

Temporal Vertex Cover with a Sliding Time Window*

Eleni C. Akrida[†] George B. Mertzios[‡] Paul G. Spirakis[§] Viktor Zamaraev[¶]

Abstract

Modern, inherently dynamic systems are usually characterized by a network structure, i.e. an underlying graph topology, which is subject to discrete changes over time. Given a static underlying graph G , a temporal graph can be represented via an assignment of a set of integer time-labels to every edge of G , indicating the discrete time steps when this edge is active. While most of the recent theoretical research on temporal graphs has focused on the notion of a temporal path and other “path-related” temporal notions, only few attempts have been made to investigate “non-path” temporal graph problems. In this paper, motivated by applications in sensor and in transportation networks, we introduce and study two natural temporal extensions of the classical problem VERTEX COVER. In our first problem, TEMPORAL VERTEX COVER, the aim is to cover every edge at least once during the lifetime of the temporal graph, where an edge can only be covered by one of its endpoints at a time step when it is active. In our second, more pragmatic variation SLIDING WINDOW TEMPORAL VERTEX COVER, we are also given a natural number Δ , and our aim is to cover every edge at *least once* at *every* Δ consecutive time steps. In both cases we wish to minimize the total number of “vertex appearances” that are needed to cover the whole graph. We present a thorough investigation of the computational complexity and approximability of these two temporal covering problems. In particular, we provide strong hardness results, complemented by various approximation and exact algorithms. Some of our algorithms are polynomial-time, while others are asymptotically almost optimal under the Exponential Time Hypothesis (ETH) and other plausible complexity assumptions.

Keywords: Temporal networks, temporal vertex cover, APX-hard, approximation algorithm, Exponential Time Hypothesis.

1 Introduction and Motivation

A great variety of both modern and traditional networks are inherently dynamic, in the sense that their link availability varies over time. Information and communication networks, social networks, transportation networks, and several physical systems are only a few examples of networks that change over time [22, 33]. The common characteristic in all these application areas is that the network structure, i.e. the underlying graph topology, is subject to *discrete changes over time*. In this paper we adopt a simple and natural model for time-varying networks which is given with time-labels on the edges of a graph, while the vertex set remains unchanged. This formalism originates in the foundational work of Kempe et al. [25].

Definition 1 (temporal graph). *A temporal graph is a pair (G, λ) , where $G = (V, E)$ is an underlying (static) graph and $\lambda : E \rightarrow 2^{\mathbb{N}}$ is a time-labeling function which assigns to every edge of G a set of discrete-time labels.*

For every edge $e \in E$ in the underlying graph G of a temporal graph (G, λ) , $\lambda(e)$ denotes the set of time slots at which e is *active* in (G, λ) . Due to its vast applicability in many areas, this notion of temporal graphs has been studied from different perspectives under various names such as *time-varying* [1, 17, 35], *evolving* [5, 11, 16], *dynamic* [8, 19], and *graphs over time* [29]; for a recent attempt to integrate existing models, concepts, and results from the distributed computing perspective see the survey papers [6–8] and the references therein. Data analytics on temporal networks have also been very recently studied in the context of summarizing networks that represent sports teams’ activity data to discover recurring

*This work was partially supported by the NeST initiative of the School of EEE and CS at the University of Liverpool and by the EPSRC Grants EP/P020372/1 and EP/P02002X/1.

[†]Department of Computer Science, University of Liverpool, Liverpool, UK. Email: e.akrida@liverpool.ac.uk

[‡]Department of Computer Science, Durham University, Durham, UK. Email: george.mertzios@durham.ac.uk

[§]Department of Computer Science, University of Liverpool, Liverpool, UK. Email: p.spirakis@liverpool.ac.uk

[¶]Department of Computer Science, Durham University, Durham, UK. Email: viktor.zamaraev@durham.ac.uk

strategies and understand team tactics [27], as well as extracting patterns from interactions between groups of entities in a social network [26].

Motivated by the fact that, due to causality, information in temporal graphs can “flow” only along sequences of edges whose time-labels are increasing, most temporal graph parameters and optimization problems that have been studied so far are based on the notion of temporal paths and other “path-related” notions, such as temporal analogues of distance, diameter, reachability, exploration, and centrality [2, 3, 14, 31, 32]. In contrast, only few attempts have been made to define “non-path” temporal graph problems. Motivated by the contact patterns among high-school students, Viard et al. [37, 38], and later Himmel et al. [21], introduced and studied Δ -cliques, an extension of the concept of cliques to temporal graphs, in which all vertices interact with each other at least once every Δ consecutive time steps within a given time interval.

In this paper we introduce and study two natural temporal extensions of the problem VERTEX COVER in static graphs, which take into account the dynamic nature of the network. In the first and simpler of these extensions, namely TEMPORAL VERTEX COVER (for short, TVC), every edge e has to be “covered” at least once during the lifetime T of the network (by one of its endpoints), and this must happen at a time step t when e is active. The goal is then to cover all edges with the minimum total number of such “vertex appearances”. On the other hand, in many real-world applications where scalability is important, the lifetime T can be arbitrarily large but the network still needs to remain sufficiently covered. In such cases, as well as in safety-critical systems (e.g. in military applications), it may not be satisfactory enough that an edge is covered just *once* during the *whole lifetime* of the network. Instead, every edge must be covered at least once within *every small Δ -window* of time (for an appropriate value of Δ), regardless of how large the lifetime is; this gives rise to our second optimization problem, namely SLIDING WINDOW TEMPORAL VERTEX COVER (for short, SW-TVC). Formal definitions of our problems TVC and SW-TVC are given in Section 2.

Our two temporal extensions of VERTEX COVER are motivated by applications in sensor networks and in transportation networks. In particular, several works in the field of sensor networks considered problems of placing sensors to cover a whole area or multiple critical locations, e.g. for reasons of surveillance. Such studies usually wish to minimize the number of sensors used or the total energy required [13, 20, 28, 34, 39]. Our temporal vertex cover notions are an abstract way to economically meet such covering demands as time progresses.

To further motivate the questions raised in this work, consider a network whose links represent transporting facilities which are not always available, while the availability schedule per link is known in advance. We wish to check each transporting facility and certify “OK” at least once per facility during every (reasonably small) window of time. It is natural to assume that the checking is done in the presence of an inspecting agent at an endpoint of the link (i.e. on a vertex), since such vertices usually are junctions with local offices. The agent can inspect more than one link at the same day, provided that these links share this vertex and that they are all alive (i.e. operating) at that day. Notice that the above is indeed an application drawn from real-life, as regular checks in roads and trucks are paramount for the correct operation of the transporting sector, according to both the European Commission¹ and the American Public Transportation Association².

1.1 Our contribution

In this paper we present a thorough investigation of the complexity and approximability of the problems TEMPORAL VERTEX COVER (TVC) and SLIDING WINDOW TEMPORAL VERTEX COVER (SW-TVC) on temporal graphs. We first prove in Section 3 that TVC remains NP-complete even on star temporal graphs, i.e. when the underlying graph G is a star. Furthermore we prove that, for any $\varepsilon < 1$, TVC on star temporal graphs cannot be optimally solved in $O(2^{\varepsilon n})$ time, assuming the Strong Exponential Time Hypothesis (SETH), as well as that it does not admit a polynomial-time $(1 - \varepsilon) \ln n$ -approximation algorithm, unless NP has $n^{O(\log \log n)}$ -time deterministic algorithms. On the positive side, we prove that

¹According to the European Commission (see https://ec.europa.eu/transport/road_safety/topics/vehicles/inspection_en), “roadworthiness checks (such as on-the-spot roadside inspections and periodic checks) not only make sure your vehicle is working properly, they are also important for environmental reasons and for ensuring fair competition in the transport sector”.

²According to the American Public Transportation Association (see <http://www.apta.com/resources/standards/Documents/APTA-RT-VIM-RP-019-03.pdf>) “developing minimum inspection, maintenance, testing and alignment procedures maintains rail transit trucks in a safe and reliable operating condition”.

TVC on star temporal graphs with n vertices can be $(H_{n-1} - \frac{1}{2})$ -approximated in polynomial time, where $H_n = \sum_{i=1}^n \frac{1}{i} \approx \ln n$ is the n th harmonic number. Moreover we prove that TVC on general temporal graphs admits a polynomial-time randomized approximation algorithm with expected ratio $O(\ln n)$.

In Section 4 and in the remainder of the paper we deal with SW-TVC. For our second problem, SW-TVC, we prove in Section 4.1 a strong complexity lower bound on arbitrary temporal graphs. More specifically we prove that, for *any* (arbitrarily growing) functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, there exists a constant $\varepsilon \in (0, 1)$ such that SW-TVC cannot be solved in $f(T) \cdot 2^{\varepsilon n \cdot g(\Delta)}$ time, assuming the Exponential Time Hypothesis (ETH). This ETH-based lower bound turns out to be asymptotically almost tight, as we present an exact dynamic programming algorithm with running time $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$. This worst-case running time can be significantly improved in certain special temporal graph classes. In particular, when the “snapshot” of (G, λ) at every time step has vertex cover number bounded by k , the running time becomes $O(T\Delta(n+m) \cdot n^{k(\Delta+1)})$. That is, for small values of Δ (say, when Δ is $O(\log n)$ or $O(\log T)$), this algorithm is polynomial in the input size on temporal graphs with bounded vertex cover number at every time step. Notably, when every snapshot is a star (i.e. a superclass of the star temporal graphs studied in Section 3) the running time of the algorithm is $O(T\Delta(n+m) \cdot 2^\Delta)$.

In Section 5 we prove strong inapproximability results for SW-TVC even when restricted to temporal graphs with length $\Delta = 2$ of the sliding window. In particular, we prove that this problem is APX-hard (and thus does not admit a *Polynomial Time Approximation Scheme (PTAS)*, unless $P = NP$), even when $\Delta = 2$, the maximum degree in the underlying graph G is at most 3, and every connected component at every graph snapshot has at most 7 vertices. Finally, in Section 6 we provide a series of approximation algorithms for the general SW-TVC problem, with respect to various incomparable temporal graph parameters. In particular, we provide polynomial-time approximation algorithms with approximation ratios (i) $O(\ln n + \ln \Delta)$, (ii) $2k$, where k is the maximum number of times that each edge can appear in a sliding Δ time window (thus implying a ratio of 2Δ in the general case), (iii) d , where d is the maximum vertex degree at every snapshot of (G, λ) . Note that, for $d = 1$, the latter result implies that SW-TVC can be optimally solved in polynomial time whenever every snapshot of (G, λ) is a matching.

2 Preliminaries and notation

A theorem proving that a problem is NP-hard does not provide much information about how efficiently (although not polynomially, unless $P = NP$) this problem can be solved. In order to prove some useful complexity lower bounds, we mostly need to rely on some complexity hypothesis that is stronger than “ $P \neq NP$ ”. The *Exponential Time Hypothesis (ETH)* is one of the established and most well-known such complexity hypotheses.

Exponential Time Hypothesis (ETH [24]). *There exists an $\varepsilon < 1$ such that 3SAT cannot be solved in $O(2^{\varepsilon n})$ time, where n is the number of variables in the input 3-CNF formula.*

In addition to formulating ETH, Impagliazzo and Paturi proved the celebrated *Sparsification Lemma* [23], which has the following theorem as a consequence. This result is quite useful for providing lower bounds assuming ETH, as it expresses the running time in terms of the size of the input 3-CNF formula, rather than only the number of its variables.

