

How to reduce temporal cliques to find sparse spanners

Sebastian Angrick, Ben Bals, Tobias Friedrich, Hans Gawendowicz, Niko Hastrich, Nicolas Klodt, Pascal Lenzner, Jonas Schmidt, George Skretas, Armin Wells

Angaben zur Veröffentlichung / Publication details:

Angrick, Sebastian, Ben Bals, Tobias Friedrich, Hans Gawendowicz, Niko Hastrich, Nicolas Klodt, Pascal Lenzner, Jonas Schmidt, George Skretas, and Armin Wells. 2024. "How to reduce temporal cliques to find sparse spanners." In 32nd Annual European Symposium on Algorithms (ESA 2024): ESA 2024, September 2-4, 2024, Royal Holloway, London, United Kingdom, edited by Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, 11:1-15. Schloss Dagstuhl: Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
<https://doi.org/10.4230/LIPIcs.ESA.2024.11>.

Nutzungsbedingungen / Terms of use:

CC BY 4.0

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

CC-BY 4.0: Creative Commons: Namensnennung

Weitere Informationen finden Sie unter: / For more information see:

<https://creativecommons.org/licenses/by/4.0/deed.de>





How to Reduce Temporal Cliques to Find Sparse Spanners

Sebastian Angrick  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Tobias Friedrich  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Niko Hastrich  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Pascal Lenzner  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

George Skretas  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Ben Bals  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Hans Gawendowicz  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Nicolas Klodt  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Jonas Schmidt  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Armin Wells  

Faculty of Digital Engineering, Hasso Plattner Institute, University of Potsdam, Germany

Abstract

Many real-world networks, such as transportation or trade networks, are dynamic in the sense that the edge-set may change over time, but these changes are known in advance. This behavior is captured by the temporal graphs model, which has recently become a trending topic in theoretical computer science. A core open problem in the field is to prove the existence of linear-size temporal spanners in temporal cliques, i.e., sparse subgraphs of complete temporal graphs that ensure all-pairs reachability via temporal paths. So far, the best known result is the existence of temporal spanners with $\mathcal{O}(n \log n)$ many edges. We present significant progress towards proving whether linear-size temporal spanners exist in all temporal cliques.

We adapt techniques used in previous works and heavily expand and generalize them. This allows us to show that the existence of a linear spanner in cliques and bi-cliques is equivalent and using this, we provide a simpler and more intuitive proof of the $\mathcal{O}(n \log n)$ bound by giving an efficient algorithm for finding linearithmic spanners. Moreover, we use our novel and efficiently computable approach to show that a large class of temporal cliques, called edge-pivotable graphs, admit linear-size temporal spanners. To contrast this, we investigate other classes of temporal cliques that do not belong to the class of edge-pivotable graphs. We introduce two such graph classes and we develop novel algorithmic techniques for establishing the existence of linear temporal spanners in these graph classes as well.

2012 ACM Subject Classification Mathematics of computing → Paths and connectivity problems

Keywords and phrases Temporal Graphs, temporal Clique, temporal Spanner, Reachability, Graph Connectivity, Graph Sparsification

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.11

Related Version *Full Version:* <https://arxiv.org/abs/2402.13624>

1 Introduction

Many real-world networks, like transportation networks with scheduled train connections or social networks with repeating meeting schedules, are dynamic, i.e., their edge-set can change over time. To address this, *temporal graphs* have gained significant attention in recent



© Sebastian Angrick, Ben Bals, Tobias Friedrich, Hans Gawendowicz, Niko Hastrich, Nicolas Klodt, Pascal Lenzner, Jonas Schmidt, George Skretas, and Armin Wells; licensed under Creative Commons License CC-BY 4.0

32nd Annual European Symposium on Algorithms (ESA 2024).

Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 11; pp. 11:1–11:15



Leibniz International Proceedings in Informatics

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

research in theoretical computer science [15, 18]. In these graphs it is assumed that the edge-set can change in every time step but the schedule of the availability of edges is known in advance. Various problems have been studied in this model, including network redesign [10, 12, 13], vertex cover [1, 11, 14], and influence maximization [9]. A general theme is that insights from the static graph setting are only of limited value in temporal graphs, since the temporal availability of the edges makes these problems much harder or even infeasible to solve. For example, even fundamental statements like Menger’s Theorem do not hold for temporal graphs [17]. Also, one of the most essential problems in graphs, that of finding a sparse spanner, is much harder in temporal graphs. A sparse spanner is a small subset of the edge-set that ensures all-pairs connectivity. In static graphs, it is well-known that linear-size spanners exist and can be computed efficiently, e.g., minimum spanning trees [8].

Finding sparse *temporal spanners* is much more intricate in temporal graphs, since there paths do not necessarily compose (i.e., concatenating two paths where the endpoint of the first is the starting point of the second does not necessarily yield a temporal path). More specifically, Kempe, Kleinberg, and Kumar [15] showed that there is a class of temporal graphs having $\Omega(n \log n)$ many edges, such that no single edge can be removed while preserving temporal connectivity. In the same work, the authors asked, what is the minimum number of edges $c(n)$, such that any temporal graph with n vertices has a temporal spanner using at most $c(n)$ many edges. Almost 15 years later, this question was answered by Axiotis and Fotakis [2] by providing a class of dense but non-complete temporal graphs having $\Theta(n^2)$ edges, such that any subgraph preserving temporal connectivity necessarily needs to contain $\Omega(n^2)$ edges as well.

This negative result raises the question if there even exists a natural class of temporal graphs that can be sparsified without losing temporal connectivity. In particular, the question if every temporal clique has a spanner of linear size has remained open for over 20 years now. In a 2019 breakthrough paper, Casteigts, Peters, and Schoeters [6] proved that temporal cliques always admit a spanner of size $O(n \log n)$. Not much progress has been made since then, with most work on the topic veering off towards the study of temporal spanners in random temporal graphs [7], or focusing on low stretch [3, 4] or a game-theoretic setting [5].

In this paper, we present significant progress on the long-standing open problem by showing that a large class of temporal cliques does admit linear-size temporal spanners.

1.1 Our Contribution

In the first part of the paper, we adapt techniques used in previous works and heavily expand and generalize them. In particular, in [6] a temporal spanner in cliques is computed by a reduction from a clique to a bi-clique. We deepen this connection by showing that the existence of a linear spanner in cliques and bi-cliques actually is equivalent. Furthermore, in the reduction of Casteigts, Peters, and Schoeters [6], the authors exploit the property that the set containing the earliest edge of each vertex is a perfect matching of the bi-clique and that this also holds for the set of latest edges. We show that this property can also be assumed when searching for bi-spanners in bi-cliques directly. This yields a simple and more intuitive algorithm for computing linearithmic spanners, thus proving the existence of size $O(n \log n)$ temporal spanners in temporal cliques.

Moreover, we reconsider the notion of a *pivot-vertex* [6, 7]. This is a vertex u that can be reached by every vertex in the graph until some time step t , and then u can reach every other vertex in the graph by a temporal path starting at or after time t . We transfer this notion to bi-cliques, extend it to *c-pivot-edges*, and we strengthen it significantly to find small spanners. If a bi-clique contains a *c-pivot-edge*, we can iteratively reduce the size of

the graph, while maintaining a linear-size temporal spanner, until the instance is solved, or the instance no longer has a c -pivot-edge. This reduction rule is widely applicable and yields linear-size temporal spanners in many cases.

In the second part of the paper, we investigate temporal cliques that are not edge-pivotable graphs. We identify two such graph classes, called *shifted matching graphs* and *product graphs*, and for both we provide novel techniques that allow us to efficiently compute linear-size spanners. Thus, we extend the class of temporal cliques that admit linear-size temporal spanners even further. It is open if our classes cover every temporal clique. However, finding other graphs that do not belong to our graph classes seems to be challenging.

