

1 Temporal Vertex Cover with a Sliding Time 2 Window*

3 **Eleni C. Akrida**

4 Department of Computer Science, University of Liverpool, UK

5 Eleni.Akrida@liverpool.ac.uk

6  <https://orcid.org/0000-0002-1126-1623>

7 **George B. Mertzios**

8 Department of Computer Science, Durham University, UK

9 George.Mertzios@durham.ac.uk

10  <https://orcid.org/0000-0001-7182-585X>

11 **Paul G. Spirakis**

12 Department of Computer Science, University of Liverpool, UK

13 Department of Computer Engineering & Informatics, University of Patras, Greece

14 P.Spirakis@liverpool.ac.uk

15  <https://orcid.org/0000-0001-5396-3749>

16 **Viktor Zamaraev**

17 Department of Computer Science, Durham University, UK

18 Viktor.Zamaraev@durham.ac.uk

19  <https://orcid.org/0000-0001-5755-4141>

20 — Abstract —

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

40 **2012 ACM Subject Classification** Mathematics of computing → Graph theory • Mathematics
41 of computing → Graph algorithms.

* 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.



© Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis and Viktor Zamaraev;
licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Don Sannella; Article No. 467,
pp. 467:1–467:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



42 **Keywords and phrases** Temporal networks, temporal vertex cover, APX-hard, approximation
 43 algorithm, Exponential Time Hypothesis.

44 **Digital Object Identifier** 10.4230/LIPIcs.ICALP.2018.467

45 **Related Version** <https://arxiv.org/abs/1802.07103>

46 **Funding** Partially supported by NeST initiative of the School of EEE and CS at the University
 47 of Liverpool and by the EPSRC Grants EP/P020372/1 and EP/P02002X/1.

48 **1 Introduction and Motivation**

49 A great variety of both modern and traditional networks are inherently dynamic, in the sense
 50 that their link availability varies over time. Information and communication networks, social
 51 networks, transportation networks, and several physical systems are only a few examples of
 52 networks that change over time [18, 27]. The common characteristic in all these application
 53 areas is that the network structure, i.e. the underlying graph topology, is subject to *discrete*
 54 *changes over time*. In this paper we adopt a simple and natural model for time-varying
 55 networks which is given with time-labels on the edges of a graph, while the vertex set remains
 56 unchanged. This formalism originates in the foundational work of Kempe et al. [20].

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

60 For every edge $e \in E$ in the underlying graph G of a temporal graph (G, λ) , $\lambda(e)$ denotes
 61 the set of time slots at which e is *active* in (G, λ) . Due to its vast applicability in many areas,
 62 this notion of temporal graphs has been studied from different perspectives under various
 63 names such as *time-varying* [1, 14, 29], *evolving* [4, 10, 13], *dynamic* [7, 15], and *graphs over*
 64 *time* [24]; for a recent attempt to integrate existing models, concepts, and results from the
 65 distributed computing perspective see the survey papers [5–7] and the references therein.
 66 Data analytics on temporal networks have also been very recently studied in the context
 67 of summarizing networks that represent sports teams’ activity data to discover recurring
 68 strategies and understand team tactics [22], as well as extracting patterns from interactions
 69 between groups of entities in a social network [21].

70 Motivated by the fact that, due to causality, information in temporal graphs can “flow” only
 71 along sequences of edges whose time-labels are increasing, most temporal graph parameters
 72 and optimization problems that have been studied so far are based on the notion of temporal
 73 paths and other “path-related” notions, such as temporal analogues of distance, diameter,
 74 reachability, exploration, and centrality [2, 3, 12, 25, 26]. In contrast, only few attempts have
 75 been made to define “non-path” temporal graph problems. Motivated by the contact patterns
 76 among high-school students, Viard et al. [31, 32], and later Himmel et al. [17], introduced
 77 and studied Δ -cliques, an extension of the concept of cliques to temporal graphs, in which
 78 all vertices interact with each other at least once every Δ consecutive time steps within a
 79 given time interval.

80 In this paper we introduce and study two natural temporal extensions of the problem
 81 VERTEX COVER in static graphs, which take into account the dynamic nature of the network.
 82 In the first and simpler of these extensions, namely TEMPORAL VERTEX COVER (for short,
 83 TVC), every edge e has to be “covered” at least once during the lifetime T of the network
 84 (by one of its endpoints), and this must happen at a time step t when e is active. The goal is

85 then to cover all edges with the minimum total number of such “vertex appearances”. On the
 86 other hand, in many real-world applications where scalability is important, the lifetime T can
 87 be arbitrarily large but the network still needs to remain sufficiently covered. In such cases,
 88 as well as in safety-critical systems (e.g. in military applications), it may not be satisfactory
 89 enough that an edge is covered just *once* during the *whole lifetime* of the network. Instead,
 90 every edge must be covered at least once within *every small Δ -window* of time (for an
 91 appropriate value of Δ), regardless of how large the lifetime is; this gives rise to our second
 92 optimization problem, namely SLIDING WINDOW TEMPORAL VERTEX COVER (for short,
 93 SW-TVC). Formal definitions of our problems TVC and SW-TVC are given in Section 2.