Theorem 1 ([23]). *3SAT can be solved in time $2^{o(n)}$ if and only if it can be solved in time $2^{o(m)}$ on 3-CNF formulas with n variables and m clauses.*

Given a (static) graph G , we denote by $V(G)$ and $E(G)$ the sets of its vertices and edges, respectively. An edge between two vertices u and v of G is denoted by uv , and in this case u and v are said to be *adjacent* in G . The maximum label assigned by λ to an edge of G , called the *lifetime* of (G, λ) , is denoted by $T(G, \lambda)$, or simply by T when no confusion arises. That is, $T(G, \lambda) = \max\{t \in \lambda(e) : e \in E\}$. For every $i, j \in \mathbb{N}$, where $i \leq j$, we denote $[i, j] = \{i, i+1, \dots, j\}$. Throughout the paper we consider temporal graphs with *finite lifetime* T , and we refer to each integer $t \in [1, T]$ as a *time point* (or *time slot*) of (G, λ) . The *instance* (or *snapshot*) of (G, λ) at time t is the static graph $G_t = (V, E_t)$, where $E_t = \{e \in E : t \in \lambda(e)\}$. For every $i, j \in [1, T]$, where $i \leq j$, we denote by $(G, \lambda)|_{[i, j]}$ the restriction of (G, λ) to the time slots $i, i+1, \dots, j$, i.e. $(G, \lambda)|_{[i, j]}$ is the sequence of the instances G_i, G_{i+1}, \dots, G_j . We assume in the remainder of the paper that every edge of G appears in at least one time slot until T , namely $\bigcup_{t=1}^T E_t = E$.

Although some optimization problems on temporal graphs may be hard to solve in the worst case, an optimal solution may be efficiently computable when the input temporal graph (G, λ) has special properties, i.e. if (G, λ) belongs to a special *temporal graph class* (or *time-varying graph class* [6, 8]). To specify a temporal graph class we can restrict (a) the *underlying topology* G , or (b) the *time-labeling* λ , i.e. the temporal pattern in which the time-labels appear, or both.

Definition 2. Let (G, λ) be a temporal graph and let \mathcal{X} be a class of (static) graphs. If $G \in \mathcal{X}$ then (G, λ) is an \mathcal{X} temporal graph. On the other hand, if $G_i \in \mathcal{X}$ for every $i \in [1, T]$, then (G, λ) is an always \mathcal{X} temporal graph.

In the remainder of the paper we denote by $n = |V|$ and $m = |E|$ the number of vertices and edges of the underlying graph G , respectively, unless otherwise stated. Furthermore, unless otherwise stated, we assume that the labeling λ is arbitrary, i.e. (G, λ) is given with an explicit list of labels for every edge. That is, the *size* of the input temporal graph (G, λ) is $O(|V| + \sum_{t=1}^T |E_t|) = O(n + mT)$. In other cases, where λ is more restricted, e.g. if λ is periodic or follows another specific temporal pattern, there may exist a more succinct representations of the input temporal graph.

For every $u \in V$ and every time slot t , we denote the *appearance of vertex u at time t* by the pair (u, t) . That is, every vertex u has T different appearances (one for each time slot) during the lifetime of (G, λ) . Similarly, for every vertex subset $S \subseteq V$ and every time slot t we denote the *appearance of set S at time t* by (S, t) . With a slight abuse of notation, we write $(S, t) = \bigcup_{v \in S} (v, t)$. A *temporal vertex subset* of (G, λ) is a set $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq t \leq T\}$ of vertex appearances in (G, λ) . Given a temporal vertex subset \mathcal{S} , for every time slot $t \in [1, T]$ we denote by $\mathcal{S}_t = \{(v, t) : (v, t) \in \mathcal{S}\}$ the set of all vertex appearances in \mathcal{S} at the time slot t . Similarly, for any pair of time slots $i, j \in [1, T]$, where $i \leq j$, $\mathcal{S}|_{[i, j]}$ is the restriction of the vertex appearances of \mathcal{S} within the time slots $i, i + 1, \dots, j$. Note that the *cardinality* of the temporal vertex subset \mathcal{S} is $|\mathcal{S}| = \sum_{1 \leq t \leq T} |\mathcal{S}_t|$.

2.1 Temporal Vertex Cover

Let \mathcal{S} be a temporal vertex subset of (G, λ) . Let $e = uv \in E$ be an edge of the underlying graph G and let (w, t) be a vertex appearance in \mathcal{S} . We say that vertex w *covers* the edge e if $w \in \{u, v\}$, i.e. w is an endpoint of e ; in that case, edge e is *covered* by vertex w . Furthermore we say that the vertex appearance (w, t) *temporally covers* the edge e if (i) w covers e and (ii) $t \in \lambda(e)$, i.e. the edge e is *active* during the time slot t ; in that case, edge e is *temporally covered* by the vertex appearance (w, t) . We now introduce the notion of a *temporal vertex cover* and the optimization problem TEMPORAL VERTEX COVER.

Definition 3. Let (G, λ) be a temporal graph. A temporal vertex cover of (G, λ) is a temporal vertex subset $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq t \leq T\}$ of (G, λ) such that every edge $e \in E$ is temporally covered by at least one vertex appearance (w, t) in \mathcal{S} .

TEMPORAL VERTEX COVER (TVC)

Input: A temporal graph (G, λ) .

Output: A temporal vertex cover \mathcal{S} of (G, λ) with the smallest cardinality $|\mathcal{S}|$.

Note that TVC is a natural temporal extension of the problem VERTEX COVER on static graphs. In fact, VERTEX COVER is the special case of TVC where $T = 1$. Thus TVC is clearly NP-complete, as it also trivially belongs to NP.

2.2 Sliding Window Temporal Vertex Cover

In the notion of a temporal vertex cover given in Section 2.1, the optimal solution actually depends on the lifetime T (and thus also on the *size*) of the input temporal graph (G, λ) . On the other hand, in many real-world applications where scalability is important, the lifetime T can be arbitrarily large. In such cases it may not be satisfactory enough that an edge is temporally covered just *once* during the whole lifetime of the temporal graph. Instead, in such cases it makes sense that every edge is temporally covered by some vertex appearance at least once during *every small period* Δ of time, regardless of how

large the lifetime T is. Motivated by this, we introduce in this section a natural *sliding window* variant of the TVC problem, which offers a greater scalability of the solution concept.

For every time slot $t \in [1, T - \Delta + 1]$, we define the *time window* $W_t = [t, t + \Delta - 1]$ as the sequence of the Δ consecutive time slots $t, t + 1, \dots, t + \Delta - 1$. We denote by $\mathcal{W}(T, \Delta) = \{W_1, W_2, \dots, W_{T - \Delta + 1}\}$ the set of all time windows in the lifetime of (G, λ) . Furthermore we denote by $E[W_t] = \bigcup_{i \in W_t} E_i$ the union of all edges appearing at least once in the time window W_t . Finally we denote by $\mathcal{S}[W_t] = \{(v, t) \in \mathcal{S} : t \in W_t\}$ the restriction of the temporal vertex subset \mathcal{S} to the window W_t . We are now ready to introduce the notion of a *sliding Δ -window temporal vertex cover* and the optimization problem SLIDING WINDOW TEMPORAL VERTEX COVER.

Definition 4. *Let (G, λ) be a temporal graph with lifetime T and let $\Delta \leq T$. A sliding Δ -window temporal vertex cover of (G, λ) is a temporal vertex subset $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq t \leq T\}$ of (G, λ) such that, for every time window W_t and for every edge $e \in E[W_t]$, e is temporally covered by at least one vertex appearance (w, t) in $\mathcal{S}[W_t]$.*

SLIDING WINDOW TEMPORAL VERTEX COVER (SW-TVC)

Input: A temporal graph (G, λ) with lifetime T , and an integer $\Delta \leq T$.

Output: A sliding Δ -window temporal vertex cover \mathcal{S} of (G, λ) with the smallest cardinality $|\mathcal{S}|$.

Whenever the parameter Δ is a fixed constant, we will refer to the above problem as the Δ -TVC (i.e. Δ is now a part of the problem name). Note that the problem TVC defined in Section 2.1 is the special case of SW-TVC where $\Delta = T$, i.e. where there is only one Δ -window in the whole temporal graph. Another special case³ of SW-TVC is the problem 1-TVC, whose optimum solution is obtained by iteratively solving the (static) problem VERTEX COVER on each of the T static instances of (G, λ) ; thus 1-TVC fails to fully capture the time dimension in temporal graphs.

2.3 Alternative models

In this section we briefly discuss two alternative variations of the problem SW-TVC, namely the problems FLEXIBLE SW-TVC and the DISJOINT WINDOW TVC. Although both these temporal variations of the (static) problem VERTEX COVER could also be considered as being natural, it turns out that each of these problems is equivalent to iteratively solving either the static VERTEX COVER problem or the TVC problem (defined in Section 2.1). We first introduce the notion of a *flexible sliding Δ -window temporal vertex cover* and the optimization problem FLEXIBLE SW-TVC.

Definition 5. *Let (G, λ) be a temporal graph with lifetime T and let $\Delta \leq T$. A flexible sliding Δ -window temporal vertex cover of (G, λ) is a temporal vertex subset $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq t \leq T\}$ of (G, λ) such that, for every time window W_t and for every edge $e \in E[W_t]$, there exists at least one vertex appearance $(w, t) \in \mathcal{S}[W_t]$ where e is covered by w .*

Note by Definitions 4 and 5 that a sliding Δ -window temporal vertex cover is also a flexible sliding Δ -window temporal vertex cover, but not vice versa. To illustrate the differences between these two notions, consider a temporal graph (G, λ) , a time window W_t , and an edge $e = uv \in E[W_t]$. Furthermore let \mathcal{S} be a temporal vertex subset of (G, λ) . In order for \mathcal{S} to be a *sliding Δ -window temporal vertex cover* (see Definition 4), it is required that \mathcal{S} contains at least one of the vertex appearances (u, s) and (v, s) , where s is a time slot in which e is *active* in W_t . In contrast, in order for \mathcal{S} to be a *flexible sliding Δ -window temporal vertex cover* (see Definition 5), it is required that \mathcal{S} contains at least one of the vertex appearances (u, s) and (v, s) , where s is *any* time slot in W_t , i.e. not necessarily a time slot of W_t in which e is active.

The problem FLEXIBLE SLIDING WINDOW TEMPORAL VERTEX COVER (for short, FLEXIBLE SW-TVC) asks to compute a flexible sliding Δ -window temporal vertex cover \mathcal{S} of (G, λ) with the smallest cardinality $|\mathcal{S}|$. The next lemma shows that FLEXIBLE SW-TVC is equivalent to iteratively solving VERTEX COVER on $\lfloor \frac{T}{\Delta} \rfloor$ static graphs that are easily derived from the input temporal graph (G, λ) , and thus FLEXIBLE SW-TVC fails to fully capture the time dimension in temporal graphs. For the sake of

³The problem 1-TVC has already been investigated under the name “evolving vertex cover” in the context of maintenance algorithms in dynamic graphs [9]; similar “evolving” variations of other graph covering problems have also been considered, e.g. the “evolving dominating set” [7].

the presentation of the next lemma, for every $1 \leq i \leq \lfloor \frac{T}{\Delta} \rfloor$ we define the interval $D_i = [(i-1)\Delta + 1, (i+1)\Delta - 1] \cap [1, T]$.