Statements where proofs or details are omitted due to space constraints are marked with “ \star ” in the full version.

1.2 Detailed Discussion of Related Work

The study of temporal spanners was initiated by Kempe, Kleinberg and Kumar [15]. In their seminal paper, they introduce the temporal graphs model and show that spanners on temporal hypercubes have $\Omega(n \log n)$ edges. Axiotis and Fotakis [2] improve on this by providing a class of graphs with $\Theta(n^2)$ edges, such that no edge can be removed while maintaining temporal connectivity. Later, Casteigts, Peters, and Schoeters [6] study temporal cliques and show that temporal spanners of size $O(n \log n)$ always exist. Their “fireworks” algorithm combines different techniques, such as dismountability and delegation, to reduce the clique to a bipartite graph, and then compute a spanner on that graph, that can then be translated to a spanner of the original graph.


Casteigts, Raskin, Renken and Zamaraev [7] take a different approach. They consider random temporal graphs, where the host graph is a clique, and each edge is added to the temporal graph with probability p . They show a sharp threshold for the generated temporal graph being temporally connected. Even more surprisingly, they show that as soon as this threshold is reached, the generated graph will contain an almost optimal spanner of size $2n + o(n)$. On a different note, Bilò, D’Angelo, Gualà, Leucci, and Rossi [4] investigate whether temporal cliques admit small temporal spanners with low stretch. Their main result shows that they can always compute a temporal spanner with stretch $O(\log n)$ and $O(n \log^2 n)$ edges.

2 Preliminaries


For two sets A, B , we write $A \otimes B := \{\{a, b\} \mid a \in A, b \in B, \text{ and } a \neq b\}$. If A and B are disjoint, we write $A \sqcup B$ for their union. For $n \in \mathbb{N}^+$, let $[n] = \{1, \dots, n\}$.

Let (V, E) be an undirected graph. This together with a labeling function¹ $\lambda: E \rightarrow \mathbb{N}$ forms the *temporal graph* $G = (V, E, \lambda)$. Semantically, an edge $e \in E$ is present at time $\lambda(e)$. For a vertex-set $U \subseteq V$, write $G[U]$ for the subgraph of G induced by U . For $v \in V$, let $N(v)$ denote the set of neighbors of v , that is, all vertices u with $\{v, u\} \in E$. For $v \in V$ and $u, u' \in N(v)$, we define the order $u \preceq_v u'$, if the edge $\{u, v\}$ is not later than the edge $\{u', v\}$, i.e., if $\lambda(\{u, v\}) \leq \lambda(\{u', v\})$. We define \prec_v, \succeq_v , and \succ_v analogously.

¹ In general each edge could have more than one label. We restrict ourselves to the case of *simple* temporal graphs where each edge has exactly one label. Both formulations are equivalent for our use-case [6].



(a) Blue edges and paths have labels earlier than $\lambda(e)$, red ones are later.



(b) A bi-clique where a_1 cannot reach a_2 using a path starting after time 2 and a_1 can not be reached by a_2 using a path arriving before 5, so $\text{In}_{a_1}^3 \cup \text{Out}_{a_1}^3 \neq V$.

■ **Figure 1** Examples for the sets In and Out.

If $P = v_1 \dots v_k$ is a path in (V, E) and for all $i \in [k - 2]$, we have $v_i \preceq_{v_{i+1}} v_{i+2}$, then P is a *temporal path* in G .² We will also write $P = v_1 \rightarrow v_2 \xrightarrow{S} v_k$, where, $v_1 \rightarrow v_2$ describes a direct edge and $v_2 \xrightarrow{S} v_k$ is a temporal path starting after $\lambda(\{v_1, v_2\})$ using only edges of $S \subseteq E$. Since only reachability is our focus, we will also construct non-simple paths. Note that such paths can be shortcut to vertex-disjoint paths while remaining temporal.

