslash B(s)$ is reachable from a vertex $y$ in $B\left(s^{\prime}\right)$ by a path departing at $t^{\prime} \succ t_{p}^{\prime}(y)$ under $\psi^{\prime}$, then the same must be true under $\psi$. By construction of a nice tree decomposition, there are no vertices in $V\left(G_{s}\right) \backslash B(s)$ that are adjacent to $v$. Therefore, any path from $v$ to a vertex in $V\left(G_{s}\right) \backslash B(s)$ must traverse another vertex $x$ in $B(s)$ and depart $x$ at time $t^{\prime \prime} \succ t_{p}^{\prime}(x)$ under $\psi^{\prime}$. Thus, if a vertex $y^{\prime}$ in $V\left(G_{s}\right) \backslash B(s)$ is reachable from a vertex $y$ in $B\left(s^{\prime}\right)$ by a path departing at $t^{\prime} \succ t_{p}^{\prime}(y)$ under $\psi$, then $y^{\prime}$ must be reachable from a (possibly different) vertex $y^{\prime \prime}$ under $\psi^{\prime}$. Therefore, a set $S$ supports $\psi$ if and only if $S \backslash\{v\}$ supports the state $\psi^{\prime}=\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}$ of $s^{\prime}$ and our calculation of $c(s, \psi)$ is correct.

Intuitively, cases (3) and (4) deal with the cases where the introduction of $v$ could change the earliest arrival time of a path from a vertex in $S$ to a vertex in $B(s) \backslash\{v\}$. In case (3) we deal with the possibility of nonstrict temporal paths when $t_{a}(w)=$ foremost $_{v}^{t_{a}(v)}(w)=t_{a}(v)$ for some vertex $w$. In this case, a child state wherein $t_{a}(w)=\perp$ for every such $w$ may be supported by a set $S$ which does not support $\psi$. For this reason, we take the minimum over the signatures of states where the $t_{a}$ value of each such neighbour is unchanged. A set $S$ supports $\psi$ if and only if there exists a $w \in W$ such that $S$ supports $\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}\left(Z^{\prime}\right) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}$. The forward implication is clear from our description. Now suppose, for contradiction, that $S$ supports
 time $t_{a}(u)$ or there is a forgotten vertex which is not reachable from $S$ or a vertex $w$ in $B(s)$ departing at $t^{\prime} \succ t_{p}(w)$. Neither case is possible, and thus we have a contradiction.


Fig. 17: A forget node with a valid state and its extension to a state of its child supported by the same partial TaRDiS.

In case (4), the above does not apply. This is either because we are in the strict setting or because there is no such nonempty set $W$, and $t_{a}(v) \in\left\{a_{v}, \perp\right\}$ under $\psi$. Then, for all vertices $w$ in $Z=\left(R_{v}^{a_{v}}(\mathcal{G}) \cap B(s)\right) \backslash\{u$ : foremost $\left.{ }_{v}^{t_{a}(v)}(u)>t_{a}(u)\right\}$, there exists a foremost path from a set $S$ that supports the state which traverses $v$. Therefore, any set $S$ which supports $\psi$ must support a child state where these vertices $w \in Z^{\prime}$ have $t_{a}$ value $\perp$ and $t_{p}$ is updated as mentioned earlier. In addition, it is clear that any set which supports $\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}$ must support $\psi$. Since we have not added any vertices to $S$, the signature of $s$ under $\psi$ must be exactly the signature of $s^{\prime}$ under $\left.\psi\right|_{B\left(s^{\prime}\right)} ^{t_{a}(Z) \rightarrow \perp, t_{p}\left(B\left(s^{\prime}\right)\right) \rightarrow f\left(B\left(s^{\prime}\right)\right)}$.

In case (5) we deal with all remaining possibilities, namely when $t_{a}(v)$ for the introduced vertex $v$ has a nonzero value and $t_{a}(v) \neq a_{v}$. Then $t_{a}(v)$ must be strictly before or after the earliest time-edge $((v, w), t)$ incident to $v$ such that, for the other endpoint $w, t_{a}(w) \prec t$. Therefore, for any set $S$ that supports $\psi$, if the foremost temporal path from a vertex in $S$ to each neighbour $w$ of $v$ arrives at $t_{a}(w)$, the foremost temporal path to $v$ from $S$ must arrive at a time which is not equal to $t_{a}(v)$. This implies that the state is inconsistent and must therefore have an infinite signature.

## Forget Nodes

Let $s$ be a forget node with child $s^{\prime}$ such that $B(s)=B\left(s^{\prime}\right) \backslash\{v\}$. Let $\Psi_{\text {strong }} \subset \Psi\left(s^{\prime}\right)$ be the set of states which extend $\psi$ to $B\left(s^{\prime}\right)$ where $t_{a}(v) \in[0, \tau]$ and $t_{a}(v)=t_{p}(v)$. Intuitively, these are the states where our partial TaRDiS already reaches the forgotten node $v$ by the promised time $t_{p}(v)$, i.e. the promise is strongly satisfied.