Lemma 1. *Let (G, λ) be a temporal graph with lifetime T , let $\Delta \leq T$, and let $k = \lfloor \frac{T}{\Delta} \rfloor$. For every $i \in [1, k]$, let S_i be a minimum vertex cover of the (static) graph $\bigcup_{j \in D_i} G_j$, where the union of the graph snapshots is understood as their edge-union. Then $\mathcal{S} = \bigcup_{t=1}^k (S_i, i\Delta)$ is an optimum solution of FLEXIBLE SW-TVC for (G, λ) .*

Proof. First we prove that \mathcal{S} is a flexible sliding Δ -window temporal vertex cover of (G, λ) . Consider an edge e of the underlying graph G such that $t_0 \in \lambda(e)$, i.e. e is active at the time slot $t_0 \in [1, T]$. Then e is active at least once at each of the time windows $W_t \in \mathcal{W}(T, \Delta)$, where $t \in [t_0 - \Delta + 1, t_0] \cap [1, T - \Delta + 1]$. Consider any of these time windows W_t . Note that W_t contains at least one time slot that is a multiple of Δ , say $i \cdot \Delta$, where $1 \leq i \leq k$. Therefore e belongs to the static graph $\bigcup_{j \in D_i} G_j$. Thus, since S_i is a vertex cover of this graph, it follows that there exists at least one vertex appearance $(w, i\Delta) \in \mathcal{S}[W_t]$ where e is covered by w . Therefore, since e is arbitrary by assumption, \mathcal{S} is a flexible sliding Δ -window temporal vertex cover of (G, λ) by Definition 5.

We now prove that \mathcal{S} is an optimum solution of FLEXIBLE SW-TVC. To this end, let \mathcal{S}' be a flexible sliding Δ -window temporal vertex cover of (G, λ) with minimum cardinality. Let $i \in [1, k]$ be arbitrary. Consider the (static) graph H_i (resp. H'_i) which is the edge-union of all snapshots G_j , where $j \in [(i-1)\Delta + 1, i\Delta]$ (resp. $j \in [i\Delta, (i+1)\Delta - 1]$). Let X_i and X'_i be minimum vertex covers of H_i and H'_i , respectively. Then, by Definition 5, \mathcal{S}' must contain at least one vertex occurrence for every vertex in $X_i \cup X'_i$. That is, $|\mathcal{S}'| \geq \sum_{1 \leq i \leq k} |X_i \cup X'_i|$. Consider now the static graph $H''_i = \bigcup_{j \in D_i} G_j$, which is the edge-union of H_i and H'_i , and let X''_i be a vertex cover of H''_i . Then note that $|X_i \cup X'_i| \geq |X''_i|$, and thus $|\mathcal{S}'| \geq \sum_{1 \leq i \leq k} |X''_i| = |\mathcal{S}|$. Therefore \mathcal{S} is an optimum solution of FLEXIBLE SW-TVC. \square

We now introduce the notion of a *disjoint Δ -window temporal vertex cover* and the optimization problem DISJOINT WINDOW TVC.

Definition 6. *Let (G, λ) be a temporal graph with lifetime T , let $\Delta \leq T$, and let $k = \lfloor \frac{T}{\Delta} \rfloor$. A disjoint Δ -window temporal vertex cover of (G, λ) is the union $\bigcup_{i=1}^k \mathcal{S}_i$, where \mathcal{S}_i is a temporal vertex cover of the temporal graph $(G, \lambda)|_{[(i-1)\Delta+1, i\Delta] \cap [1, T]}$.*

The problem DISJOINT WINDOW TEMPORAL VERTEX COVER (for short, DISJOINT WINDOW TVC) asks to compute a disjoint Δ -window temporal vertex cover \mathcal{S} of (G, λ) with the smallest cardinality $|\mathcal{S}|$. As it immediately follows by Definition 6, an optimum solution to DISJOINT WINDOW TVC is obtained by iteratively solving TVC (as defined in Section 2.1) on each of the temporal graphs $(G, \lambda)|_{[(i-1)\Delta+1, i\Delta] \cap [1, T]}$, where $1 \leq i \leq \lfloor \frac{T}{\Delta} \rfloor$. Therefore DISJOINT WINDOW TVC does not provide any additional temporal insights compared to the problems TVC and SW-TVC (cf. Sections 2.1 and 2.2, respectively).

3 Hardness and approximability of TVC

In this section we investigate the complexity of TEMPORAL VERTEX COVER (TVC). First we prove in Section 3.1 that TVC remains NP-complete even on star temporal graphs, i.e. where the underlying graph G is a star. Furthermore we prove that, for any $\varepsilon < 1$, TVC on star temporal graphs cannot be optimally solved in $O(2^{\varepsilon n})$ time, unless the Strong Exponential Time Hypothesis (SETH) fails, as well as that it does not admit a polynomial-time $(1 - \varepsilon) \ln n$ -approximation algorithm, unless NP has $n^{O(\log \log n)}$ -time deterministic algorithms.

In contrast, we prove that TVC on star temporal graphs can be $(H_{n-1} - \frac{1}{2})$ -approximated in polynomial time, where $H_n = \sum_{i=1}^n \frac{1}{i} \approx \ln n$ is the n th harmonic number. In Section 3.2 we use randomized rounding to prove that TVC on general temporal graphs admits a polynomial-time randomized approximation algorithm with expected ratio $O(\ln n)$. This result is complemented by our results in Section 6.1 where we prove that SW-TVC (and thus also TVC) can be deterministically approximated with ratio $H_{2n\Delta} - \frac{1}{2} \approx \ln n + \ln 2\Delta - \frac{1}{2}$ in polynomial time.

3.1 Hardness on star temporal graphs

In the next theorem we prove that TVC on star temporal graphs (i.e. when the underlying graph G is a star, cf. Definition 2) is equivalent to SET COVER, and thus the known (in)approximability results for SET COVER carry over to TVC. This hardness result of TVC is in wide contrast to the (trivial) solution of VERTEX COVER on a static star graph.

Theorem 2. *TVC on star temporal graphs is NP-complete and it admits a polynomial-time $(H_{n-1} - \frac{1}{2})$ -approximation algorithm. Furthermore, for any $\varepsilon > 0$, TVC on star temporal graphs does not admit any polynomial-time $(1 - \varepsilon) \ln n$ -approximation algorithm, unless NP has $n^{O(\log \log n)}$ -time deterministic algorithms.*

Proof. First we reduce SET COVER to TVC on star temporal graphs, and vice versa.

SET COVER

Input: A universe $U = \{1, 2, \dots, n\}$ and a collection of $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ of m subsets of U such that $\bigcup_{i=1}^m C_i = U$.

Output: A subset $\mathcal{C}' \subseteq \mathcal{C}$ with the smallest cardinality such that $\bigcup_{C_i \in \mathcal{C}'} C_i = U$.

Given an instance (U, \mathcal{C}) of SET COVER, we construct an equivalent instance (G, λ) of TVC on star temporal graphs as follows. We set $T = m$ and we let G be a star graph on $n + 1$ vertices, with center c and leaves v_1, v_2, \dots, v_n . The labeling λ is such that, at every time slot $i \in [1, m]$, G_i is a star centered at c and with leaves all the vertices v_j such that $j \in C_i$.

Conversely, let (G, λ) be an instance of TVC on star temporal graphs, where G is a star on $n + 1$ vertices, with center c and leaves v_1, v_2, \dots, v_n . From (G, λ) we construct an equivalent instance (U, \mathcal{C}) of SET COVER as follows. We set $m = T$ and, for every $i = 1, 2, \dots, T$, we define C_i to be the set of all elements j such that $v_j c \in E_i$.

In both reductions, we will now prove that there exists a temporal vertex cover \mathcal{S} in (G, λ) such that $|\mathcal{S}| \leq k$ if and only if there exists a set cover \mathcal{C}' of (U, \mathcal{C}) such that $|\mathcal{C}'| \leq k$.

(\Rightarrow) Let \mathcal{S} be a minimum-cardinality temporal vertex cover of (G, λ) , and let $|\mathcal{S}| \leq k$. Since \mathcal{S} has minimum cardinality and G is a star, it follows that, for every $i = 1, 2, \dots, m$, either $\mathcal{S}_i = \{(c, i)\}$ or $\mathcal{S}_i = \emptyset$. Then the collection $\mathcal{C}' = \{C_i \in \mathcal{C} : \mathcal{S}_i \neq \emptyset\}$ is a set cover of U . Indeed, if $\mathcal{S}_i \neq \emptyset$ then the appearance (c, i) in \mathcal{S} covers all the edges cv_j of G , where $j \in C_i$. Thus, as the sequence of all non-empty sets \mathcal{S}_i covers all edges of (G, λ) , it follows that the union of all sets $C_i \in \mathcal{C}'$ covers all elements of $U = \{1, 2, \dots, n\}$. Finally, since $|\mathcal{S}| \leq k$, it follows by construction that $|\mathcal{C}'| \leq k$.

(\Leftarrow) Let \mathcal{C}' be an optimal solution to SET COVER, and let $|\mathcal{C}'| \leq k$. We define the temporal vertex set $\mathcal{S} = \{(c, i) : C_i \in \mathcal{C}'\}$. For every $C_i \in \mathcal{C}'$, the appearance of (c, i) in \mathcal{S} covers all edges cv_j of G , where $j \in C_i$. Thus, as the sets of \mathcal{C}' cover all elements of $U = \{1, 2, \dots, n\}$, it follows that \mathcal{S} is a temporal vertex cover of (G, λ) . Finally, since $|\mathcal{C}'| \leq k$, it follows by construction that $|\mathcal{S}| \leq k$.

Therefore TVC on star temporal graphs is equivalent to SET COVER, and thus in particular TVC on star temporal graphs is NP-complete.

The polynomial-time greedy algorithm of [10] for SET COVER achieves an approximation ratio of $H_k - \frac{1}{2}$, where k is the maximum size of a set in the input instance and $H_k = \sum_{i=1}^k \frac{1}{i}$ is the k th harmonic number. Therefore, since in an input instance (U, \mathcal{C}) of SET COVER with n elements in the universe U , the maximum size of a set in \mathcal{C} is at most n , the approximation algorithm of [10] has ratio $H_n - \frac{1}{2}$ on such an instance. Furthermore, due to the above polynomial-time reduction of TVC on star temporal graphs to SET COVER, it follows that TVC on star temporal graphs with n vertices can be approximated in polynomial time within ratio $H_{n-1} - \frac{1}{2}$.

Finally, it is known that, for any $\varepsilon > 0$, SET COVER cannot be approximated in polynomial time within a factor of $(1 - \varepsilon) \ln n$ unless NP has $n^{O(\log \log n)}$ -time deterministic algorithms [15]. Therefore, due to the above polynomial-time reduction of SET COVER to TVC on star temporal graphs, it follows that TVC on star temporal graphs does not admit such polynomial-time approximation algorithms as well. \square

In the next theorem we complement our hardness results for TVC on star temporal graphs by reducing HITTING SET to it.

Theorem 3. For every $\varepsilon < 1$, TVC on star temporal graphs cannot be optimally solved in $O(2^{\varepsilon n})$ time, unless the Strong Exponential Time Hypothesis (SETH) fails.