94 Our two temporal extensions of VERTEX COVER are motivated by applications in sensor
 95 networks and in transportation networks. In particular, several works in the field of sensor
 96 networks considered problems of placing sensors to cover a whole area or multiple critical
 97 locations, e.g. for reasons of surveillance. Such studies usually wish to minimize the number
 98 of sensors used or the total energy required [11, 16, 23, 28, 33]. Our temporal vertex cover
 99 notions are an abstract way to economically meet such covering demands as time progresses.

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

111 1.1 Our contribution

112 In this paper we present a thorough investigation of the complexity and approximability of the
 113 problems TEMPORAL VERTEX COVER (TVC) and SLIDING WINDOW TEMPORAL VERTEX
 114 COVER (SW-TVC) on temporal graphs. We first prove in Section 3 that SET COVER is
 115 equivalent to a special case of TVC on star temporal graphs (i.e. when the underlying graph
 116 G is a star), which immediately provides several complexity and algorithmic consequences for
 117 TVC. In particular, TVC remains NP-complete even on star temporal graphs, and it does not
 118 admit a polynomial-time $(1 - \varepsilon) \ln n$ -approximation algorithm, unless NP has $n^{O(\log \log n)}$ -time
 119 deterministic algorithms. On the positive side, TVC on star temporal graphs with n vertices
 120 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
 121 harmonic number. Similar equivalence with HITTING SET yields that for any $\varepsilon < 1$, TVC
 122 on star temporal graphs cannot be optimally solved in $O(2^{\varepsilon n})$ time, assuming the Strong
 123 Exponential Time Hypothesis (SETH). We complement these results by showing that TVC

¹ According 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”.

² According 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”.

124 on general temporal graphs admits a polynomial-time randomized approximation algorithm
125 with expected ratio $O(\ln n)$.

126 In Section 4 and in the remainder of the paper we deal with our second problem, SW-TVC.
127 We prove in Section 4.1 a strong complexity lower bound on arbitrary temporal graphs. More
128 specifically we prove that, for *any* (arbitrarily growing) functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$,
129 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,
130 assuming the Exponential Time Hypothesis (ETH). This ETH-based lower bound turns out
131 to be asymptotically almost tight, as we present an exact dynamic programming algorithm
132 with running time $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$. This worst-case running time can be significantly
133 improved in certain special temporal graph classes. In particular, when the “snapshot” of
134 (G, λ) at every time step has vertex cover number bounded by k , the running time becomes
135 $O(T\Delta(n+m) \cdot n^{k(\Delta+1)})$. That is, when Δ is a constant, this algorithm is polynomial in
136 the input size on temporal graphs with bounded vertex cover number at every time step.
137 Notably, when every snapshot is a star (i.e. a superclass of the star temporal graphs studied
138 in Section 3) the running time of the algorithm is $O(T\Delta(n+m) \cdot 2^\Delta)$.

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

152 **2 Preliminaries and notation**

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

158 ► **Exponential Time Hypothesis (ETH [19]).** There exists an $\varepsilon < 1$ such that 3SAT cannot be
159 solved in $O(2^{\varepsilon n})$ time, where n is the number of variables in the input 3-CNF formula.

160 Given a (static) graph G , we denote by $V(G)$ and $E(G)$ the sets of its vertices and edges,
161 respectively. An edge between two vertices u and v of G is denoted by uv , and in this case
162 u and v are said to be *adjacent* in G . The maximum label assigned by λ to an edge of G ,
163 called the *lifetime* of (G, λ) , is denoted by $T(G, \lambda)$, or simply by T when no confusion arises.
164 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
165 $[i, j] = \{i, i+1, \dots, j\}$. Throughout the paper we consider temporal graphs with *finite*
166 *lifetime* T , and we refer to each integer $t \in [1, T]$ as a *time slot* of (G, λ) . The *instance* (or
167 *snapshot*) of (G, λ) at time t is the static graph $G_t = (V, E_t)$, where $E_t = \{e \in E : t \in \lambda(e)\}$.
168 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
169 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

170 assume in the remainder of the paper that every edge of G appears in at least one time slot
 171 until T , namely $\bigcup_{t=1}^T E_t = E$.

172 Although some optimization problems on temporal graphs may be hard to solve in the
 173 worst case, an optimal solution may be efficiently computable when the input temporal
 174 graph (G, λ) has special properties, i.e. if (G, λ) belongs to a special *temporal graph class*
 175 (or *time-varying graph class* [5, 7]). To specify a temporal graph class we can restrict (a)
 176 the *underlying topology* G , or (b) the *time-labeling* λ , i.e. the temporal pattern in which the
 177 time-labels appear, or both.