Let $t^{\prime}$ be the earliest time such that there exists a vertex $w$ in $B(s)$ and a temporal path in $B\left(s^{\prime}\right)$ from $w$ to $v$ departing at some time $t$ such that $t_{p}(w) \prec t$ under $\psi$ and arriving by $t^{\prime}$. Let $\Psi_{\text {weak }} \subset \Psi\left(s^{\prime}\right)$ be the set of states which extend $\psi$ to $B\left(s^{\prime}\right)$ where either $t^{\prime} \leq t_{p}(v)$ or $t_{p}(v)=\perp$. Intuitively, $\Psi_{\text {weak }}$ is the set of states where the the requirement that a forgotten vertex is reached by a path departing at time $t \succ t_{p}(v)$ is automatically satisfied by a path from $w$ departing at $t \succ t_{p}(w)$ for some $w \neq v$. That is, the promise is weakly satisfied. An example of a valid state of a forget node can be seen in Figure 17.
Lemma 36. Let $s$ be a forget node with state $\psi$ and child $s^{\prime}$ where $v$ is the vertex forgotten at $s$. Then

$$
c(s, \psi)=\min \left(\left\{c\left(s^{\prime}, \psi^{\prime}\right): \psi^{\prime} \in \Psi_{\text {weak }} \cup \Psi_{\text {strong }}\right\} \cup\{\infty\}\right) .
$$

Proof. We begin by showing that if a set $S$ supports any state $\psi^{\prime} \in \Psi_{\text {weak }} \cup \Psi_{\text {strong }}$ of $s^{\prime}$ then the same set $S$ supports $\psi$. This shows us that the signature of $\psi$ is at most what we have calculated. Following this, we show that there is no smaller set that supports $\psi$ and hence that we have calculated the signature correctly.

Recall that set $S$ supports a state $\psi$ of a node $s$ if and only if, under $\psi$ :

- For all vertices $u$ in $B(s) \backslash t_{a}^{-1}(\perp)$, the foremost path from $S$ arrives at time $t_{a}(u)$, and
- for all vertices $x$ in $G_{s} \backslash B(s)$ which are not reachable from $S$, there is a temporal path to $x$ which departs from a vertex $w$ in $B(s)$ at some time $t \succ t_{p}(w)$.
To start, consider the case where $S$ supports $\psi^{\prime} \in \Psi_{\text {weak. }}$. That is, there exists a vertex $w \in B\left(s^{\prime}\right)$ and a temporal path from $w$ to $v$ that departs at some time $t^{\prime} \succ t_{p}(w)$ and arrives by time $t_{p}(v)$. As a result, any forgotten vertex $x$ temporally reachable from $v$ by a path departing at some time $t^{\prime \prime} \succ t_{p}(v)$ must be temporally reachable from $w$ by a path departing at time $t^{\prime}$ by prefixing the path from $v$ to $x$ with the temporal path from $w$ to $v$. If $\psi^{\prime}$ is consistent, then all vertices $y$ in $B\left(s^{\prime}\right) \backslash\{v\}=B(s)$ that are reachable
from a vertex in $S$ are reached by time $t_{a}(y)$ if $t_{a}(y) \neq \perp$. Therefore, the restriction of $\psi^{\prime}$ to $B(s)$ is consistent and supported by $S$.

Now consider the case where $S$ supports a state $\psi^{\prime} \in \Psi_{\text {strong }}$. Recall that $\Psi_{\text {strong }}$ is the set of states of $s^{\prime}$ where $t_{a}(v) \in[0, \tau]$ and $t_{a}(v)=t_{p}(v)$. If such a state $\psi^{\prime}$ is consistent for $s^{\prime}$, then $v$ is reached from a vertex in $S$ by time $t_{a}(v)$. Using that $t_{p}(v)=t_{a}(v)$, we obtain that any forgotten vertex $x$ reachable from $B(s)$ by a path from $v$ departing at some time $t^{\prime} \succ t_{p}(v)$ is reachable from $S$ as well. Namely, we can prefix the temporal path from $v$ to $x$ which departs at time $t^{\prime}$ with the temporal path from $S$ to $v$ which arrives by time $t_{a}(v)=t_{p}(v)$. Hence we obtain that $\psi$ is consistent for $s$ and supported by $S$.

We now suppose for contradiction that our calculation of the signature is incorrect. We have already shown that the signature must be at most $|S|$ for some set $S$ supporting a state in $\Psi_{\text {strong }} \cup \Psi_{\text {weak }}$. Therefore, we suppose that there is a smaller set $S^{*}$ which supports $\psi$ whose intersection with $B(s)$ is exactly the set of vertices $t_{a}^{-1}(0)$ under $\psi$.

Consider the state $\psi^{*}$ extending $\psi$ to $B\left(s^{\prime}\right)$ in which $t_{a}(v)$ is the arrival time of the foremost path from $S^{*}$ to $v$ (or $\perp$ if there is no such path) and $t_{p}(v)=\tau$. Clearly $\psi^{*} \in \Psi_{\text {weak }} \cup \Psi_{\text {strong }}$, and $S^{*}$ supports $\psi^{*}$. This contradicts our assumption that $S^{*}$ was smaller than any set supporting a state in $\Psi_{\text {weak }} \cup \Psi_{\text {strong }}$.

## Join Nodes

Let $s$ be a join node with children $s_{1}$ and $s_{2}$.
Lemma 37. Let $\psi_{1}=\left(t_{a, 1}, t_{p, 1}\right)$ and $\psi_{2}=\left(t_{a, 2}, t_{p, 2}\right)$ be states of the children $s_{1}, s_{2}$ of a join node $s$. We say that $\psi_{1}$ and $\psi_{2}$ coincide with the state $\psi$ if

- $\psi(v)=(0,0)$ if and only if $\psi_{1}(v)=(0,0)$ and $\psi_{2}(v)=(0,0)$;
- for $i \in\{a, p\}$ and all $v \in B(s)$,

$$
t_{i}(v)= \begin{cases}\min \left\{t_{i, 1}(v), t_{i, 2}(v)\right\} & \text { if } t_{i, 1}(v), t_{i, 2}(v) \neq \perp \\ t_{i, 1}(v) & \text { if } t_{i, 2}(v)=\perp \neq t_{i, 1}(v) \\ t_{i, 2}(v) & \text { if } t_{i, 1}(v)=\perp \neq t_{i, 2}(v) \\ \perp, & \text { otherwise }\end{cases}
$$