Proof. The proof is done via a reduction of HITTING SET to TVC on star temporal graphs. This reduction has a similar flavor as the one presented in Theorem 2, since the problems SET COVER and HITTING SET are in a sense dual to each other⁴. We first present the definition HITTING SET.

HITTING SET

Input: A universe $U = \{1, 2, \dots, n\}$ and a collection of $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ of m subsets of U such that $\bigcup_{i=1}^m C_i = U$.

Output: A subset $U' \subseteq U$ with the smallest cardinality such that U' contains at least one element from each set in \mathcal{C} .

Given an instance (U, \mathcal{C}) of HITTING SET, we construct an equivalent instance of TVC on star temporal graphs as follows. We set $T = n$ and let G be a star on $m + 1$ vertices with center vertex c and leaves v_1, v_2, \dots, v_m . The labeling λ is such that, at every time slot $i \in [1, n]$, G_i is a star centered at c and with leaves all the vertices v_j such that $i \in C_j$. Following a similar argumentation as in Theorem 2, it follows that there exists a temporal vertex cover \mathcal{S} in (G, λ) such that $|\mathcal{S}| \leq k$ if and only if there exists a hitting set $U' \subseteq U$ of (U, \mathcal{C}) such that $|U'| \leq k$.

Assume now that there exists an $O(2^{\varepsilon T})$ -time algorithm for optimally solving TVC on star temporal graphs, for some $\varepsilon < 1$. Then, due to the above reduction from HITTING SET, we can use this algorithm to optimally solve HITTING SET in $O(2^{\varepsilon n})$ -time. This is a contradiction, unless the Strong Exponential Time Hypothesis (SETH) fails [12]. \square

3.2 A randomized rounding algorithm for TVC

In this section we provide a linear programming relaxation of TVC, and then, with the help of a randomized rounding technique, we construct a feasible solution whose expected size is within a factor of $O(\ln n)$ of the optimal size.

Let (G, λ) be a temporal graph with lifetime T ; let $G = (V, E)$, where $|V| = n$ and $|E| = m$. First we give an integer programming formulation of TVC. We introduce a variable x_{vt} , for every vertex and time slot pair, that will be 1 if the vertex appearance (v, t) is selected in the candidate temporal vertex cover and 0 otherwise. The objective then is to minimize the sum of vertex appearances selected in the candidate temporal vertex cover, while covering every edge at a time slot in which it appears. Since every edge $e = uv$ must be covered by at least one of its endpoints' appearances at a time slot $t \in \lambda(e)$, we require for every edge $e = uv$ the inequality:

$$\sum_{t \in \lambda(e)} (x_{ut} + x_{vt}) \geq 1$$

This inequality constraint together with the integrality constraint for all variables yields the following integer programming formulation of TVC:

$$\begin{aligned} & \text{minimize} && \sum_{v \in V, t \in [T]} x_{vt} \\ & \text{subject to} && \sum_{t \in \lambda(uv)} (x_{ut} + x_{vt}) \geq 1, \quad \text{for all } uv \in E \\ & && x_{vt} \in \{0, 1\}, \quad \text{for all } v \in V, t \in [T] \end{aligned}$$

If Z_{IP}^* is the optimal value of this integer program, then it is not hard to see that $Z_{IP}^* = \text{OPT}$, where OPT is the value of an optimal solution to TVC.

The corresponding linear programming relaxation of this integer program is:

⁴That is seen by observing that an instance of SET COVER can be viewed as an arbitrary bipartite graph, with sets represented by vertices on the left, the universe represented by vertices on the right, and edges representing the inclusion of elements in sets. The task is then to find a minimum-cardinality subset of left-vertices which covers all of the right-vertices. In the HITTING SET problem, the objective is to cover the left-vertices using the smallest number of right vertices. Converting from one problem to the other is therefore achieved by interchanging the two sets of vertices.

$$\begin{aligned}
& \text{minimize} && \sum_{v \in V, t \in [T]} x_{vt} \\
& \text{subject to} && \sum_{t \in \lambda(uv)} (x_{ut} + x_{vt}) \geq 1, \quad \text{for all } uv \in E \\
& && x_{vt} \geq 0, \quad \text{for all } v \in V, t \in [T]
\end{aligned}$$

We could also add the constraints $x_{vt} \leq 1$, for all $v \in V, t \in [T]$ but they would be redundant, since in any optimal solution to the problem all the variables are at most 1; indeed, notice that any $x_{vt} > 1$ can be reduced to $x_{vt} = 1$ without affecting the feasibility of the solution. If Z_{LP}^* is the optimal value of this linear program, then clearly $Z_{LP}^* \leq Z_{IP}^* = \text{OPT}$.

Theorem 4. *There exists a polynomial-time randomized approximation algorithm for TEMPORAL VERTEX COVER with expected approximation factor $O(\ln n)$.*

Proof. We apply randomized rounding to our linear program. Let x^* be an optimal solution to the linear programming relaxation. For every vertex appearance (v, t) we pick (v, t) to our solution (equivalently, we set x_{vt} to 1) with probability x_{vt}^* independently.

We begin by analyzing the probability that a given edge $e = uv$ is covered. Pick an arbitrary edge $e = uv$. Let $\lambda(e) = \{l_1, \dots, l_k\}$, for some $k \geq 1$. For all $i = 1, \dots, k$ denote by α_i the value $1 - x_{ul_i}^*$ and denote by β_i the value $1 - x_{vl_i}^*$. Then, by applying the arithmetic-geometric mean inequality⁵, we see that:

$$\begin{aligned}
Pr[\text{edge } e = uv \text{ is not covered}] &= \prod_{t \in \lambda(e)} (1 - x_{ut}^*) (1 - x_{vt}^*) \\
&= \prod_{i=1}^k (\alpha_i \beta_i) \\
&\leq \left(\frac{1}{2k} \left(\sum_{i=1}^k \alpha_i + \sum_{i=1}^k \beta_i \right) \right)^{2k}. \tag{1}
\end{aligned}$$

Since x^* is an optimal solution to the linear programming relaxation, it holds that:

$$\begin{aligned}
\sum_{t \in \lambda(e)} (x_{ut}^* + x_{vt}^*) &\geq 1 \Leftrightarrow \sum_{i=1}^k (x_{ul_i}^* + x_{vl_i}^*) \geq 1 \Leftrightarrow \sum_{i=1}^k ((1 - \alpha_i) + (1 - \beta_i)) \geq 1 \Leftrightarrow \\
\sum_{i=1}^k \alpha_i + \sum_{i=1}^k \beta_i &\leq 2k - 1.
\end{aligned}$$

Hence, equation (1) becomes:

$$Pr[\text{edge } e = uv \text{ is not covered}] \leq \left(1 - \frac{1}{2k} \right)^{2k} \leq e^{-1}. \tag{2}$$

We now repeat the above random experiment $\gamma \ln n$ times independently, for some constant $\gamma \geq 4$. Let C^* denote the union of all of the sets of selected vertex appearances per experiment, i.e. C^* contains all vertex appearances (v, t) which have been picked in at least one experiment. The probability that there is an edge of the graph that is not covered by C^* is, by the union bound, bounded from above by the product of the number of edges and the probability that a particular edge is not covered by C^* . By equation (2), we get that the probability of a particular edge not being covered by C^* is bounded from above by $e^{-\gamma \ln n}$. Therefore, we have that the probability that C^* is a Temporal Vertex Cover is:

⁵For any nonnegative $\alpha_1, \dots, \alpha_k$, $(\prod_{i=1}^k \alpha_i)^{1/k} \leq \frac{1}{k} \sum_{i=1}^k \alpha_i$

$$\begin{aligned}
Pr[C^* \text{ is a temporal vertex cover}] &= 1 - Pr[\exists e = uv \in E \text{ that is not covered by } C^*] \\
&\geq 1 - n^2 e^{-\gamma \ln n} \\
&= 1 - \frac{1}{n^{\gamma-2}}.
\end{aligned}$$

Consider a particular experiment (out of the $\gamma \ln n$). Let X_{vt} be a Bernoulli random variable that takes value 1 if (v, t) was picked in this experiment, and 0 otherwise. Then $E[\sum_{v,t} X_{vt}] = \sum_{v,t} E[X_{vt}] = \sum_{v,t} x_{vt}^* = Z_{LP}^*$. Now, notice that the expected size of C^* is:

$$E[|C^*|] \leq \gamma \ln n E[\sum_{v,t} X_{vt}] = \gamma \ln n Z_{LP}^* \leq \gamma \ln n Z_{IP}^* = \gamma \ln n \text{ OPT}.$$

If C^* is not a temporal vertex cover, then one may select as a temporal vertex cover the trivial temporal vertex cover, C' , which selects for every edge $e = uv \in E(G)$ an appearance (u, t) or (v, t) at a time slot t with $t \in \lambda(e)$; C' has size at most $m \leq n^2$. Then the expected temporal vertex cover size that we get is:

$$\begin{aligned}
E[\text{size of temporal vertex cover}] &\leq \gamma \ln n \text{ OPT} \left(1 - \frac{1}{n^{\gamma-2}}\right) + n^2 \frac{1}{n^{\gamma-2}} \\
&= \gamma \ln n \text{ OPT} + \frac{1}{n^{\gamma-4}} - \frac{\gamma \ln n}{n^{\gamma-2}} \text{ OPT} \\
&\leq \gamma \ln n \text{ OPT} + 1 \quad (\text{since } \gamma \geq 4).
\end{aligned}$$

Since we can assume that the input temporal graph has at least one edge, we have $\text{OPT} \geq 1$, and so the expected approximation ratio is:

$$\frac{E[\text{size of temporal vertex cover}]}{\text{OPT}} \leq \gamma \ln n + 2. \quad (3)$$

The time complexity of our proposed algorithm is at most $\mathcal{T}_{LP} + nT\gamma \ln n + mT + m = O(\mathcal{T}_{LP} + T(n\gamma \ln n + m))$, where \mathcal{T}_{LP} is the time required to solve the linear programming relaxation of the problem. Indeed, performing a single randomized rounding experiment takes time nT , then verifying if C^* is a temporal vertex cover takes time at most mT and finding a trivial temporal vertex cover C' (if needed) takes time m . □

4 An almost tight algorithm for SW-TVC

In this section we investigate the complexity of SLIDING WINDOW TEMPORAL VERTEX COVER (SW-TVC). First we prove in Section 4.1 a strong lower bound on the complexity of optimally solving this problem on arbitrary temporal graphs. More specifically we prove that, for *any* (arbitrarily growing) functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, there exists a constant $\varepsilon \in (0, 1)$ such that SW-TVC cannot be solved in $f(T) \cdot 2^{\varepsilon n \cdot g(\Delta)}$ time, assuming the Exponential Time Hypothesis (ETH). This ETH-based lower bound turns out to be asymptotically almost tight. In fact, we present in Section 4.2 an exact dynamic programming algorithm for SW-TVC whose running time on an arbitrary temporal graph is $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$, which is asymptotically almost optimal, assuming ETH. In Section 4.3 we prove that our algorithm can be refined so that, when the vertex cover number of each snapshot G_i is bounded by a constant k , the running time becomes $O(T\Delta(n+m) \cdot n^{k(\Delta+1)})$. That is, for small values of Δ (say, when Δ is $O(\log n)$ or $O(\log T)$), this algorithm is polynomial in the input size on temporal graphs with bounded vertex cover number at every slot. Notably, for the class of always star temporal graphs (i.e. a superclass of the star temporal graphs studied in Section 3.1) the running time of the algorithm is $O(T\Delta(n+m) \cdot 2^\Delta)$.