178 ► **Definition 2.** Let (G, λ) be a temporal graph and let \mathcal{X} be a class of (static) graphs.
 179 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
 180 $i \in [1, T]$, then (G, λ) is an *always \mathcal{X} temporal graph*.

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

188 For every $u \in V$ and every time slot t , we denote the *appearance of vertex u at time t* by
 189 the pair (u, t) . That is, every vertex u has T different appearances (one for each time slot)
 190 during the lifetime of (G, λ) . Similarly, for every vertex subset $S \subseteq V$ and every time slot t we
 191 denote the *appearance of set S at time t* by (S, t) . With a slight abuse of notation, we write
 192 $(S, t) = \bigcup_{v \in S} (v, t)$. A *temporal vertex subset* of (G, λ) is a set $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq t \leq T\}$
 193 of vertex appearances in (G, λ) . Given a temporal vertex subset \mathcal{S} , for every time slot
 194 $t \in [1, T]$ we denote by $\mathcal{S}_t = \{(v, t) : (v, t) \in \mathcal{S}\}$ the set of all vertex appearances in \mathcal{S} at
 195 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
 196 restriction of the vertex appearances of \mathcal{S} within the time slots $i, i + 1, \dots, j$. Note that the
 197 *cardinality* of the temporal vertex subset \mathcal{S} is $|\mathcal{S}| = \sum_{1 \leq t \leq T} |\mathcal{S}_t|$.

198 2.1 Temporal Vertex Cover

199 Let \mathcal{S} be a temporal vertex subset of (G, λ) . Let $e = uv \in E$ be an edge of the underlying
 200 graph G and let (w, t) be a vertex appearance in \mathcal{S} . We say that vertex w *covers* the edge e if
 201 $w \in \{u, v\}$, i.e. w is an endpoint of e ; in that case, edge e is *covered* by vertex w . Furthermore
 202 we say that the vertex appearance (w, t) *temporally covers* the edge e if (i) w covers e and
 203 (ii) $t \in \lambda(e)$, i.e. the edge e is *active* during the time slot t ; in that case, edge e is *temporally*
 204 *covered* by the vertex appearance (w, t) . We now introduce the notion of a *temporal vertex*
 205 *cover* and the optimization problem TEMPORAL VERTEX COVER.

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

TEMPORAL VERTEX COVER (TVC)

209 **Input:** A temporal graph (G, λ) .

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

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

213 2.2 Sliding Window Temporal Vertex Cover

214 In the notion of a temporal vertex cover given in Section 2.1, the requirement is that every
 215 edge is temporally covered at least once during the lifetime T of the input temporal graph
 216 (G, λ) . On the other hand, in many real-world applications where scalability is important,
 217 the lifetime T can be arbitrarily large. In such cases it may not be satisfactory enough that
 218 an edge is temporally covered just *once* during the whole lifetime of the temporal graph.
 219 Instead, in such cases it makes sense that every edge is temporally covered by some vertex
 220 appearance at least once during *every small period* Δ of time, regardless of how large the
 221 lifetime T is. Motivated by this, we introduce in this section a natural *sliding window* variant
 222 of the TVC problem, which offers a greater scalability of the solution concept.

223 For every time slot $t \in [1, T - \Delta + 1]$, we define the *time window* $W_t = [t, t + \Delta - 1]$ as
 224 the sequence of the Δ consecutive time slots $t, t + 1, \dots, t + \Delta - 1$. Furthermore we denote
 225 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 .
 226 Finally we denote by $\mathcal{S}[W_t] = \{(v, i) \in \mathcal{S} : i \in W_t\}$ the restriction of the temporal vertex
 227 subset \mathcal{S} to the window W_t . We are now ready to introduce the notion of a *sliding Δ -window*
 228 *temporal vertex cover* and the optimization problem SLIDING WINDOW TEMPORAL VERTEX
 229 COVER.

230 ► **Definition 4.** Let (G, λ) be a temporal graph with lifetime T and let $\Delta \leq T$. A *sliding*
 231 *Δ -window temporal vertex cover* of (G, λ) is a temporal vertex subset $\mathcal{S} \subseteq \{(v, t) : v \in V, 1 \leq$
 232 $t \leq T\}$ of (G, λ) such that, for every time window W_t and for every edge $e \in E[W_t]$, e is
 233 *temporally covered* by at least one vertex appearance (w, t) in $\mathcal{S}[W_t]$.

SLIDING WINDOW TEMPORAL VERTEX COVER (SW-TVC)

234 **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 $|\mathcal{S}|$.

235 Whenever the parameter Δ is a fixed constant, we will refer to the above problem as the
 236 Δ -TVC (i.e. Δ is now a part of the problem name). Note that the problem TVC defined
 237 in Section 2.1 is the special case of SW-TVC where $\Delta = T$, i.e. where there is only one
 238 Δ -window in the whole temporal graph. Another special case³ of SW-TVC is the problem
 239 1-TVC, whose optimum solution is obtained by iteratively solving the (static) problem
 240 VERTEX COVER on each of the T static instances of (G, λ) ; thus 1-TVC fails to fully capture
 241 the time dimension in temporal graphs.