Vertex u can *reach* vertex v if there is a temporal path from u to v . Also, G is called *temporally connected*, if every vertex can reach every other vertex. A *spanner* for G is a set $S \subseteq E$, such that $(V, S, \lambda|_S)$ is temporally connected, where $\lambda|_S$ is λ restricted to S .

The edges with the smallest and largest label take a special role. For $v \in V$, let $\pi^-(v) := \arg \min_{u \in N(v)} \lambda(\{v, u\})$ be v 's earliest neighbor and let $\pi^+(v) := \arg \max_{u \in N(v)} \lambda(\{v, u\})$ be v 's latest neighbor. We do not consider graphs with isolated vertices and we justify in Section 3 that we only use injective labelings, so this is well-defined. Denote the set of all earliest edges as $\pi^- := \{\{v, \pi^-(v)\} \mid v \in V\}$ and the set of latest edges as $\pi^+ := \{\{v, \pi^+(v)\} \mid v \in V\}$. If both sets π^- and π^+ form a perfect matching for V , then G is *extremally matched*. Also, for $v, w \in V$, and $e = \{v, w\} \in E$, we write $\sigma_v(e)$ or simply $\sigma_v(w)$ for the index of e in $[n]$, when ordering all edges incident to v ascendingly by their label (e.g., if $w = \pi^-(v)$, then $\sigma_v(e) = 1$).

Let $v \in V$ be a vertex and $t \in \mathbb{N}$ be a time step. We define the set $\text{In}_v^t \subseteq V$ to be the vertices reaching v via a temporal path with every label being at most t . Respectively, we define the set $\text{Out}_v^t \subseteq V$ to be the vertices that v can reach via any temporal path starting at or after t . As a shorthand, for an edge $e = \{u, v\} \in E$, we define $\text{In}(e) = \text{In}_u^{\lambda(e)} \cup \text{In}_v^{\lambda(e)}$ and $\text{Out}(e) = \text{Out}_u^{\lambda(e)} \cup \text{Out}_v^{\lambda(e)}$. Note, that we always have $u, v \in \text{In}(e) \cap \text{Out}(e)$.

We focus only on temporal graphs where the underlying graph is a clique or a complete bipartite graph. In these cases, the set of edges is self-evident and for cliques we simply omit it and write $C = (V, \lambda)$. We, similarly, denote complete bipartite temporal graphs by a triple $D = (A, B, \lambda)$, where A and B are the two parts and $\lambda: A \otimes B \rightarrow \mathbb{N}$ is the corresponding labeling function. When necessary, we refer to the set of edges as $E(C) := V \otimes V$ or $E(D) := A \otimes B$. We call C a *temporal clique* and D a *temporal bi-clique*. Finally, we define an asymmetric analogue of spanners for temporal bi-cliques. We call a set $S \subseteq A \otimes B$ a *bi-spanner* for D , if every $a \in A$ can reach every $b \in B$ through a temporal path that is contained in S . Note that here, only one side needs to reach the other.

For $e = \{u, v\}$ in temporal cliques or bi-cliques, notice that all vertices must be in $\text{In}(e)$ or $\text{Out}(e)$, since all vertices have an edge to u or v , which puts this vertex in $\text{In}(e)$ or $\text{Out}(e)$, depending on the time-label. See Figure 1a for a visualization of In and Out.

► **Observation 2.1.** *Let $G = (V, E, \lambda)$ be a temporal clique or bi-clique. For all $e \in E$, we have $\text{In}(e) \cup \text{Out}(e) = V$.*

² This is also referred to as non-strict temporal paths, where strict temporal paths ask that $v_i \prec_{v_{i+1}} v_{i+1}$.

Note that for the case of bi-cliques considering both endpoints of e is crucial to ensure that every vertex has a direct edge to an endpoint. Otherwise, there exist simple instances with a vertex v and a time t with $\text{In}_v^t \cup \text{Out}_v^t \neq A \sqcup B$. See Figure 1b.