Then we calculate $c(s, \psi)$ by

$$
c(s, \psi)=\min _{\psi_{1}, \psi_{2}}\left\{c\left(s_{1}, \psi_{1}\right)+c\left(s_{2}, \psi_{2}\right)-\left|t_{a}^{-1}(0)\right|\right\}
$$

where the minimum is taken over all pairs of states $\psi_{1}, \psi_{2}$ which coincide with $\psi$.
Proof. We begin by showing that a state $\psi$ of a join node $s$ is consistent if and only if there are consistent states $\psi_{1}$ and $\psi_{2}$ of its children which coincide with $\psi$. We then give a proof of correctness of our calculation of the signature.

Suppose that the states $\psi_{1}$ and $\psi_{2}$ are consistent for the nodes $s_{1}$ and $s_{2}$ respectively and they coincide with the state $\psi$. Then, let $S$ be the union of sets $S_{1}$ and $S_{2}$ which support $\psi_{1}$ and $\psi_{2}$ respectively. By our assumption of consistency, all vertices in $V\left(G_{s}\right) \backslash B(s)$ are temporally reachable from a vertex in $S \cup B(s)$ and a foremost temporal path from a vertex in $S$ arrives at each vertex $u$ in $B(s)$ at time $t_{a}(u)$. In addition, if any forgotten vertex $w$ which is not reached from $S$ is temporally reachable from a vertex $v \in B(s)$ by a path departing at some time $t \succ t_{p, j}(v)$ for $j \in\{1,2\}$, then $w$ must be temporally reached from $v$ by a temporal path departing at some time $t \succ t_{p}(v)$. Therefore, $\psi$ must be a consistent state of $s$.

We now assume that $\psi$ is a consistent state of $s$ supported by the set $S$. Suppose, for a contradiction, that $S$ supports $\psi$ and there are no child states supported by $S$ which coincide with $\psi$. There must be states for which $t_{a}$ describes the earliest time of arrival of a path from $S$ to any vertex in $B\left(s_{1}\right)$ and $B\left(s_{2}\right)$. Therefore, if there do not exist states $\psi_{1}$ and $\psi_{2}$ supported by $S$, then there must be a vertex $w$ in $G_{s_{1}} \backslash B\left(s_{1}\right)$ or $G_{s_{2}} \backslash B\left(s_{2}\right)$ which is not temporally reachable from a vertex in $S \cup B(s)$. If this is the case, then $\psi$ must also be inconsistent. This is because, by construction of a tree decomposition, the only (temporal) path from a vertex $G_{s_{1}} \backslash B\left(s_{1}\right)$ to $G_{s_{2}} \backslash B\left(s_{2}\right)$ or vice versa must traverse at least one vertex in $B(s)=B\left(s_{1}\right)=B\left(s_{2}\right)$. Therefore, if there is a temporal path from $S$ or $B(s)$ to each vertex in $G_{s} \backslash B(s)$, there must be a temporal path from $S \cup B\left(s_{1}\right)$ and $S \cap B\left(s_{2}\right)$ to each vertex in $G_{s_{1}} \backslash B\left(s_{1}\right)$ and $G_{s_{2}} \backslash B\left(s_{2}\right)$ respectively. Therefore, there exist child states supported by $S$.

The smallest set $S$ that supports a consistent state of $s$ must therefore be the smallest set that is the union of sets $S_{1}$ and $S_{2}$ which support consistent states $\psi_{1}, \psi_{2}$ respectively of its children $s_{1}$ and $s_{2}$ which coincide. To find the cardinality of this set, we employ the inclusion-exclusion principle. The only vertices which $G_{s_{1}}$
and $G_{s_{2}}$ have in common are those in $B(s)$. Therefore, $c(s, \psi)=\min _{\psi_{1}, \psi_{2}}\left\{c\left(s_{1}, \psi_{1}\right)+c\left(s_{2}, \psi_{2}\right)-\left|t_{a}^{-1}(0)\right|\right\}$ where the minimum is taken over pairs of states $\psi_{1}, \psi_{2}$ which coincide with $\psi$.

### 4.3.3 Running Time and Extensions

We note that for every vertex $v$ in a bag, there are strictly fewer than $(\tau+2)^{2}$ values of $\psi(v)$. The number of vertices in a bag is bounded by treewidth $\omega$ plus 1 . Thus, there are less than $(\tau+2)^{2(\omega+1)}$ states per node. We now explore the running time of calculating the signature at each type of node. For leaf nodes, the signature can only be one value; therefore this can be found in constant time. If we have the signatures of all child states of an introduce node and we want to verify validity of a given state of that node, we need to compute the value of $a_{v}$ for the introduced vertex $v$, the reachability set $R_{v}^{t_{a}(v)}$ and the earliest time of arrival at each other vertex in the bag from $v$ by a path departing at times $t \succ t_{a}(v)$ and $t^{\prime} \succ t_{p}(v)$.

The reachability set $R_{v}^{t}(\mathcal{G})$ can be computed in time $O\left(n^{2}\right)$ by a modified breadth first search algorithm for any $v$ and $t$. Hence it takes $O\left(\tau n^{3}\right)$ time to calculate $R_{v}^{t}(\mathcal{G})$ for every vertex $v \in V$ and time $t \in[\tau]$. This can be done as a preprocessing step before we begin working up the tree decomposition.