242 3 Hardness and approximability of TVC

243 In this section we investigate the complexity of TEMPORAL VERTEX COVER (TVC). First
 244 we prove in Section 3.1 that TVC on star temporal graphs is equivalent to both SET COVER
 245 and HITTING SET, and derive several complexity and algorithmic consequences for TVC.

³ The problem 1-TVC has already been investigated under the name “evolving vertex cover” in the
 context of maintenance algorithms in dynamic graphs [8]; similar “evolving” variations of other graph
 covering problems have also been considered, e.g. the “evolving dominating set” [6].

246 In Section 3.2 we use randomized rounding technique to prove that TVC on general
 247 temporal graphs admits a polynomial-time randomized approximation algorithm with ex-
 248 pected ratio $O(\ln n)$. This result is complemented by our results in Section 6.1 where we
 249 prove that SW-TVC (and thus also TVC) can be deterministically approximated with ratio
 250 $H_{2n\Delta} - \frac{1}{2} \approx \ln n + \ln 2\Delta - \frac{1}{2}$ in polynomial time.

251 3.1 Hardness on star temporal graphs

252 In the next theorem we reduce SET COVER to TVC on star temporal graphs, and vice versa.
 253 Our hardness results are complemented in Theorem 6 by reducing from HITTING SET.

254 ▶ **Theorem 5.** *TVC on star temporal graphs is NP-complete and it admits a polynomial-time*
 255 *$(H_{n-1} - \frac{1}{2})$ -approximation algorithm. Furthermore, for any $\varepsilon > 0$, TVC on star temporal*
 256 *graphs does not admit any polynomial-time $(1 - \varepsilon) \ln n$ -approximation algorithm, unless NP*
 257 *has $n^{O(\log \log n)}$ -time deterministic algorithms.*

258 ▶ **Theorem 6.** *For every $\varepsilon < 1$, TVC on star temporal graphs cannot be optimally solved*
 259 *in $O(2^{\varepsilon n})$ time, unless the Strong Exponential Time Hypothesis (SETH) fails.*

260 3.2 A randomized rounding algorithm for TVC

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

264 ▶ **Theorem 7.** *There exists a polynomial-time randomized approximation algorithm for TVC*
 265 *with expected approximation factor $O(\ln n)$.*

266 4 An almost tight algorithm for SW-TVC

267 In this section we investigate the complexity of SLIDING WINDOW TEMPORAL VERTEX
 268 COVER (SW-TVC). First we prove in Section 4.1 a strong lower bound on the complexity
 269 of optimally solving this problem on arbitrary temporal graphs. More specifically we
 270 prove that, for *any* (arbitrarily growing) functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$, there
 271 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,
 272 assuming the Exponential Time Hypothesis (ETH). This ETH-based lower bound turns
 273 out to be asymptotically almost tight. In fact, we present in Section 4.2 an exact dynamic
 274 programming algorithm for SW-TVC whose running time on an arbitrary temporal graph is
 275 $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$, which is asymptotically almost optimal, assuming ETH. In Section 4.3
 276 we prove that our algorithm can be refined so that, when the vertex cover number of each
 277 snapshot G_i is bounded by a constant k , the running time becomes $O(T\Delta(n+m) \cdot n^{k(\Delta+1)})$.
 278 That is, when Δ is a constant, this algorithm is polynomial in the input size on temporal
 279 graphs with bounded vertex cover number at every slot. Notably, for the class of always star
 280 temporal graphs (i.e. a superclass of the star temporal graphs studied in Section 3.1) the
 281 running time of the algorithm is $O(T\Delta(n+m) \cdot 2^\Delta)$.

282 4.1 A complexity lower bound

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

287 ► **Theorem 8.** For any two (arbitrarily growing) functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, there exists a
 288 constant $\varepsilon \in (0, 1)$ such that SW-TVC cannot be solved in $f(T) \cdot 2^{\varepsilon n \cdot g(\Delta)}$ time assuming
 289 ETH, where n is the number of vertices in the underlying graph G of the temporal graph.

290 4.2 An exact dynamic programming algorithm

291 The main idea of our dynamic programming algorithm for SW-TVC is to scan the temporal
 292 graph from left to right with respect to time (i.e. to scan the snapshots G_i increasingly on i),
 293 and at every time slot to consider all possibilities for the vertex appearances at the previous Δ
 294 time slots. Let (G, λ) be a temporal graph with n vertices and lifetime T , and let $\Delta \leq T$. For
 295 every $t = 1, 2, \dots, T - \Delta + 1$ and every Δ -tuple of vertex subsets A_1, \dots, A_Δ of G , we define
 296 $f(t; A_1, A_2, \dots, A_\Delta)$ to be the smallest cardinality of a sliding Δ -window temporal vertex cover
 297 \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 , $\mathcal{S}_{t+\Delta-1} = (A_\Delta, t+\Delta-1)$.
 298 If there exists no sliding Δ -window temporal vertex cover \mathcal{S} of $(G, \lambda)|_{[1, t+\Delta-1]}$ with these
 299 prescribed vertex appearances in the time slots $t, t+1, \dots, t+\Delta-1$, then we define
 300 $f(t; A_1, A_2, \dots, A_\Delta) = \infty$. Note that, once we have computed all possible values of the
 301 function $f(\cdot)$, then the optimum solution of SW-TVC on (G, λ) has cardinality