4.1 A complexity lower bound

In the classic textbook NP-hardness reduction from 3SAT to VERTEX COVER (see e.g. [18]), the produced instance of VERTEX COVER is a graph whose number of vertices is linear in the number of variables and clauses of the 3SAT instance. Therefore the next theorem follows by Theorem 1 (for a discussion see also [30]).

Theorem 5. *Assuming ETH, there exists an $\varepsilon_0 < 1$ such that VERTEX COVER cannot be solved in $O(2^{\varepsilon_0 n})$, where n is the number of vertices.*

In the the following theorem we prove a strong ETH-based lower bound for SW-TVC. This lower bound is asymptotically almost tight, as we present in Section 4.2 a dynamic programming algorithm for SW-TVC with running time $O(T\Delta(n+m) \cdot 2^{n\Delta})$, where n and m are the numbers of vertices and edges in the underlying graph G , respectively.

Theorem 6. *For any two (arbitrarily growing) functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, there exists a constant $\varepsilon \in (0, 1)$ such that SW-TVC cannot be solved in $f(T) \cdot 2^{\varepsilon n \cdot g(\Delta)}$ time assuming ETH, where n is the number of vertices in the underlying graph G of the temporal graph.*

Proof. We reduce VERTEX COVER to SW-TVC as follows. Let G be an instance graph of VERTEX COVER with n vertices. We construct the temporal graph (G, λ) with $T = \Delta = 2$ such that $G_1 = G$ and G_2 is an independent set on n vertices. Then, clearly the optimum solutions of VERTEX COVER on G coincide with the optimum solutions of SW-TVC on (G, λ) .

For the sake of contradiction assume that, for every $\varepsilon < 1$, there exists an algorithm \mathcal{A}_ε that solves SW-TVC in $f(T) \cdot 2^{\varepsilon n \cdot g(\Delta)}$ time. Then we optimally solve SW-TVC on the temporal graph (G, λ) by applying the algorithm \mathcal{A}_ε , where $\varepsilon = \frac{\varepsilon_0}{2g(2)}$ and ε_0 is the constant of Theorem 5 for VERTEX COVER. Note that ε is a constant, since both ε_0 and $g(2)$ are constants. Furthermore note that the result of the algorithm is also a minimum vertex cover in the original (static) graph G . The running time of \mathcal{A}_ε is by assumption $f(2) \cdot 2^{\frac{\varepsilon_0}{2g(2)} n \cdot g(\Delta)} = f(2) \cdot 2^{\varepsilon_0 n}$. Therefore, since $f(2)$ is also a constant, the existence of the algorithm \mathcal{A}_ε for SW-TVC implies an algorithm for VERTEX COVER with running time $O(2^{\varepsilon_0 n})$, which is a contradiction, assuming ETH, due to Theorem 5. \square

4.2 An exact dynamic programming algorithm

The main idea of our dynamic programming algorithm for SW-TVC is to scan the temporal graph from left to right with respect to time (i.e. to scan the snapshots G_i increasingly on i), and at every time slot to consider all possibilities for the vertex appearances at the previous Δ time slots. Before we proceed with the presentation and analysis of our algorithm, we start with a simple but useful observation.

Observation 1. *Let (G, λ) be a temporal graph with lifetime T . Let \mathcal{S} be a sliding Δ -window temporal vertex cover of (G, λ) . Then, for every $\Delta \leq t \leq T$, the temporal vertex subset $\mathcal{S}|_{[1,t]} = \{(v, i) \in \mathcal{S} : i \leq t\}$ is a sliding Δ -window temporal vertex cover of $(G, \lambda)|_{[1,t]}$.*

Let (G, λ) be a temporal graph with n vertices and lifetime T , and let $\Delta \leq T$. For every $t = 1, 2, \dots, T - \Delta + 1$ and every Δ -tuple of vertex subsets A_1, \dots, A_Δ of G , we define $f(t; A_1, A_2, \dots, A_\Delta)$ to be the smallest cardinality of a sliding Δ -window temporal vertex cover \mathcal{S} of $(G, \lambda)|_{[1,t+\Delta-1]}$, such that $\mathcal{S}_t = (A_1, t)$, $\mathcal{S}_{t+1} = (A_2, t+1)$, \dots , $\mathcal{S}_{t+\Delta-1} = (A_\Delta, t+\Delta-1)$. If there exists no sliding Δ -window temporal vertex cover \mathcal{S} of $(G, \lambda)|_{[1,t+\Delta-1]}$ with these prescribed vertex appearances in the time slots $t, t+1, \dots, t+\Delta-1$, then we define $f(t; A_1, A_2, \dots, A_\Delta) = \infty$. Note that, once we have computed all possible values of the function $f(\cdot)$, then the optimum solution of SW-TVC on (G, λ) has cardinality

$$\text{OPT}_{\text{SW-TVC}}(G, \lambda) = \min_{A_1, A_2, \dots, A_\Delta \subseteq V} \{f(T - \Delta + 1; A_1, A_2, \dots, A_\Delta)\}. \quad (4)$$

Observation 2. *If the temporal vertex set $\bigcup_{i=1}^{\Delta} (A_i, t+i-1)$ is not a temporal vertex cover of $(G, \lambda)|_{[t,t+\Delta-1]}$ then $f(t; A_1, A_2, \dots, A_\Delta) = \infty$.*

Due to Observation 2 we assume below without loss of generality that $\bigcup_{i=1}^{\Delta} (A_i, t+i-1)$ is a temporal vertex cover of $(G, \lambda)|_{[t,t+\Delta-1]}$. We are now ready to present our main recursion formula in the next lemma.

Lemma 2. Let (G, λ) be a temporal graph, where $G = (V, E)$. Let $2 \leq t \leq T - \Delta + 1$ and let $A_1, A_2, \dots, A_\Delta$ be a Δ -tuple of vertex subsets of the underlying graph G . Suppose that $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$ is a temporal vertex cover of $(G, \lambda)|_{[t, t + \Delta - 1]}$. Then

$$f(t; A_1, A_2, \dots, A_\Delta) = |A_\Delta| + \min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\}. \quad (5)$$

Proof. First consider the case where $\min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\} = \infty$. Assume that $f(t; A_1, A_2, \dots, A_\Delta) \neq \infty$ and let \mathcal{S} be a sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 1]}$, in which the vertex appearances in the the last Δ time slots $t, t + 1, \dots, t + \Delta - 1$ are given by $\mathcal{S}_t = (A_1, t)$, $\mathcal{S}_{t+1} = (A_2, t + 1)$, \dots , $\mathcal{S}_{t+\Delta-1} = (A_\Delta, t + \Delta - 1)$. Then, by Observation 1, $\mathcal{S}|_{[1, t + \Delta - 2]}$ is a sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 2]}$. Moreover, the vertex appearances of $\mathcal{S}|_{[1, t + \Delta - 2]}$ in the the last $\Delta - 1$ time slots $t, t + 1, \dots, t + \Delta - 2$ are given by $\mathcal{S}_t = (A_1, t)$, $\mathcal{S}_{t+1} = (A_2, t + 1)$, \dots , $\mathcal{S}_{t+\Delta-2} = (A_{\Delta-1}, t + \Delta - 2)$. Now let X be the set of vertices of G which are active in $\mathcal{S}|_{[1, t + \Delta - 2]}$ during the time slot $t - 1$. Then note that $f(t - 1; X, A_1, \dots, A_{\Delta-1})$ is upper-bounded by the cardinality of $\mathcal{S}|_{[1, t + \Delta - 2]}$, and thus $f(t - 1; X, A_1, \dots, A_{\Delta-1}) \neq \infty$, which is a contradiction. That is, if $\min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\} = \infty$ then also $f(t; A_1, A_2, \dots, A_\Delta) = \infty$, and in this case the value of $f(t; A_1, A_2, \dots, A_\Delta)$ is correctly computed by (5).

Now consider the case where $\min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\} \neq \infty$, and let $X \subseteq V$ be a vertex subset for which $f(t - 1; X, A_1, \dots, A_{\Delta-1})$ is minimized. Furthermore let \mathcal{S} be a minimum-cardinality sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 2]}$, in which the vertex appearances in the the last Δ time slots $t - 1, t, \dots, t + \Delta - 2$ are given by $\mathcal{S}_{t-1} = (X, t - 1)$, $\mathcal{S}_t = (A_1, t)$, \dots , $\mathcal{S}_{t+\Delta-2} = (A_{\Delta-1}, t + \Delta - 2)$. Then $\mathcal{S} \cup (A_\Delta, t + \Delta - 1)$ is a sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 1]}$, since $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$ is a temporal vertex cover of $(G, \lambda)|_{[t, t + \Delta - 1]}$ by the assumption of the lemma. Thus

$$\begin{aligned} f(t; A_1, A_2, \dots, A_\Delta) &\leq |\mathcal{S} \cup (A_\Delta, t + \Delta - 1)| = |A_\Delta| + |\mathcal{S}| \\ &= |A_\Delta| + \min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\}, \end{aligned} \quad (6)$$

and thus, in particular, $f(t; A_1, A_2, \dots, A_\Delta) \neq \infty$. Now let \mathcal{S}' be a minimum-cardinality sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 1]}$, in which the vertex appearances in the the last Δ time slots $t, t + 1, \dots, t + \Delta - 1$ are given by $\mathcal{S}'_t = (A_1, t)$, $\mathcal{S}'_{t+1} = (A_2, t + 1)$, \dots , $\mathcal{S}'_{t+\Delta-1} = (A_\Delta, t + \Delta - 1)$. Note that $|\mathcal{S}'| = f(t; A_1, A_2, \dots, A_\Delta)$, since \mathcal{S}' has minimum cardinality by assumption. Observation 1 implies that the temporal vertex subset $\mathcal{S}'' = \mathcal{S}'|_{[1, t + \Delta - 2]}$ is a sliding Δ -window temporal vertex cover of the instance $(G, \lambda)|_{[1, t + \Delta - 2]}$, and thus $|\mathcal{S}''| \geq \min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\}$. Furthermore, since $|\mathcal{S}'| = |A_\Delta| + |\mathcal{S}''|$ by construction, it follows that

$$f(t; A_1, A_2, \dots, A_\Delta) = |A_\Delta| + |\mathcal{S}''| \geq |A_\Delta| + \min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\}. \quad (7)$$

Summarizing, equations (6) and (7) imply that $f(t; A_1, A_2, \dots, A_\Delta) = |A_\Delta| + \min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\}$, whenever $\min_{X \subseteq V} \{f(t - 1; X, A_1, \dots, A_{\Delta-1})\} \neq \infty$. This completes the proof of the lemma. \square

Using the recursive computation of Lemma 2, we are now ready to present Algorithm 1 for computing the value of an optimal solution of SW-TVC on a given arbitrary temporal graph (G, λ) . Note that Algorithm 1 can be easily modified such that it also computes the actual optimum solution of SW-TVC (instead of only its optimum cardinality). The proof of correctness and running time analysis of Algorithm 1 are given in the next theorem.

Theorem 7. Let (G, λ) be a temporal graph, where $G = (V, E)$ has n vertices and m edges. Let T be its lifetime and let Δ be the length of the sliding window. Algorithm 1 computes in $O(T\Delta(n + m) \cdot 2^{n(\Delta+1)})$ time the value of an optimal solution of SW-TVC on (G, λ) .