► **Remark 2.2.** For $v \in A \sqcup B$ and $t \in \mathbb{N}$, the analogue to Observation 2.1 (i.e., $\text{In}_v^t \cup \text{Out}_v^t = A \sqcup B$) does not necessarily hold, as one can see in Figure 1b.

Finally, for $n \in \mathbb{N}$, let $\mathcal{C}(n)$ be the minimum number such that every temporal clique with n vertices admits a spanner with at most $\mathcal{C}(n)$ edges. Analogously, let $\mathcal{D}(n)$ be the minimum number such that any bi-clique in which both sides have n vertices admits a bi-spanner with at most $\mathcal{D}(n)$ edges. We study the open problem whether $\mathcal{C}(n) \in \mathcal{O}(n)$ and whether $\mathcal{D}(n) \in \mathcal{O}(n)$ holds.

3 Cliques and Bi-Cliques

In this section, we prove that the asymptotic behavior of spanners in cliques differs from the asymptotic behavior of bi-spanners in bi-cliques only by a constant factor (i.e., we prove $\mathcal{C}(n) \in \Theta(\mathcal{D}(n))$). First, we show that $\mathcal{C}(n) \in \mathcal{O}(\mathcal{D}(n))$. After this, we show $\mathcal{C}(n) \in \Omega(\mathcal{D}(n))$ by proving each inequality of $\mathcal{C}(n) \geq \mathcal{D}(\frac{n}{2}) \geq \frac{1}{4} \mathcal{D}(n)$ separately. Bi-cliques and bi-spanners were already discussed in [6], when deriving $\mathcal{C}(n) \in \mathcal{O}(n \log n)$. They called these instances *residual* as they remained once their other techniques were not applicable anymore. Our result can be seen as evidence that the connection between cliques and bi-cliques is fundamental, beyond being a step in their algorithm. First, we state the main result of this section.

► **Theorem 3.1.** *The worst-case minimum size of a spanner for temporal cliques and temporal bi-cliques only differs by a constant factor, that is, we have $\mathcal{C}(n) \in \Theta(\mathcal{D}(n))$.*

We start with a technical lemma that allows us to focus on globally injective labeling functions for the rest of the paper. This slightly extends a result by [6] which proved the generality of locally injective labelings. Note that the statement holds for all temporal graphs.

► **Lemma 3.2** (★). *Let $G = (V, E, \lambda)$ be a temporal graph. Then there is an injective labeling function λ' s.t. every spanner $S \subseteq E$ for (V, E, λ') is also a spanner for the original graph G . In case the original graph $D = (A, B, \lambda)$ is a bi-clique, the same holds for a bi-spanner S .*

We can now relate minimum spanners in cliques to minimum bi-spanners in bi-cliques.

► **Lemma 3.3** (★). *Let $C = (V, \lambda)$ be a temporal clique. There is a temporal bi-clique $D = (A, B, \lambda')$ with $|V| = |A| = |B|$ s.t. for any bi-spanner S_D in D there is a spanner S_C in C with $|S_C| \leq |S_D|$.*

This is enough to arrive at the first half of the tight bound we aim to prove in the section.

► **Theorem 3.4** (★). *For $n \in \mathbb{N}$, we have $\mathcal{C}(n) \leq \mathcal{D}(n)$.*

We now show that $\mathcal{C}(2n) \geq \mathcal{D}(n)$. To this end, we take a bi-clique and add edges within the parts that yield no benefit in the construction of a spanner.