$$302 \quad \text{OPT}_{\text{SW-TVC}}(G, \lambda) = \min_{A_1, A_2, \dots, A_\Delta \subseteq V} \{f(T - \Delta + 1; A_1, A_2, \dots, A_\Delta)\}. \quad (1)$$

303 ► **Lemma 9.** Let (G, λ) be a temporal graph, where $G = (V, E)$. Let $2 \leq t \leq T - \Delta + 1$ and
 304 let $A_1, A_2, \dots, A_\Delta$ be a Δ -tuple of vertex subsets of the underlying graph G . Suppose that
 305 $\bigcup_{i=1}^{\Delta} (A_i, t+i-1)$ is a temporal vertex cover of $(G, \lambda)|_{[t, t+\Delta-1]}$. Then

$$306 \quad f(t; A_1, A_2, \dots, A_\Delta) = |A_\Delta| + \min_{X \subseteq V} \{f(t-1; X, A_1, \dots, A_{\Delta-1})\}. \quad (2)$$

307 Using the recursive computation of Lemma 9, we are now ready to present Algorithm 1
 308 for computing the value of an optimal solution of SW-TVC on a given arbitrary temporal
 309 graph (G, λ) . Note that Algorithm 1 can be easily modified such that it also computes the
 310 actual optimum solution of SW-TVC (instead of only its optimum cardinality). The proof
 311 of correctness and running time analysis of Algorithm 1 are given in the next theorem.