Proof. In its main part (lines 1-9), the algorithm iterates over all time slots $1 \leq t \leq T - \Delta + 1$ and over all vertex subsets $A_1, A_2, \dots, A_\Delta \subseteq V$. Whenever it detects a tuple $(t; A_1, A_2, \dots, A_\Delta)$ such that $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$ is not a temporal vertex cover of $(G, \lambda)|_{[t, t + \Delta - 1]}$, then it sets $f(t; A_1, A_2, \dots, A_\Delta) = \infty$ in line 9. This is correct by Observation 2.

Algorithm 1 SW-TVC

Input: A temporal graph (G, λ) with lifetime T , where $G = (V, E)$, and a natural $\Delta \leq T$.

Output: The smallest cardinality of a sliding Δ -window temporal vertex cover in (G, λ) .

```
1: for  $t = 1$  to  $T - \Delta + 1$  do
2:   for all  $A_1, A_2, \dots, A_\Delta \subseteq V$  do
3:     if  $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$  is a temporal vertex cover of  $(G, \lambda)|_{[t, t+\Delta-1]}$  then
4:       if  $t = 1$  then
5:          $f(t; A_1, A_2, \dots, A_\Delta) \leftarrow \sum_{i=1}^\Delta |A_i|$ 
6:       else
7:          $f(t; A_1, A_2, \dots, A_\Delta) \leftarrow |A_\Delta| + \min_{X \subseteq V} \{f(t-1; X, A_1, \dots, A_{\Delta-1})\}$ 
8:       else
9:          $f(t; A_1, A_2, \dots, A_\Delta) \leftarrow \infty$ 
10: return  $\min_{A_1, \dots, A_\Delta \subseteq V} \{f(T - \Delta + 1; A_1, \dots, A_\Delta)\}$ 
```

For all other tuples $(t; A_1, A_2, \dots, A_\Delta)$, the algorithm distinguishes in lines 4-7 the cases $t = 1$ and $t \geq 2$. If $t \geq 2$ the algorithm recursively computes in line 7 the value of $f(t; A_1, A_2, \dots, A_\Delta)$ using values that have been previously computed. The correctness of this recursive computation follows by Lemma 2. If $t = 1$, then clearly the optimum solution of SW-TVC on $(G, \lambda)|_{[1, \Delta]}$ has value equal to the total number of vertex appearances in the time slots $1, 2, \dots, \Delta$, i.e. $f(1; A_1, A_2, \dots, A_\Delta) = \sum_{i=1}^\Delta |A_i|$, as it is computed in line 5. Finally, the algorithm correctly returns in line 10 the smallest value of $f(T - \Delta + 1; A_1, A_2, \dots, A_\Delta)$ among all possible Δ -tuples $A_1, A_2, \dots, A_\Delta$. The correctness of this step has been discussed above, in equation (4).

With respect to running time, Algorithm 1 iterates for each value $t = 1, 2, \dots, T$ and for each of the $2^{n\Delta}$ different Δ -tuples $A_1, A_2, \dots, A_\Delta \subseteq V$ in lines 1-9. The only non-trivial computations within these lines are in lines 3 and 7. In line 3 the algorithm checks whether $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$ is a temporal vertex cover of $(G, \lambda)|_{[t, t+\Delta-1]}$. This can be done in $O(\Delta(n+m))$ time, where m is the number of edges in the (static) underlying graph G , by simply enumerating all edges that are covered by the vertex appearances in $\bigcup_{i=1}^\Delta (A_i, t + i - 1)$ and comparing their number with the total number of edges that are active at least once in the time window $W_t = [t, t + \Delta - 1]$. On the other hand, to execute line 7 we need at most $O(2^n)$ time for computing the minimum among at most 2^n different known values. Similarly, to execute line 10 we need at most $O(2^{n\Delta})$ time for computing the minimum among at most $2^{n\Delta}$ different known values. Therefore the total running time of Algorithm 1 is upper-bounded by $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$ time. \square

4.3 Always bounded vertex cover number temporal graphs

Let (G, λ) be a temporal graph of lifetime T , and let \mathcal{S} be a minimum-cardinality sliding Δ -window temporal vertex cover of (G, λ) . Note that the minimality of $|\mathcal{S}|$ implies that, for every $i = 1, 2, \dots, T$, $|\mathcal{S}_i|$ is upper-bounded by the (static) vertex cover number of G_i . Therefore, in the recursive relation of Lemma 2, it is enough to only consider subsets $X, A_1, A_2, \dots, A_\Delta \subseteq V$ whose cardinalities are bounded by the vertex cover numbers of the corresponding snapshots. Thus, for small values of Δ (say, when Δ is $O(\log n)$ or $O(\log T)$), Algorithm 1 can be modified to a polynomial time algorithm for the class of always bounded vertex cover number temporal graphs, when Δ . Formally, let k be a constant and let \mathcal{C}_k be the class of graphs with the vertex cover number at most k . The next theorem follows now from the analysis of Theorem 7.

Theorem 8. SW-TVC on always \mathcal{C}_k temporal graphs can be solved in $O(T\Delta(n+m) \cdot n^{k(\Delta+1)})$ time.

In particular, in the special, yet interesting, case of always star temporal graphs (i.e. where every snapshot G_i is a star graph), our search at every step reduces to just one binary choice for each of the previous Δ time windows, of whether to include the central vertex of a star in a snapshot or not. Hence we have the following theorem as a direct implication of Theorem 8.

Theorem 9. SW-TVC on always star temporal graphs can be solved in $O(T\Delta(n+m) \cdot 2^\Delta)$ time.

5 Approximation hardness of 2-TVC

In this section we study the complexity of Δ -TVC where Δ is constant. We start with an intuitive observation that, for every fixed Δ , the problem $(\Delta + 1)$ -TVC is at least as hard as Δ -TVC. Indeed, let \mathcal{A} be an algorithm that computes a minimum-cardinality sliding $(\Delta + 1)$ -window temporal vertex cover of (G, λ) . It is easy to see that a minimum-cardinality sliding Δ -window temporal vertex cover of (G, λ) can also be computed using \mathcal{A} , if we amend the input temporal graph by inserting one edgeless snapshot after every Δ consecutive snapshots of (G, λ) , see Figure 1.




Figure 1: Inserting “empty” time slots to compute a minimum-cardinality sliding Δ -window temporal vertex cover on (G, λ) using algorithm \mathcal{A} for $(\Delta + 1)$ -TVC.

Since the 1-TVC problem is equivalent to solving T instances of VERTEX COVER (on static graphs), the above reduction demonstrates in particular that, for any natural Δ , Δ -TVC is at least as hard as VERTEX COVER. Therefore, if VERTEX COVER is hard for a class \mathcal{X} of static graphs, then Δ -TVC is also hard for the class of always \mathcal{X} temporal graphs. In this section, we show that the converse is not true. Namely, we reveal a class \mathcal{X} of graphs, for which VERTEX COVER can be solved in *linear* time, but 2-TVC is NP-hard on always \mathcal{X} temporal graphs. In fact, we show the even stronger result that 2-TVC is APX-hard (and thus does not admit a PTAS, unless $P = NP$) on always \mathcal{X} temporal graphs⁶.

To prove the main result (in Theorem 10) we start with an auxiliary lemma, showing that VERTEX COVER is APX-hard on the class \mathcal{Y} of graphs which can be obtained from a cubic graph by subdividing every edge exactly 4 times.

Lemma 3. VERTEX COVER is APX-hard on \mathcal{Y} .

Proof. VERTEX COVER is known to be APX-hard on cubic graphs [4]. By a simple reduction we will show that the problem remains APX-hard on the class \mathcal{Y} .

Given a cubic graph G , let $H \in \mathcal{Y}$ be the graph obtained from G by subdividing each edge 4 times. It is well known and can be easily verified that a double subdivision of an edge increases the size of a minimum vertex cover exactly by one. Hence, denoting by $\tau(G)$ and $\tau(H)$ the sizes of minimum vertex covers of G and H , respectively, we have:

$$\tau(H) = \tau(G) + 2m, \quad (8)$$

where m is the number of edges in G .

Further, we will show that for every vertex cover S_H for H we can efficiently obtain a vertex cover S_G for G of size at most $|S_H| - 2m$. To this end we show how to construct S_G from S_H by decreasing the cardinality of S_H by at least two for every edge of G . Initially, we set $S_G = S_H$. Let $uv \in E(G)$ be an edge in G , and let $ua_1, a_1a_2, a_2a_3, a_3a_4, a_4v$ be the edges of H corresponding to the 4-subdivision of uv . Note that at least two of the vertices a_1, a_2, a_3, a_4 belong to S_H . If S_H contains at least three of these vertices, then we remove them from S_G , and add either u or v to S_G if none of them was already in S_G . Otherwise, if exactly two of a_1, a_2, a_3, a_4 belong to S_H , then at least one of u and v must also belong to S_H , and we just remove the two vertices from S_G . After repeating this procedure for every edge of G , we obtain a set S_G that covers all edges of G , and $|S_G| \leq |S_H| - 2m$.

Now suppose, for the sake of contradiction, that there exists a PTAS for VERTEX COVER in \mathcal{Y} . That is, for every $\epsilon > 0$, we can find in polynomial time a vertex cover S_H of H such that $|S_H| \leq (1 + \epsilon)\tau(H)$.

⁶More specifically, 2-TVC on always \mathcal{X} temporal graphs is APX-complete. In fact, since the maximum vertex degree of any static graph in the class \mathcal{X} is at most 3, this problem can be 3-approximated in polynomial time by Algorithm 3 (see also Lemma 7 in Section 6.2), and thus the problem also belongs to the class APX.

By the above discussion, we can find a vertex cover S_G of G such that

$$\begin{aligned}
|S_G| &\leq |S_H| - 2m \\
&\leq (1 + \epsilon)\tau(H) - 2m \\
&= (1 + \epsilon)(\tau(G) + 2m) - 2m \\
&= (1 + \epsilon)\tau(G) + 2\epsilon \cdot m \\
&\leq (1 + \epsilon)\tau(G) + 6\epsilon \cdot \tau(G) \\
&= (1 + 7\epsilon)\tau(G),
\end{aligned} \tag{9}$$

where in the first equality we used (8), and in the last inequality we used the fact that $m \leq 3\tau(G)$, because every vertex in the cubic graph G covers exactly 3 edges. Summarizing, the existence of a PTAS for VERTEX COVER in the class \mathcal{Y} would imply the existence of a PTAS in the class of cubic graphs, which would be a contradiction by [4] unless $P = NP$. \square

Let now \mathcal{X} be the class of graphs whose connected components are induced subgraphs of graph Ψ (see Figure 1). Clearly, VERTEX COVER is linearly solvable on graphs from \mathcal{X} . We will show that 2-TVC is APX-hard on always \mathcal{X} temporal graphs by using a reduction from VERTEX COVER on \mathcal{Y} .




Figure 2: The graph Ψ .

Theorem 10. *2-TVC is APX-hard on always \mathcal{X} temporal graphs.*

Proof. To prove the theorem we will reduce VERTEX COVER on \mathcal{Y} to 2-TVC on always \mathcal{X} temporal graphs. Let $H = (V, E)$ be a graph in \mathcal{Y} . First we will show how to construct an always \mathcal{X} temporal graph (G, λ) of lifetime 2. Then we will prove that the size τ of a minimum vertex cover of H is equal to the size σ of a minimum-cardinality sliding 2-window temporal vertex cover of (G, λ) .