The earliest time of arrival of the two paths can be computed using a variation of breadth-first search and thus takes $O\left(\omega^{2}\right)$ time. The value $a_{v}$ is computable in $O(\omega)$ time. In addition, for some cases in Nonstrict TARDIS we find the minimum signature of restrictions of the state where we change the value of $t_{a}$ for some neighbours of $v$. Assuming all signatures of descendant nodes are calculated, this takes $O(\omega)$ time. Therefore, computing the signature of all states of an introduce node requires $O\left(\omega^{2}(\tau)^{2(\omega+1)}\right)$ time.

For forget nodes, we must compare the signatures of multiple child states. There are $O\left(\tau^{2}\right)$ of these extensions for each state of the forget node. In addition, we compute the earliest time of arrival of a path from any vertex in $B(s)$ to $v$ in the subgraph induced by $B(s)$ with restrictions on its departure time. This can be achieved by a variant of a breadth-first search which takes at most $O\left(\omega^{3}\right)$ time. Thus, finding the signature of all states of a forget node requires $O\left(\omega^{3}(\tau)^{2 \omega+4}\right)$ time.

Finally, to compute the signature of the state of a join node, we must compare the states of both child nodes which coincide with this state. For a given vertex in the bag of a join node, the number of values of $\psi_{1}$ and $\psi_{2}$ which coincide with $\psi(v)$ are bounded by $O\left(\tau^{4}\right)$. Therefore, there are $O\left(\tau^{4(\omega+1)}\right)$ tuples of states to consider when calculating the signature of $\psi$ on $s$. Therefore, we calculate the signature of a state of a join node in $O\left(\tau^{4(\omega+1)} \cdot \omega\right)$ time. Note that this dominates the time needed to generate all possible states for a given bag.

We now combine the lemmas from this section to get the following theorem. We can find a tree decomposition of width at most $\omega$ if one exists in linear time by Theorem 6. Lemma 34 states that we can find a nice tree decomposition of width $\omega$ given a tree decomposition of width $\omega$ in linear time. We can recursively compute the signature of a given state of a node using Lemmas 35,36 and 37 . We solve TARDIS by finding the signature $c(r, \psi)$ of the root where the state $\psi$ is the empty function which gives the cardinality of a minimal TaRDiS.
Theorem 7. The algorithm described takes as input a temporal graph $\mathcal{G}$ consisting of $n$ vertices with a nice tree decomposition of width at most $\omega$ and solves Strict and Nonstrict TaRDiS on $\mathcal{G}$ in time $O\left(\tau^{4(\omega+1)} \cdot \omega \cdot n+\tau n^{3}\right)$.

We emphasize that $\tau n^{3}$ is polynomial in the input size because $\tau \leq \mathcal{E}$ (recall, no snapshot in a temporal graph is empty). The algorithm allows for edges to be active multiple times. That is, it is not restricted to simple temporal graphs.

## 5 Parameterized complexity results for MaxMinTaRDiS

Having shown MaxMinTaRDIS is in $\Sigma_{2}^{P}$, finding instances of tractability is even more surprising than with the variants of the TARDIS problem. We begin the following result, closely related to Lemma 33 which gives us tractability when each component of the input to Strict MaxMinTaRDiS is restricted.
Lemma 38. Strict MaxMintarDis is in FPT when parameterized by maximum degree $\Delta$ in $H$ and $k$.
Proof. Recall from Lemma 19 that $(H, k)$ is a yes-instance of Strict MaxMinTaRDiS if and only if $(H, k-1)$ is a no-instance of Dominating Set. Also, any pair $(H, k)$ where $H$ has maximum degree $\Delta$ satisfying $|V(H)|>k(\Delta+1)$ is trivially a no-instance of Dominating Set. Applying the same reasoning as in our proof of Lemma 33, we obtain that Strict MaxMinTaRDiS is solvable in polynomial time when $\Delta$ and $k$ are bounded. If the input graph has at least $k(\Delta+1)$ vertices then output YES, and otherwise solve the problem (necessarily of bounded size) by brute-force.

### 5.1 Existence of an fpt algorithm for MaxMinTaRDiS

We show tractability of our problems by expressing them in EMSO logic and applying the variant of Courcelle's theorem given by Arnborg, Lagergren and Seese [55]. This result states that an optimisation problem which is definable in EMSO can be solved in polynomial time when parameterized by treewidth and length of the formula. The theorem is as follows.
Theorem 8 (adapted from [56], Theorem 30). Let $P$ be an EMSO-definable problem, then one can solve $P$ on graphs $G=(V, E)$ of order $n:=|V|$ and treewidth at most $w$ in time $O\left(f_{P}(w) \cdot \operatorname{poly}(n)\right)$.

An extended monadic second order formula (EMSO) over a static graph $H$ is a formula that uses:

1. the logical operators $\vee, \wedge, \neg,=$ and parentheses;
2. a finite set of variables, each of which takes an element of $V(H)$ or $E(H)$;
3. the quantifiers $\forall$ and $\exists$;
4. a finite set of variables which take subsets of the sets of edges or vertices;
5. integers.

We make use of the predicates $\neq, \Longrightarrow, \Longleftarrow, \Longleftrightarrow, \subseteq, \in$ and $\backslash$, which can be implemented using the above. We note that formulas consisting only of the first 3 components is a first-order formula (FO). A formula which is first order with the addition of the fourth bullet point is referred to as an MSO formula. A more formal definition of an EMSO-definable problem is given by a survey by Langer, Reidl, Rossmanith and Sikdar [56].
Theorem 9. MaxMinTaRDiS is fixed-parameter tractable when parameterized by lifetime, $k$ and treewidth of the graph.