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  $\min_{A_1, \dots, A_\Delta \subseteq V} \{f(T - \Delta + 1; A_1, \dots, A_\Delta)\}$ 

```

312 ► **Theorem 10.** Let (G, λ) be a temporal graph, where $G = (V, E)$ has n vertices and m edges.
 313 Let T be its lifetime and let Δ be the length of the sliding window. Algorithm 1 computes
 314 in $O(T\Delta(n+m) \cdot 2^{n(\Delta+1)})$ time the value of an optimal solution of SW-TVC on (G, λ) .

4.3 Always bounded vertex cover number temporal graphs

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 10.

► **Theorem 11.** *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, our search at every step reduces to just one binary choice for each of the previous Δ time slots, 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 11.

► **Theorem 12.** *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) .

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 graphs.

To prove the main result (in Theorem 14) 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 13.** *VERTEX COVER is APX-hard on \mathcal{Y} .*

Let now \mathcal{X} be the class of graphs whose connected components are induced subgraphs of the graph obtained from the star with three leaves by subdividing each of its edges exactly once. 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} .

► **Theorem 14.** *2-TVC is APX-hard on always \mathcal{X} temporal graphs.*

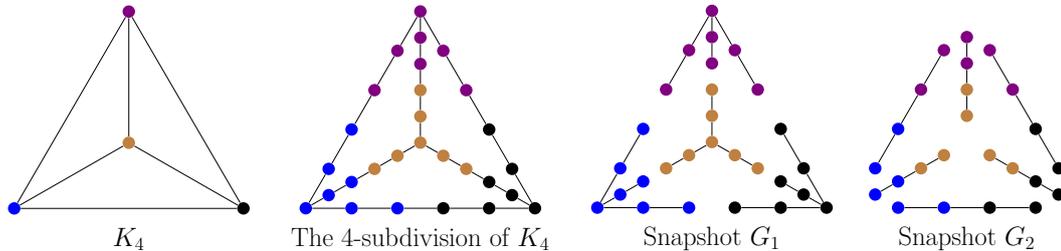
Proof. To prove the theorem we will reduce VERTEX COVER on \mathcal{Y} to 2-TVC on always \mathcal{X} temporal graphs. Let $H = (V, E)$ be a graph in \mathcal{Y} . First we will show how to construct an always \mathcal{X} temporal graph (G, λ) of lifetime 2. Then we will prove that the size τ of a minimum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

357 Let $R \subseteq V$ be the set of vertices of degree 3 in H . We define (G, λ) to be a temporal
 358 graph of lifetime 2, where snapshot G_1 is obtained from H by removing the edges with both
 359 ends being at distance exactly 2 from R , and snapshot $G_2 = H - R$. Figure 1 illustrates the
 360 reduction for $H = K_4$.

361 Let $\mathcal{S} = (S_1, 1) \cup (S_2, 2)$ be an arbitrary sliding 2-window temporal vertex cover of (G, λ)
 362 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 ,
 363 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
 364 chosen arbitrarily we further conclude that $\tau \leq \sigma$.

365 To show the converse inequality, let $C \subseteq V$ be a minimum vertex cover of H . Let S_1
 366 be those vertices in C which either have degree 3, or have a neighbor of degree 3. Let also
 367 $S_2 = C \setminus S_1$. We claim that $(S_1, 1) \cup (S_2, 2)$ is a sliding 2-window temporal vertex cover
 368 of (G, λ) . First, let $e \in E$ be an edge in H incident to a vertex of degree 3. Then, by the
 369 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
 370 belongs to S_1 . Hence, e is temporally covered by $(v, 1)$ in (G, λ) . Let now $e \in E$ be an edge
 371 in H whose both end vertices have degree 2. If one of the end vertices of e is adjacent to
 372 a vertex of degree 3 in H , then, by the construction, e is active in both time slots 1 and 2.
 373 Therefore, since $C = S_1 \cup S_2$, edge e will be temporally covered in (G, λ) in at least one of
 374 the time slots. Finally, if none of the end vertices of e is adjacent to a vertex of degree 3
 375 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
 376 vertex v in C covering e belongs to S_2 . Hence, e is temporally covered by $(v, 2)$ in (G, λ) .
 377 This shows that $(S_1, 1) \cup (S_2, 2)$ is a sliding 2-window temporal vertex cover of (G, λ) , and
 378 therefore $\sigma \leq |S_1| + |S_2| = |C| = \tau$.

379 Note that the size of a minimum vertex cover of H is equal to the size of a minimum-
 380 cardinality sliding 2-window temporal vertex cover of (G, λ) and that any feasible solution to
 381 2-TVC on (G, λ) of size r defines a vertex cover of H of size at most r . Thus, since VERTEX
 382 COVER is APX-hard on \mathcal{Y} by Lemma 13 and the reduction is approximation-preserving, it
 383 follows that 2-TVC is APX-hard as well. ◀



■ **Figure 1** A cubic graph K_4 , its 4-subdivision, and the corresponding snapshots G_1 and G_2

384 6 Approximation algorithms

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

389 6.1 Approximations in terms of T , Δ , and the largest edge frequency

390 We begin by presenting a reduction from SW-TVC to SET COVER, which proves useful
 391 for deriving approximation algorithms for the original problem. Consider an instance,

392 (G, λ) and $\Delta \leq T$, of the SW-TVC problem. Construct an instance of SET COVER as
 393 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
 394 pairs (e, t) of an edge e and a time slot t such that e appears (and so must be temporally
 395 covered) within window W_t . For every vertex appearance (v, s) we define $C_{v,s}$ to be the
 396 set of elements (e, t) in the universe U , such that (v, s) temporally covers e in the window
 397 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
 398 family of all sets $C_{v,s}$, where $v \in V, s \in [1, T]$. The following lemma shows that finding
 399 a minimum-cardinality sliding Δ -window temporal vertex cover of (G, λ) is equivalent to
 400 finding a minimum-cardinality family of sets $C_{v,s}$ that covers the universe U .

401 **► Lemma 15.** *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 $\mathcal{S} =$
 402 $\{(v_1, t_1), \dots, (v_k, t_k)\}$ is a sliding Δ -window temporal vertex cover of (G, λ) .*

403 **$O(\ln n + \ln \Delta)$ -approximation.** In the instance of SET COVER constructed by the above
 404 reduction, every set $C_{v,s}$ in \mathcal{C} contains at most $n\Delta$ elements of the universe U . Indeed, the
 405 vertex appearance (v, s) temporally covers at most $n-1$ edges, each in at most Δ windows
 406 (namely from window $W_{s-\Delta+1}$ up to window W_s). Thus we can apply the polynomial-
 407 time greedy algorithm from [9] for SET COVER which achieves an approximation ratio of
 408 $H_{n\Delta} - \frac{1}{2} = \sum_{i=1}^{n\Delta} \frac{1}{i} - \frac{1}{2} \approx \ln n + \ln \Delta - \frac{1}{2}$.

409 **$2k$ -approximation, where k is the maximum edge frequency.** Given a temporal graph
 410 (G, λ) and an edge e of G , the Δ -frequency of e is the maximum number of time
 411 slots at which e appears within a Δ -window. Let k denote the maximum Δ -frequency
 412 over all edges of G . Clearly, for a particular Δ -window W_t , an edge $e \in E[W_t]$ can be
 413 temporally covered in W_t by at most $2k$ vertex appearances. So in the above reduction
 414 to SET COVER, every element $(e, t) \in U$ belongs to at most $2k$ sets in \mathcal{C} . Therefore, the
 415 optimal solution of the constructed instance of SET COVER can be approximated within
 416 a factor of $2k$ in polynomial time [30], yielding a $2k$ -approximation for SW-TVC.

417 **2Δ -approximation.** Since the maximum Δ -frequency of an edge is always upper-bounded
 418 by Δ , the previous algorithm gives a worst-case polynomial-time 2Δ -approximation for
 419 SW-TVC on arbitrary temporal graphs.

420 6.2 Approximation in terms of maximum degree of snapshots

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

428 Single-edge temporal graphs

429 Consider a temporal graph (G_0, λ) where G_0 is the single-edge graph, i.e. $V(G_0) = \{u, v\}$
 430 and $E(G_0) = \{uv\}$. We reduce SW-TVC on (G_0, λ) to an instance of INTERVAL COVERING.

INTERVAL COVERING

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

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

467:12 Temporal Vertex Cover with a Sliding Time Window

432 An easy linear-time greedy algorithm for the INTERVAL COVERING picks at each iteration,
 433 among the intervals that cover the leftmost uncovered point, the one with largest finishing
 434 time. Algorithm 2 implements this simple rule in the context of the SW-TVC problem.

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  $\mathcal{S}$ 