Let $R \subseteq V$ be the set of vertices of degree 3 in H . We define (G, λ) to be a temporal graph of lifetime 2, where snapshot G_1 is obtained from H by removing the edges with both ends being at distance exactly 2 from R , and snapshot $G_2 = H - R$. Figure 3 illustrates the reduction for $H = K_4$.

Let $\mathcal{S} = (S_1, 1) \cup (S_2, 2)$ be an arbitrary sliding 2-window temporal vertex cover of (G, λ) for some $S_1, S_2 \subseteq V$. Since every edge of H belongs to at least one of the graphs G_1 and G_2 , the set $S_1 \cup S_2$ covers all the edges of H . Hence, $\tau \leq |S_1 \cup S_2| \leq |S_1| + |S_2| = |\mathcal{S}|$. As \mathcal{S} was chosen arbitrarily we further conclude that $\tau \leq \sigma$.

To show the converse inequality, let $C \subseteq V$ be a minimum vertex cover of H . Let S_1 be those vertices in C which either have degree 3, or have a neighbor of degree 3. Let also $S_2 = C \setminus S_1$. We claim that $(S_1, 1) \cup (S_2, 2)$ is a sliding 2-window temporal vertex cover of (G, λ) . First, let $e \in E$ be an edge in H incident to a vertex of degree 3. Then, by the construction, e is active only in time slot 1, i.e. $e \in E_1 \setminus E_2$, and a vertex v in C covering e belongs to S_1 . Hence, e is temporally covered by $(v, 1)$ in (G, λ) . Let now $e \in E$ be an edge in H whose both end vertices have degree 2. If one of the end vertices of e is adjacent to a vertex of degree 3 in H , then, by the construction, e is active in both time slots 1 and 2. Therefore, since $C = S_1 \cup S_2$, edge e will be temporally covered in (G, λ) in at least one of the time slots. Finally, if none of the end vertices of e is adjacent to a vertex of degree 3 in H , then e is active only in time slot 2, i.e. $e \in E_2 \setminus E_1$. Moreover, by the construction a vertex v in C covering e belongs to S_2 . Hence, e is temporally covered by $(v, 2)$ in (G, λ) . This shows that $(S_1, 1) \cup (S_2, 2)$ is a sliding 2-window temporal vertex cover of (G, λ) , and therefore $\sigma \leq |S_1| + |S_2| = |C| = \tau$.

Note that the size of a minimum vertex cover of H is equal to the size of a minimum-cardinality sliding 2-window temporal vertex cover of (G, λ) and that any feasible solution to 2-TVC on (G, λ) of size r defines a vertex cover of H of size at most r . Thus, since VERTEX COVER is APX-hard on \mathcal{Y} by Lemma 3 and the reduction is approximation-preserving, it follows that 2-TVC is APX-hard as well. \square




Figure 3: A cubic graph K_4 , its 4-subdivision, and the corresponding snapshots G_1 and G_2

6 Approximation algorithms

In this section we provide several approximation algorithms for SW-TVC with respect to different temporal graph parameters. As the various approximation factors that are achieved are incomparable, the best option for approximating an optimal solution depends on the specific application domain and the specific values of those parameters.

6.1 Approximations in terms of T , Δ , and the largest edge frequency

We begin by presenting a reduction from SW-TVC to SET COVER, which proves useful for deriving approximation algorithms for the original problem. Consider an instance, (G, λ) and $\Delta \leq T$, of the SW-TVC problem. Construct an instance of SET COVER as follows: Let the universe be $U = \{(e, t) : e \in E[W_t], t \in [1, T - \Delta + 1]\}$, i.e. the set of all pairs (e, t) of an edge e and a time slot t such that e appears (and so must be temporally covered) within window W_t . For every vertex appearance (v, s) we define $C_{v,s}$ to be the set of elements (e, t) in the universe U , such that (v, s) temporally covers e in the window W_t . Formally, $C_{v,s} = \{(e, t) : v \text{ is an endpoint of } e, e \in E_s, \text{ and } s \in W_t\}$. Let \mathcal{C} be the family of all sets $C_{v,s}$, where $v \in V, s \in [1, T]$. The following lemma shows that finding a minimum-cardinality sliding Δ -window temporal vertex cover of (G, λ) is equivalent to finding a minimum-cardinality family of sets $C_{v,s}$ that covers the universe U .

Lemma 4. *A family $\mathcal{C} = \{C_{v_1, t_1}, \dots, C_{v_k, t_k}\}$ is a set cover of U if and only if $\mathcal{S} = \{(v_1, t_1), \dots, (v_k, t_k)\}$ is a sliding Δ -window temporal vertex cover of (G, λ) .*

Proof. First assume that \mathcal{C} is a set cover of U , but \mathcal{S} is not a sliding Δ -window temporal vertex cover of (G, λ) . Then there exists a window W_r for some $r \in [1, T - \Delta + 1]$, such that an edge uv appears in W_r but is not temporally covered in W_r by \mathcal{S} . This means that, for every $j \in W_r$ such that $uv \in E_j$, neither (u, j) nor (v, j) belongs to \mathcal{S} . Therefore, $C_{u,j}, C_{v,j} \notin \mathcal{C}$ for all $j \in W_r$ such that $uv \in E_j$. But then \mathcal{C} does not cover $(uv, r) \in U$, which is a contradiction.

Conversely, assume that \mathcal{S} is a sliding Δ -window temporal vertex cover of (G, λ) , but \mathcal{C} is not a set cover of U , i.e. there exists some $(uv, r) \in U$ which belongs to none of the sets in \mathcal{C} . The latter means that $C_{u,j}, C_{v,j} \notin \mathcal{C}$, and therefore $(u, j), (v, j) \notin \mathcal{S}$, for every $j \in W_r$ such that $uv \in E_j$. Therefore uv is not covered in W_r , which is a contradiction. \square

$O(\ln n + \ln \Delta)$ -approximation. In the instance of SET COVER constructed by the above reduction, every set $C_{v,s}$ in \mathcal{C} contains at most $n\Delta$ elements of the universe U . Indeed, the vertex appearance (v, s) temporally covers at most $n - 1$ edges, each in at most Δ windows (namely from window $W_{s-\Delta+1}$ up to window W_s). Thus we can apply the polynomial-time greedy algorithm of [10] for SET COVER which achieves an approximation ratio of $H_{n\Delta} - \frac{1}{2}$, where $n\Delta$ is the maximum size of a set in the input instance and $H_{n\Delta} = \sum_{i=1}^{n\Delta} \frac{1}{i} \approx \ln n + \ln \Delta$ is the $n\Delta$ -th harmonic number.

$2k$ -approximation, where k is the maximum edge frequency. Given a temporal graph (G, λ) and an edge e of G , the Δ -frequency of e is the maximum number of time slots at which e appears within a Δ -window. Let k denote the maximum Δ -frequency over all edges of G . Clearly, for a particular Δ -window W_t , an edge $e \in E[W_t]$ can be temporally covered in W_t by at most $2k$ vertex appearances. So in the above reduction to SET COVER, every element $(e, t) \in U$ belongs to at most $2k$ sets in \mathcal{C} . Therefore, the optimal solution of the constructed instance of SET COVER

can be approximated within a factor of $2k$ in polynomial time [36], yielding a polynomial-time $2k$ -approximation for SW-TVC.

2Δ -approximation. Since the maximum Δ -frequency of an edge is always upper-bounded by Δ , the previous algorithm gives a worst-case polynomial-time 2Δ -approximation for SW-TVC on arbitrary temporal graphs.

6.2 Approximation in terms of maximum degree of snapshots

In this section we give a polynomial-time d -approximation algorithm for the SW-TVC problem on *always degree at most d* temporal graphs, that is, temporal graphs where the maximum degree in each snapshot is at most d . In particular, the algorithm computes an optimum solution (i.e. with approximation ratio $d = 1$) for always matching (i.e. always degree at most 1) temporal graphs. As a building block, we first provide an exact $O(T)$ -time algorithm for optimally solving SW-TVC in the class of single-edge temporal graphs, namely temporal graphs whose underlying graph is a single edge.

Single-edge temporal graphs Consider a temporal graph (G_0, λ) where G_0 is the single-edge graph, i.e. $V(G_0) = \{u, v\}$ and $E(G_0) = \{uv\}$. We reduce SW-TVC on (G_0, λ) to an instance of INTERVAL COVERING, which has a known greedy algorithm that we then translate to an algorithm for SW-TVC on single-edge temporal graphs.

INTERVAL COVERING

Input: A family \mathcal{I} of intervals in the line.

Output: A minimum-cardinality subfamily $\mathcal{I}' \subseteq \mathcal{I}$ such that $\bigcup_{I \in \mathcal{I}} I = \bigcup_{I \in \mathcal{I}'} I$.

We construct the family \mathcal{I} as follows. For every $i = 1, 2, \dots, T$ such that $uv \in E_i$ we include into \mathcal{I} the interval $I_i = [i - \Delta + 1, i] \cap [1, T - \Delta + 1]$, which contains the first time slot of all those Δ -windows that include time slot i .

Lemma 5. *Let i_1, i_2, \dots, i_k be such that $uv \in E_{i_j}$ for every $j = 1, 2, \dots, k$. Then $\mathcal{I}' = \{I_{i_1}, \dots, I_{i_k}\}$ is an interval covering of \mathcal{I} if and only if $\mathcal{S} = \{(u, i_1), (u, i_2), \dots, (u, i_k)\}$ is a sliding Δ -window temporal vertex cover of (G_0, λ) .*

Proof. Assume first that \mathcal{I}' is an interval covering of \mathcal{I} , but \mathcal{S} is not a sliding Δ -window temporal vertex cover of (G_0, λ) . The latter means that there exists a Δ -window W_t such that uv exists at some time slot s in W_t , but uv is not temporally covered by any vertex appearance in $\mathcal{S}[W_t]$. Therefore $t \notin I_{i_1} \cup I_{i_2} \cup \dots \cup I_{i_k}$, but $t \in I_s$, which contradicts the assumption that \mathcal{I}' is an interval covering of \mathcal{I} .

Conversely, assume that \mathcal{S} a sliding Δ -window temporal vertex cover of (G_0, λ) , but \mathcal{I}' is not an interval covering of \mathcal{I} , that is, there exists $I_i \in \mathcal{I}$ and $t \in I_i$ such that $t \notin \bigcup_{I \in \mathcal{I}'} I$. By the construction, this means that $uv \in E_i$ is not temporally covered by any vertex appearances in $\mathcal{S}[W_t]$, which is a contradiction. \square

Lemma 5 shows that finding a minimum-cardinality sliding Δ -window temporal vertex cover of (G_0, λ) is equivalent to finding a minimum interval covering of \mathcal{I} . An easy linear-time greedy algorithm for the INTERVAL COVERING picks at each iteration, among the intervals that cover the leftmost uncovered point, the one with largest finishing time. Algorithm 2 implements this simple rule in the context of the SW-TVC problem.