Proof. The formal definition of EMSO given by Langer Reidl, Rossmanith and Sikdar [56] requires a weight function bounded by a constant. Our weight function will be the cardinality of a TaRDiS, which is bounded by $k$. Since the temporal assignment is not part of the input, we encode it as a partition of edges into sets which correspond to the time at which they are active. The EMSO formula is constructed using the following auxiliary subformulae

- $\operatorname{card}(k, X)$ tests whether a set $X$ has cardinality at least $k$ :

$$
\operatorname{card}(k, X):=\exists x_{1}, \ldots, x_{k} \in X: \bigwedge_{1 \leq i<j \leq k} x_{i} \neq x_{j}
$$

- $\operatorname{geq}\left(X_{1}, X_{2}\right)$ tests whether $\left|X_{1}\right| \geq\left|X_{2}\right|$ for sets $X_{1}$ and $X_{2}$ :

$$
\operatorname{geq}\left(X_{1}, X_{2}\right):=\exists k: \operatorname{card}\left(k, X_{1}\right) \wedge \neg \operatorname{card}\left(k+1, X_{2}\right)
$$

- $\operatorname{part}\left(S_{1}, \ldots, S_{\tau}\right)$ tests whether the sets of edges $S_{1} \ldots S_{\tau}$ partition the edges of $H$ :

$$
\begin{aligned}
& \operatorname{part}\left(S_{1}, \ldots, S_{\tau}\right):=\forall e \in E: \\
& \left(\bigvee_{1 \leq i \leq \tau} e \in S_{i} \wedge\left(\bigwedge_{1 \leq j<i} e \notin S_{j}\right) \wedge\left(\bigwedge_{i<\ell \leq \tau} e \notin S_{\ell}\right)\right)
\end{aligned}
$$

The two right-most brackets can be ignored if we do not require that the temporal assignment is simple. To enforce a happy temporal assignment, we can add the requirement that no two edges in the same set share and endpoint. Recall that tractability of Happy MaxMinTaRDiS is shown with respect to lifetime and $k$ combined in Lemma 38. That is a stronger result than what we have here since Courcelle's theorem only gives tractability on graphs with bounded treewidth.
The following subformulae can be adapted to write TARDIS in MSO logic.

- mconn $\left(X, S_{t}\right)$ tests whether all vertices in a set $X$ are in the same connected component of $G_{t}$.

$$
\operatorname{mconn}\left(X, S_{t}\right):=\forall Y \subset X, Y \neq \emptyset, \exists x \in X, \exists y \in Y \backslash X: x y \in S_{t}
$$

- $\operatorname{mvconn}(v, w, t)$ tests whether two vertices are in the same connected component of $G_{t}$ :

$$
\operatorname{mvconn}(v, w, t):=\exists X \subset V: v \in X \wedge w \in X \wedge \operatorname{mconn}\left(X, S_{t}\right)
$$

- mtadj tests whether two vertices $v, w \in V$ are adjacent at time $t$ :

$$
\operatorname{mtadj}\left(v, w, S_{t}\right):=\exists e \in S_{t}: v \in e \wedge w \in e
$$

- mpath tests whether there is a temporal path from $v$ to $w \in V$ with latest time $\tau$ :

$$
\begin{array}{r}
\operatorname{mpath}\left(v, w, S_{1}, \ldots, S_{\tau}\right):=\exists v_{0}, \ldots, v_{\tau} \in V: \\
v=v_{0} \wedge v_{\tau}=w \wedge \bigwedge_{t=0}^{\tau-1}\left(v_{t}=v_{t+1} \vee \mathrm{a}\left(v_{t}, v_{t+1}, S_{t+1}\right)\right)
\end{array}
$$

where $\mathrm{a}\left(v, w, S_{t}\right)$ can be substituted for $\operatorname{mvconn}(v, w, t)$ or $\operatorname{mtadj}\left(v, w, S_{t}\right)$ depending on whether we are testing for nonstrict or strict temporal paths respectively.

- $\operatorname{TaRDiS}\left(S_{1}, \ldots, S_{\tau}, X\right)$ which tests if every vertex is temporally reachable from $X$ by a temporal path of lifetime $\tau$ :

$$
\operatorname{TaRDiS}\left(S_{1}, \ldots, S_{\tau}, X\right):=\forall v \in V(H), \exists s \in X: \operatorname{mpath}\left(s, v, S_{1}, \ldots, S_{\tau}\right)
$$

- $m \operatorname{TaRDiS}\left(S_{1}, \ldots, S_{\tau}, X\right)$ tests whether a set $X$ is a minimum TaRDiS:

$$
\begin{aligned}
m \operatorname{TaRDiS}\left(S_{1}, \ldots, S_{\tau}, X\right):= & \forall X^{\prime} \subset V(H): \operatorname{TaRDiS}\left(S_{1}, \ldots, S_{\tau}, X\right) \wedge \\
& \left(\operatorname{TaRDIS}\left(S_{1}, \ldots, S_{\tau}, X^{\prime}\right) \Longrightarrow \operatorname{geq}\left(X^{\prime}, X\right)\right)
\end{aligned}
$$

- MinTaRDiS $(H, \tau, k)$ which tests if there is temporal assignment with lifetime at most $\tau$ such that there exists a minimum TaRDiS of size at least $k$ on $H$ :

$$
\begin{aligned}
& \operatorname{MinTARDIS}(H, \tau, k):=\exists X \subset V, \exists S_{1}, \ldots, S_{\tau} \subset E: \\
& \operatorname{part}\left(S_{1}, \ldots, S_{\tau}\right) \wedge m \operatorname{TaRDIS}\left(S_{1}, \ldots, S_{\tau}, X\right) \wedge \operatorname{card}(k, X)
\end{aligned}
$$

Therefore MaxMinTaRDiS can be expressed in EMSO and the theorem holds.
Corollary 9. Temporal reachability is expressible in MSO logic.
Furthermore, we can use the expression to express TARDIS in EMSO. This, however, gives a weaker tractability result than Theorem 7.