```

435 ► **Lemma 16.** *Algorithm 2 solves SW-TVC on a single-edge temporal graph and can be*
 436 *implemented to work in time $O(T)$.*

437 Always degree at most d temporal graphs

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

442 Let (G, λ) be a temporal graph, where $G = (V, E)$, $|V| = n$, and $|E| = m$. For every edge
 443 $e = uv \in E$, let $(G[\{u, v\}], \lambda)$ denote the temporal graph where the underlying graph is the
 444 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  $\mathcal{S}$ 

```

445 ► **Lemma 17.** *Algorithm 3 is a $O(mT)$ -time d -approximation algorithm for SW-TVC on*
 446 *always degree at most d temporal graphs.*

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

449 ► **Corollary 18.** *SW-TVC can be optimally solved in $O(mT)$ time on the class of always*
 450 *matching temporal graphs.*

References

- 451 ———
- 452 1 Eric Aaron, Danny Krizanc, and Elliot Meyerson. DMVP: foremost waypoint coverage of
453 time-varying graphs. In *Proceedings of the 40th International Workshop on Graph-Theoretic
454 Concepts in Computer Science (WG)*, pages 29–41, 2014.
- 455 2 Eleni C. Akrida, Leszek Gasieniec, George B. Mertzios, and Paul G. Spirakis. Ephemeral
456 networks with random availability of links: The case of fast networks. *Journal of Parallel
457 and Distributed Computing*, 87:109–120, 2016.
- 458 3 Eleni C. Akrida, Leszek Gasieniec, George B. Mertzios, and Paul G. Spirakis. The com-
459 plexity of optimal design of temporally connected graphs. *Theory of Computing Systems*,
460 61(3):907–944, 2017.
- 461 4 Binh-Minh Bui-Xuan, Afonso Ferreira, and Aubin Jarry. Computing shortest, fastest, and
462 foremost journeys in dynamic networks. *International Journal of Foundations of Computer
463 Science*, 14(2):267–285, 2003.
- 464 5 Arnaud Casteigts and Paola Flocchini. Deterministic Algorithms in Dynamic Networks:
465 Formal Models and Metrics. Technical report, Defence R&D Canada, April 2013. URL:
466 <https://hal.archives-ouvertes.fr/hal-00865762>.
- 467 6 Arnaud Casteigts and Paola Flocchini. Deterministic Algorithms in Dynamic Networks:
468 Problems, Analysis, and Algorithmic Tools. Technical report, Defence R&D Canada, April
469 2013. URL: <https://hal.archives-ouvertes.fr/hal-00865764>.
- 470 7 Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-
471 varying graphs and dynamic networks. *International Journal of Parallel, Emergent and
472 Distributed Systems (IJPEDES)*, 27(5):387–408, 2012.
- 473 8 Arnaud Casteigts, Bernard Mans, and Luke Mathieson. On the feasibility of maintenance
474 algorithms in dynamic graphs. *CoRR*, abs/1107.2722, 2011. URL: [http://arxiv.org/
475 abs/1107.2722](http://arxiv.org/abs/1107.2722).
- 476 9 Rong chii Duh and Martin Fürer. Approximation of k -set cover by semi-local optimization.
477 In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*,
478 pages 256–264, 1997.
- 479 10 Andrea E. F. Clementi, Claudio Macci, Angelo Monti, Francesco Pasquale, and Riccardo
480 Silvestri. Flooding time of edge-markovian evolving graphs. *SIAM Journal on Discrete
481 Mathematics (SIDMA)*, 24(4):1694–1712, 2010.
- 482 11 Stefan Dobrev, Stephane Durocher, Mohsen Eftekhari Hesari, Konstantinos Georgiou,
483 Evangelos Kranakis, Danny Krizanc, Lata Narayanan, Jaroslav Opatrny, Sunil M. Shende,
484 and Jorge Urrutia. Complexity of barrier coverage with relocatable sensors in the plane.
485 *Theoretical Computer Science*, 579:64–73, 2015.
- 486 12 Thomas Erlebach, Michael Hoffmann, and Frank Kammer. On temporal graph explo-
487 ration. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and
488 Programming (ICALP)*, pages 444–455, 2015.
- 489 13 Afonso Ferreira. Building a reference combinatorial model for MANETs. *IEEE Network*,
490 18(5):24–29, 2004.
- 491 14 Paola Flocchini, Bernard Mans, and Nicola Santoro. Exploration of periodically varying
492 graphs. In *Proceedings of the 20th International Symposium on Algorithms and Computa-
493 tion (ISAAC)*, pages 534–543, 2009.