Lemma 6. *Algorithm 2 solves SW-TVC on a single-edge temporal graph and can be implemented to work in time $O(T)$.*

Proof. The time complexity of the algorithm is dominated by the running time of the while-loop. We provide an implementation of the while-loop, which works in time $O(T)$: in each iteration, we inspect the current Δ -window $[t, t + \Delta - 1]$ from the rightmost time slot moving to the left. As we go through the time slots, we mark the ones in which edge uv does not appear as “NO”. When we move to the next iteration (the next Δ -window), we do not need to revisit any time slots that have been marked as “NO” and we immediately move to the next iteration whenever we meet such a slot. This way we visit every time slot at most once, and hence we exit the while-loop after $O(T)$ operations. \square

Algorithm 2 SW-TVC on single-edge temporal graphs

Input: A temporal graph (G_0, λ) of lifetime T with $V(G_0) = \{u, v\}$, and $\Delta \leq T$.

Output: A minimum-cardinality sliding Δ -window temporal vertex cover \mathcal{S} of (G_0, λ) .

```
1:  $\mathcal{S} \leftarrow \emptyset$ 
2:  $t = 1$ 
3: while  $t \leq T - \Delta + 1$  do
4:   if  $\exists r \in [t, t + \Delta - 1]$  such that  $uv \in E_t$  then
5:     choose maximum such  $r$  and add  $(u, r)$  to  $\mathcal{S}$ 
6:      $t \leftarrow r + 1$ 
7:   else
8:      $t \leftarrow t + 1$ 
9: return  $\mathcal{S}$ 
```

Always degree at most d temporal graphs We present now the main algorithm of this section, the idea of which is to independently solve SW-TVC for every possible single-edge temporal subgraph of a given temporal graph by Algorithm 2, and take the union of these solutions. We will show that this algorithm is a d -approximation algorithm for SW-TVC on always degree at most d temporal graphs.

Let (G, λ) be a temporal graph, where $G = (V, E)$, $|V| = n$, and $|E| = m$. For every edge $e = uv \in E$, let $(G[\{u, v\}], \lambda)$ denote the temporal graph where the underlying graph is the induced subgraph $G[\{u, v\}]$ of G and the labels of e are exactly the same as in (G, λ) .

Algorithm 3 d -approximation of SW-TVC on always degree at most d temporal graphs

Input: An always degree at most d temporal graph (G, λ) of lifetime T , and $\Delta \leq T$.

Output: A sliding Δ -window temporal vertex cover \mathcal{S} of (G, λ) .

```
1: for  $i = 1$  to  $T$  do
2:    $\mathcal{S}_i \leftarrow \emptyset$ 
3:   for every edge  $e = uv \in E(G)$  do
4:     Compute an optimal solution  $\mathcal{S}'(uv)$  of the problem for  $(G[\{u, v\}], \lambda)$  by Algorithm 2
5:     for  $i = 1$  to  $T$  do
6:        $\mathcal{S}_i \leftarrow \mathcal{S}_i \cup \mathcal{S}'_i(uv)$ 
7: return  $\mathcal{S}$ 
```

Lemma 7. *Algorithm 3 is a $O(mT)$ -time d -approximation algorithm for SW-TVC on always degree at most d temporal graphs.*

Proof. Let (G, λ) be an always degree at most d temporal graph of lifetime T , and let \mathcal{S}^* be a minimum-cardinality sliding Δ -window temporal vertex cover of (G, λ) . We will show that $|\mathcal{S}| \leq d \cdot |\mathcal{S}^*|$. To this end we apply a double counting argument to the set C of all triples $(v, e, t) \in V \times E \times [1, T]$ such that $v \in \mathcal{S}_t^*$, $e \in E_t$, and v is incident to e .

On the one hand

$$|C| = \sum_{t=1}^T \sum_{v \in \mathcal{S}_t^*} |\{(v, e, t) : e \in E_t \text{ and } v \text{ is incident to } e\}| \leq \sum_{t=1}^T \sum_{v \in \mathcal{S}_t^*} d = d \cdot |\mathcal{S}^*|,$$

where the inequality follows from the assumption that every snapshot of (G, λ) has maximum degree at most d .

On the other hand

$$|C| = \sum_{e \in E} \sum_{t=1}^T |\{(v, e, t) : e \in E_t, v \in \mathcal{S}_t^*, \text{ and } v \text{ is incident to } e\}| \geq \sum_{e \in E} |\mathcal{S}'(e)| = |\mathcal{S}|,$$

where the inequality follows from the fact that the restriction of any sliding Δ -window temporal vertex cover of (G, λ) to the temporal subgraph induced by the endpoints of e is a sliding Δ -window temporal vertex cover of the temporal subgraph, and therefore has cardinality at least $|\mathcal{S}'|$.

We conclude that $|\mathcal{S}| \leq |C| \leq d \cdot |\mathcal{S}^*|$, as required.

The time-complexity of Algorithm 3 is dominated by the time needed to execute the for-loop of lines 3-6. The latter requires time $O(mT)$. \square

Note that in the case of always matching temporal graphs, i.e. where every snapshot is a matching, the maximum degree in each snapshot is $d = 1$, so the above d -approximation actually yields an exact algorithm.

Corollary 1. *SW-TVC can be optimally solved in $O(mT)$ time on the class of always matching temporal graphs.*

References

- [1] E. Aaron, D. Krizanc, and E. Meyerson. DMVP: foremost waypoint coverage of time-varying graphs. In *Proceedings of the 40th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 29–41, 2014.
- [2] E. C. Akrida, L. Gasieniec, G. B. Mertzios, and P. G. Spirakis. Ephemeral networks with random availability of links: The case of fast networks. *Journal of Parallel and Distributed Computing*, 87:109–120, 2016.
- [3] E. C. Akrida, L. Gasieniec, G. B. Mertzios, and P. G. Spirakis. The complexity of optimal design of temporally connected graphs. *Theory of Computing Systems*, 61(3):907–944, 2017.
- [4] P. Alimonti and V. Kann. Some APX-completeness results for cubic graphs. *Theoretical Computer Science*, 237(1-2):123–134, 2000.
- [5] B.-M. Bui-Xuan, A. Ferreira, and A. Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *International Journal of Foundations of Computer Science*, 14(2):267–285, 2003.
- [6] A. Casteigts and P. Flocchini. Deterministic Algorithms in Dynamic Networks: Formal Models and Metrics. Technical report, Defence R&D Canada, April 2013.
- [7] A. Casteigts and P. Flocchini. Deterministic Algorithms in Dynamic Networks: Problems, Analysis, and Algorithmic Tools. Technical report, Defence R&D Canada, April 2013.
- [8] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems (IJPEDS)*, 27(5):387–408, 2012.
- [9] A. Casteigts, B. Mans, and L. Mathieson. On the feasibility of maintenance algorithms in dynamic graphs. *CoRR*, abs/1107.2722, 2011.
- [10] R. chii Duh and M. Fürer. Approximation of k -set cover by semi-local optimization. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 256–264, 1997.
- [11] A. E. F. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time of edge-markovian evolving graphs. *SIAM Journal on Discrete Mathematics (SIDMA)*, 24(4):1694–1712, 2010.
- [12] M. Cygan, H. Dell, D. Lokshtanov, D. Marx, J. Nederlof, Y. Okamoto, R. Paturi, S. Saurabh, and M. Wahlström. On problems as hard as CNF-SAT. *ACM Transactions on Algorithms*, 12(3):41:1–41:24, 2016.
- [13] S. Dobrev, S. Durocher, M. E. Hesari, K. Georgiou, E. Kranakis, D. Krizanc, L. Narayanan, J. Opatrny, S. M. Shende, and J. Urrutia. Complexity of barrier coverage with relocatable sensors in the plane. *Theoretical Computer Science*, 579:64–73, 2015.
- [14] T. Erlebach, M. Hoffmann, and F. Kammer. On temporal graph exploration. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 444–455, 2015.

- [15] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [16] A. Ferreira. Building a reference combinatorial model for MANETs. *IEEE Network*, 18(5):24–29, 2004.
- [17] P. Flocchini, B. Mans, and N. Santoro. Exploration of periodically varying graphs. In *Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC)*, pages 534–543, 2009.
- [18] M. R. Garey and D. S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- [19] G. Giakkoupis, T. Sauerwald, and A. Stauffer. Randomized rumor spreading in dynamic graphs. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 495–507, 2014.
- [20] M. E. Hesari, E. Kranakis, D. Krizanc, O. M. Ponce, L. Narayanan, J. Opatrny, and S. M. Shende. Distributed algorithms for barrier coverage using relocatable sensors. *Distributed Computing*, 29(5):361–376, 2016.
- [21] A. Himmel, H. Molter, R. Niedermeier, and M. Sorge. Adapting the bron-kerbosch algorithm for enumerating maximal cliques in temporal graphs. *Social Network Analysis and Mining*, 7(1):35:1–35:16, 2017.
- [22] P. Holme and J. Saramäki, editors. *Temporal Networks*. Springer, 2013.
- [23] R. Impagliazzo and R. Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- [24] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
- [25] D. Kempe, J. M. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of the 32nd annual ACM symposium on Theory of computing (STOC)*, pages 504–513, 2000.
- [26] O. Kostakis and A. Gionis. On mining temporal patterns in dynamic graphs, and other unrelated problems. In *Proceedings of the 6th International Conference on Complex Networks and Their Applications (COMPLEX NETWORKS)*, pages 516–527, 2017.
- [27] O. Kostakis, N. Tatti, and A. Gionis. Discovering recurring activity in temporal networks. *Data Mining and Knowledge Discovery*, 31(6):1840–1871, 2017.
- [28] E. Kranakis, D. Krizanc, F. L. Luccio, and B. Smith. Maintaining intruder detection capability in a rectangular domain with sensors. In *Proceedings of the 11th International Symposium on Algorithms and Experiments for Wireless Sensor Networks (ALGOSENSORS)*, pages 27–40, 2015.
- [29] J. Leskovec, J. M. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data*, 1(1), 2007.
- [30] D. Lokshtanov, D. Marx, and S. Saurabh. Lower bounds based on the Exponential Time Hypothesis. *Bulletin of the EATCS*, pages 41–71, 2011.
- [31] G. B. Mertzios, O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Temporal network optimization subject to connectivity constraints. In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP), Part II*, pages 657–668, 2013.
- [32] O. Michail and P. G. Spirakis. Traveling salesman problems in temporal graphs. *Theoretical Computer Science*, 634:1–23, 2016.
- [33] O. Michail and P. G. Spirakis. Elements of the theory of dynamic networks. *Communications of the ACM*, 61(2):72–72, Jan. 2018.

- [34] S. E. Nikolettseas and P. G. Spirakis. Probabilistic distributed algorithms for energy efficient routing and tracking in wireless sensor networks. *Algorithms*, 2(1):121–157, 2009.
- [35] J. K. Tang, M. Musolesi, C. Mascolo, and V. Latora. Characterising temporal distance and reachability in mobile and online social networks. *ACM Computer Communication Review*, 40(1):118–124, 2010.
- [36] V. V. Vazirani. *Approximation algorithms*. Springer, 2003.
- [37] J. Viard, M. Latapy, and C. Magnien. Revealing contact patterns among high-school students using maximal cliques in link streams. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1517–1522, 2015.
- [38] T. Viard, M. Latapy, and C. Magnien. Computing maximal cliques in link streams. *Theoretical Computer Science*, 609:245–252, 2016.
- [39] C. Zhu, C. Zheng, L. Shu, and G. Han. A survey on coverage and connectivity issues in wireless sensor networks. *Journal of Network and Computer Applications*, 35(2):619–632, 2012.