## 6 Conclusions and open questions

In this paper, we introduce the TARDiS and MaxMinTaRDiS problems and study their (parameterized) complexity. We show a bound on the lifetime $\tau$ and a restriction to planar inputs combined are insufficient to obtain tractability (Theorems 1, 4, and Corollary 7) and moreover tightly characterize the minimum lifetime $\tau$ for which each problem becomes intractable. Further, we give an algorithm on a nice tree decomposition of a temporal graph which gives tractability of TARDIS with respect to lifetime and treewidth of the footprint of the graph. In addition, we show that $\tau, k$ and the treewidth of the input graph combined are sufficient to yield tractability in all cases of MaxMinTaRDIS by leveraging Courcelle's theorem.

These results leave open the questions of the exact complexity of Nonstrict MaxMinTaRDiS with lifetime $\tau \geq 3$ and Happy MaxMinTaRDiS with lifetime $\tau \geq 4$. An interesting extension of our work would be to find approximability results for these problems. From the parameterized side, it remains to be shown whether parameterization by a structural parameter of the footprint (e.g. treewidth) alone is sufficient to obtain tractability for any of the considered variants. Another interesting dimension is the comparison of Nonstrict MaxMinTaRDiS and Happy MaxMinTaRDiS when $\tau$ is lower-bounded by a function of the number of edges $m$. With the constraint $\tau=m$ the two problems become equivalent, and their computational complexity in this case is an interesting open question. Analogously to $t$-Dominating Set [57], $t$-TARDIS, in which $t$ individuals must be reached provides a natural generalisation of our problem and the potential for parameterization by $t$. Recall that, in Strict MaxMinTaRDiS, it is always optimal to choose the constant function as our temporal assignment $\lambda$. It may be interesting to consider restrictions on $\lambda$ other than happiness which require the use of a non-constant temporal assignment (as in [58]) to make the problem more interesting.

## Declarations

The authors have no relevant financial or non-financial interests to disclose. The authors have no conflicts of interest to declare that are relevant to the content of this article. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no financial or proprietary interests in any material discussed in this article.