- 494 15 George Giakkoupis, Thomas Sauerwald, and Alexandre Stauffer. Randomized rumor
495 spreading in dynamic graphs. In *Proceedings of the 41st International Colloquium on
496 Automata, Languages and Programming (ICALP)*, pages 495–507, 2014.
- 497 16 Mohsen Eftekhari Hesari, Evangelos Kranakis, Danny Krizanc, Oscar Morales Ponce, Lata
498 Narayanan, Jaroslav Opatrny, and Sunil M. Shende. Distributed algorithms for barrier
499 coverage using relocatable sensors. *Distributed Computing*, 29(5):361–376, 2016.

- 500 **17** Anne-Sophie Himmel, Hendrik Molter, Rolf Niedermeier, and Manuel Sorge. Adapting
501 the bron-kerbosch algorithm for enumerating maximal cliques in temporal graphs. *Social*
502 *Network Analysis and Mining*, 7(1):35:1–35:16, 2017.
- 503 **18** P. Holme and J. Saramäki, editors. *Temporal Networks*. Springer, 2013.
- 504 **19** Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly
505 exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
- 506 **20** David Kempe, Jon M. Kleinberg, and Amit Kumar. Connectivity and inference problems
507 for temporal networks. In *Proceedings of the 32nd annual ACM symposium on Theory of*
508 *computing (STOC)*, pages 504–513, 2000.
- 509 **21** Orestis Kostakis and Aristides Gionis. On mining temporal patterns in dynamic graphs, and
510 other unrelated problems. In *Proceedings of the 6th International Conference on Complex*
511 *Networks and Their Applications (COMPLEX NETWORKS)*, pages 516–527, 2017.
- 512 **22** Orestis Kostakis, Nikolaj Tatti, and Aristides Gionis. Discovering recurring activity in
513 temporal networks. *Data Mining and Knowledge Discovery*, 31(6):1840–1871, 2017.
- 514 **23** Evangelos Kranakis, Danny Krizanc, Flaminia L. Luccio, and Brett Smith. Maintaining
515 intruder detection capability in a rectangular domain with sensors. In *Proceedings of the*
516 *11th International Symposium on Algorithms and Experiments for Wireless Sensor Net-*
517 *works (ALGOSENSORS)*, pages 27–40, 2015.
- 518 **24** Jure Leskovec, Jon M. Kleinberg, and Christos Faloutsos. Graph evolution: Densification
519 and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data*, 1(1),
520 2007.
- 521 **25** George B. Mertzios, Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Tem-
522 poral network optimization subject to connectivity constraints. In *Proceedings of the 40th*
523 *International Colloquium on Automata, Languages and Programming (ICALP), Part II*,
524 pages 657–668, 2013.
- 525 **26** Othon Michail and Paul G. Spirakis. Traveling salesman problems in temporal graphs.
526 *Theoretical Computer Science*, 634:1–23, 2016.
- 527 **27** Othon Michail and Paul G. Spirakis. Elements of the theory of dynamic networks. *Com-*
528 *munications of the ACM*, 61(2):72–72, January 2018.
- 529 **28** Sotiris E. Nikolettseas and Paul G. Spirakis. Probabilistic distributed algorithms for energy
530 efficient routing and tracking in wireless sensor networks. *Algorithms*, 2(1):121–157, 2009.
531 URL: <https://doi.org/10.3390/a2010121>.
- 532 **29** John Kit Tang, Mirco Musolesi, Cecilia Mascolo, and Vito Latora. Characterising temporal
533 distance and reachability in mobile and online social networks. *ACM Computer Commu-*
534 *nication Review*, 40(1):118–124, 2010.
- 535 **30** Vijay V. Vazirani. *Approximation algorithms*. Springer, 2003.
- 536 **31** Jordan Viard, Matthieu Latapy, and Clémence Magnien. Revealing contact patterns among
537 high-school students using maximal cliques in link streams. In *Proceedings of the 2015*
538 *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*
539 *(ASONAM)*, pages 1517–1522, 2015.
- 540 **32** Tiphaine Viard, Matthieu Latapy, and Clémence Magnien. Computing maximal cliques in
541 link streams. *Theoretical Computer Science*, 609:245–252, 2016.
- 542 **33** Chuan Zhu, Chunlin Zheng, Lei Shu, and Guangjie Han. A survey on coverage and connec-
543 tivity issues in wireless sensor networks. *Journal of Network and Computer Applications*,
544 35(2):619–632, 2012.