## References

[1] Kempe, D., Kleinberg, J., Kumar, A.: Connectivity and inference problems for temporal networks. In: Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing. STOC ' 00 , pp. 504-513. Association for Computing Machinery, New York, NY, USA (2000). https://doi.org/10.1145/ 335305.335364
[2] Deligkas, A., Potapov, I.: Optimizing reachability sets in temporal graphs by delaying. Information and Computation 285, 104890 (2022) https://doi.org/10.1016/j.ic.2022.104890
[3] Enright, J., Meeks, K., Skerman, F.: Assigning times to minimise reachability in temporal graphs. Journal of Computer and System Sciences 115, 169-186 (2021) https://doi.org/10.1016/j.jcss.2020.08. 001
[4] Molter, H., Renken, M., Zschoche, P.: Temporal reachability minimization: Delaying vs. deleting. arXiv preprint arXiv:2102.10814 (2021)
[5] Klobas, N., Mertzios, G.B., Molter, H., Spirakis, P.G.: The complexity of computing optimum labelings for temporal connectivity. arXiv preprint arXiv:2202.05880 (2022)
[6] Casteigts, A., Peters, J.G., Schoeters, J.: Temporal cliques admit sparse spanners. Journal of Computer and System Sciences 121, 1-17 (2021) https://doi.org/10.1016/j.jcss.2021.04.004
[7] Erlebach, T., Spooner, J.T.: A game of cops and robbers on graphs with periodic edge-connectivity. CoRR (2019) 1908.06828
[8] Erlebach, T., Spooner, J.T.: Faster Exploration of Degree-Bounded Temporal Graphs. In: Potapov, I., Spirakis, P., Worrell, J. (eds.) 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018). Leibniz International Proceedings in Informatics (LIPIcs), vol. 117, pp. 3613613. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2018). https://doi.org/ 10.4230/LIPIcs.MFCS.2018.36
[9] Whitbeck, J., Amorim, M.D., Conan, V., Guillaume, J.-L.: Temporal Reachability Graphs. arXiv (2012). https://doi.org/10.48550/arXiv.1207.7103
[10] Casteigts, A., Corsini, T., Sarkar, W.: Simple, strict, proper, happy: A study of reachability in temporal graphs. arXiv (2022). http://arxiv.org/abs/2208.01720
[11] Casteigts, A.: A Journey through Dynamic Networks (with Excursions). Thesis, Université de Bordeaux (June 2018). https:/ /hal.science/tel-01883384
[12] Pybus, O., Rambaut, A., COG-UK-Consortium, et al.: Preliminary analysis of SARS-CoV-2 importation \& establishment of UK transmission lineages. Virological. org. (2020)
[13] Deligkas, A., Eiben, E., Skretas, G.: Minimizing reachability times on temporal graphs via shifting labels. In: Elkind, E. (ed.) Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23, Macao, pp. 5333-5340 (2023). https://doi.org/10.24963/ijcai.2023/592
[14] 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) https://doi.org/ 10.1016/j.jcss.2021.01.007
[15] Munroe, R.: xkcd: Post Vaccine Social Scheduling. xkcd (2021). https://xkcd.com/2450/
[16] Eto, H., Guo, F., Miyano, E.: Distance- $d$ independent set problems for bipartite and chordal graphs. Journal of Combinatorial Optimization 27(1), 88-99 (2014) https://doi.org/10.1007/s10878-012-9594-4
[17] Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21275-3
[18] Courcelle, B.: The expression of graph properties and graph transformations in monadic second-order logic. In: Handbook Of Graph Grammars And Computing By Graph Transformation: Volume 1: Foundations, pp. 313-400. World Scientific, Singapore (1997)
[19] 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) https://doi.org/10.1016/j.jcss. 2019.08.002
[20] Bilò, D., D'Angelo, G., Gualà, L., Leucci, S., Rossi, M.: Blackout-Tolerant Temporal Spanners. In: Erlebach, T., Segal, M. (eds.) Algorithmics of Wireless Networks. Lecture Notes in Computer Science, pp. 31-44. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22050-0_3
[21] Füchsle, E., Molter, H., Niedermeier, R., Renken, M.: Delay-Robust Routes in Temporal Graphs. arXiv. arXiv:2201.05390 [cs] (2022). https://doi.org/10.48550/arXiv.2201.05390
[22] Braunstein, A., Ingrosso, A.: Inference of causality in epidemics on temporal contact networks. Scientific Reports 6(1), 27538 (2016) https://doi.org/10.1038/srep27538
[23] Valdano, E., Ferreri, L., Poletto, C., Colizza, V.: Analytical Computation of the Epidemic Threshold on Temporal Networks. Physical Review X 5(2), 021005 (2015) https://doi.org/10.1103/PhysRevX.5. 021005
[24] Tang, J., Leontiadis, I., Scellato, S., Nicosia, V., Mascolo, C., Musolesi, M., Latora, V.: Applications of Temporal Graph Metrics to Real-World Networks. In: Holme, P., Saramäki, J. (eds.) Temporal Networks. Understanding Complex Systems, pp. 135-159. Springer, Berlin, Heidelberg (2013). https: / /doi.org/10.1007/978-3-642-36461-7_7
[25] Peleg, D.: Time-efficient broadcasting in radio networks: A review. In: International Conference on Distributed Computing and Internet Technology, pp. 1-18 (2007). Springer
[26] Erlebach, T., Hall, A.: NP-Hardness of Broadcast Scheduling and Inapproximability of Single-Source Unsplittable Min-Cost Flow. Journal of Scheduling 7(3), 223-241 (2004) https://doi.org/10.1023/B: JOSH.0000019682.75022.96
[27] Jakoby, A., Reischuk, R., Schindelhauer, C.: The complexity of broadcasting in planar and decomposable graphs. In: Mayr, E.W., Schmidt, G., Tinhofer, G. (eds.) Graph-Theoretic Concepts in Computer Science. Lecture Notes in Computer Science, pp. 219-231. Springer, Berlin, Heidelberg (1995). https: //doi.org/10.1007/3-540-59071-4_50
[28] Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-Varying Graphs and Dynamic Networks. In: Frey, H., Li, X., Ruehrup, S. (eds.) Ad-hoc, Mobile, and Wireless Networks. Lecture Notes in Computer Science, pp. 346-359. Springer, Berlin, Heidelberg (2011). https://doi.org/10.1007/ 978-3-642-22450-8_27
[29] Chia, M.-L., Kuo, D., Tung, M.-F.: The multiple originator broadcasting problem in graphs. Discrete Applied Mathematics 155(10), 1188-1199 (2007) https://doi.org/10.1016/j.dam.2006.10.011
[30] Grigoryan, H.: Problems related to broadcasting in graphs. phd, Concordia University (September 2013). https://spectrum.library.concordia.ca/id/eprint/977773/ Accessed 2023-09-20
[31] Akrida, E.C., Mertzios, G.B., Spirakis, P.G.: The Temporal Explorer Who Returns to the Base. In: Heggernes, P. (ed.) Algorithms and Complexity. Lecture Notes in Computer Science, pp. 13-24. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17402-6_2
[32] Bumpus, B.M., Meeks, K.: Edge exploration of temporal graphs. arXiv. arXiv:2103.05387 [cs, math]
(2021). https://doi.org/10.48550/arXiv.2103.05387
[33] Mertzios, G.B., Molter, H., Niedermeier, R., Zamaraev, V., Zschoche, P.: Computing maximum matchings in temporal graphs. Journal of Computer and System Sciences 137, 1-19 (2023) https://doi.org/ 10.1016/j.jcss.2023.04.005
[34] Klobas, N., Mertzios, G.B., Molter, H., Niedermeier, R., Zschoche, P.: Interference-free walks in time: Temporally disjoint paths. Autonomous Agents and Multi-Agent Systems 37(1), 1 (2023)
[35] Enright, J., Meeks, K.: Deleting Edges to Restrict the Size of an Epidemic: A New Application for Treewidth. Algorithmica 80(6), 1857-1889 (2018) https://doi.org/10.1007/s00453-017-0311-7
[36] Erlebach, T., Hoffmann, M., Kammer, F.: On Temporal Graph Exploration. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) Automata, Languages, and Programming. Lecture Notes in Computer Science, pp. 444-455. Springer, Berlin, Heidelberg (2015). https://doi.org/10.1007/ 978-3-662-47672-7_36
[37] Haag, R., Molter, H., Niedermeier, R., Renken, M.: Feedback edge sets in temporal graphs. Discrete Applied Mathematics 307, 65-78 (2022) https://doi.org/10.1016/j.dam.2021.09.029
[38] Courcelle, B., Engelfriet, J.: Monadic second-order logic. In: Courcelle, B., Engelfriet, J. (eds.) Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach. Encyclopedia of Mathematics and its Applications, pp. 315-426. Cambridge University Press, Cambridge (2012). https: //doi.org/10.1017/CBO9780511977619.007
[39] Downey, R.G., Fellows, M.R.: Fixed-Parameter Tractability and Completeness I: Basic Results. SIAM Journal on Computing 24(4), 873-921 (1995) https://doi.org/10.1137/S0097539792228228
[40] Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: Fourier meets Möbius: fast subset convolution. arXiv e-prints, 0611101 (2006) https://doi.org/10.48550/arXiv.cs/0611101
[41] Bhadra, S., Ferreira, A.: Complexity of Connected Components in Evolving Graphs and the Computation of Multicast Trees in Dynamic Networks. In: Pierre, S., Barbeau, M., Kranakis, E. (eds.) Ad-Hoc, Mobile, and Wireless Networks. Lecture Notes in Computer Science, pp. 259-270. Springer, Berlin, Heidelberg (2003). https://doi.org/10.1007/978-3-540-39611-6_23
[42] Verheije, M.: Algorithms for Domination Problems on Temporal Graphs. PhD thesis, Utrecht University (2021). Accepted: 2021-08-26. https://studenttheses.uu.nl/handle/20.500.12932/41240
[43] Balev, S., Pigné, Y., Sanlaville, E., Schoeters, J.: Temporally connected components (2023). https: //hal.science/hal-03966327
[44] Garey, M.R., Johnson, D.S., Stockmeyer, L.: Some simplified NP-complete problems. In: Proceedings of the Sixth Annual ACM Symposium on Theory of Computing. STOC '74, pp. 47-63. Association for Computing Machinery, New York, NY, USA (1974). https://doi.org/10.1145/800119.803884
[45] Tippenhauer, S., Muzler, W.: On planar 3-sat and its variants. Fachbereich Mathematik und Informatik der Freien Universitat Berlin (2016)
[46] Karp, R.M.: In: Miller, R.E., Thatcher, J.W., Bohlinger, J.D. (eds.) Reducibility among Combinatorial Problems, pp. 85-103. Springer, Boston, MA (1972). https://doi.org/10.1007/978-1-4684-2001-2_9
[47] Holyer, I.: The NP-completeness of edge-coloring. SIAM Journal on computing 10(4), 718-720 (1981)
[48] Arora, S., Barak, B.: Computational Complexity: A Modern Approach. Cambridge University Press, Cambridge (2009). https://doi.org/10.1017/CBO9780511804090
[49] Papadimitriou, C.H.: Computational Complexity. Addison-Wesley, Reading, Mass (1994)
[50] Gutner, S.: The complexity of planar graph choosability. arXiv. arXiv:0802.2668 [cs] (2008). https: //doi.org/10.48550/arXiv.0802.2668
[51] Eto, H., Ito, T., Liu, Z., Miyano, E.: Approximability of the Distance Independent Set Problem on Regular Graphs and Planar Graphs. In: Chan, T.-H.H., Li, M., Wang, L. (eds.) Combinatorial Optimization and Applications. Lecture Notes in Computer Science, pp. 270-284. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48749-6_20
[52] Agnarsson, G., Damaschke, P., Halldórsson, M.M.: Powers of geometric intersection graphs and dispersion algorithms. Discrete Applied Mathematics 132(1), 3-16 (2003) https://doi.org/10.1016/ S0166-218X(03)00386-X
[53] Niedermeier, R.: Invitation to Fixed-Parameter Algorithms. Oxford University Press, England (2006). https://doi.org/10.1093/acprof:oso/9780198566076.001.0001
[54] Bodlaender, H.L.: A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth. SIAM Journal on Computing 25(6), 1305-1317 (1996) https://doi.org/10.1137/S0097539793251219
[55] Arnborg, S., Lagergren, J., Seese, D.: Easy problems for tree-decomposable graphs. Journal of Algorithms 12(2), 308-340 (1991) https://doi.org/10.1016/0196-6774(91)90006-K
[56] Langer, A., Reidl, F., Rossmanith, P., Sikdar, S.: Practical algorithms for MSO model-checking on treedecomposable graphs. Computer Science Review 13-14, 39-74 (2014) https://doi.org/10.1016/j.cosrev. 2014.08.001
[57] Kneis, J., Mölle, D., Rossmanith, P.: Partial vs. Complete Domination: t-Dominating Set. In: Leeuwen, J., Italiano, G.F., Hoek, W., Meinel, C., Sack, H., Plášil, F. (eds.) SOFSEM 2007: Theory and Practice of Computer Science. Lecture Notes in Computer Science, pp. 367-376. Springer, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-69507-3_31
[58] Christiann, E., Sanlaville, E., Schoeters, J.: On inefficiently connecting temporal networks. arXiv preprint arXiv:2312.07117 (2023)


[^0]:    ${ }^{1}$ We use fpt (lowercase) as a descriptor for algorithms witnessing the inclusion of a problem in the parameterized complexity class FPT. A full definition is given in Section 4.
    ${ }^{2}$ Informally, the treewidth of a graph is a measure of its likeness to a tree. A formal definition is given in Section 4.

[^1]:    ${ }^{3}$ Our definition of equivalence can be found in Section 3.

[^2]:    Set Cover
    Input: Universe $\mathcal{U}=\left\{x_{1}, \ldots x_{n}\right\}$, a family $\mathcal{S}=\left\{s_{i} \mid s_{i} \subseteq \mathcal{U}\right\}$ of subsets in $\mathcal{U}$ and an integer $k$. Question: Is there a set $\left\{s_{i \in I}\right\}$ with at most $k$ elements such that $\cup_{i \in I} s_{i}=\mathcal{U}$ ?