► **Lemma 3.5.** *Let $D = (A, B, \lambda)$ be a temporal bi-clique with $|A| = |B|$. Then there is a temporal clique $C = (A \sqcup B, \lambda')$ such that for any spanner S in C the set $S \cap (A \otimes B)$ is a bi-spanner in D .*

11:6 How to Reduce Temporal Cliques to Find Sparse Spanners

Proof. We define λ' such that for distinct $v, w \in A \sqcup B$, we have

$$\lambda'(\{v, w\}) = \begin{cases} 0, & \text{if } v, w \in B; \\ 1 + \lambda(\{v, w\}), & \text{if } v \in A, w \in B \text{ or } v \in B, w \in A; \\ \mu & \text{if } v, w \in A, \end{cases}$$

where $\mu > 1 + \max_{e \in A \otimes B} \lambda(e)$. Intuitively, all early edges are between vertices in B and all late edges are between vertices in A . The original edges keep their ordering.

Let S be a spanner of C . To show $S \cap (A \otimes B)$ is a bi-spanner in D , let $a \in A$ and $b \in B$. As S is a spanner, there is a temporal ab -path $P \subseteq S$. If P contains any edges within A , no edge to B can be used afterwards. Similarly, once the path arrives at B , it is already too late to use any edge within B . Thus, $P \subseteq A \otimes B$ and $S \cap (A \otimes B)$ is a bi-spanner in D . ◀

This yields the following relationship between clique and bi-clique spanner sizes.

► **Theorem 3.6** (★). *For $n \in \mathbb{N}$, we have $\mathcal{C}(2n) \geq \mathcal{D}(n)$.*

To conclude this section, we just need to show, that $\mathcal{D}(n) \geq \frac{1}{4} \mathcal{D}(2n)$. To that end, we show the slightly more general statement that for any $k \in \mathbb{N}^+$, we have $k^2 \mathcal{D}(n) \geq \mathcal{D}(kn)$, which we achieve by splitting the instances into smaller subinstances.

► **Lemma 3.7** (★). *For all $k, n \in \mathbb{N}^+$, we have $k^2 \mathcal{D}(n) \geq \mathcal{D}(kn)$.*

► **Remark 3.8.** We later strengthen this statement using tools developed in Section 4. Concretely, considering $D_i := D[A_i \cup B]$, we show $2k(k-1)n + k \mathcal{D}(n) \geq \mathcal{D}(kn)$.

To plug all pieces together, we need a final property about \mathcal{D} .

► **Lemma 3.9** (★). *The worst-case size of spanners for bi-cliques is monotonic in the size of the bi-clique, that is for all $n \in \mathbb{N}$, we have $\mathcal{D}(n) \leq \mathcal{D}(n+1)$. The same holds for \mathcal{C} .*

Now all pieces are in place to puzzle together the proof of our main theorem.

Proof of Theorem 3.1. Let $n \in \mathbb{N}$ and $c \in \{0, 1, 2, 3\}$ be such that $n \geq 3$ and $n+c$ is a multiple of 4. In the following equation, we use Lemma 3.9 in the first and last step, apply Lemma 3.7 in the third step and Theorem 3.6 in the fourth step, to obtain

$$\mathcal{D}(n) \leq \mathcal{D}(n+c) = \mathcal{D}\left(4 \cdot \frac{n+c}{4}\right) \leq 16 \cdot \mathcal{D}\left(\frac{n+c}{4}\right) \leq 16 \cdot \mathcal{C}\left(\frac{n+c}{2}\right) \leq 16 \cdot \mathcal{C}(n).$$

Theorem 3.4 directly gives us $\mathcal{C}(n) \leq \mathcal{D}(n)$, allowing us to conclude $\mathcal{C}(n) \in \Theta(\mathcal{D}(n))$. ◀

4 Structure of Bi-Clique Instances

Using Theorem 3.1, from now on we only consider bi-cliques. Our first structural insight for this type of instance was already used by Casteigts, Peters, and Schoeters [6] on special reduced instances of bi-cliques. Remember, that we call an instance extremally matched if the set of all earliest edges π^- and the set of all latest edges π^+ each form a perfect matching. We give a general reduction rule which reduces the graph in case the instance is not extremally matched. This structural property, that we can assume the existence of such perfect matchings or reduce the instance, is useful for almost all further constructions. For example, it enables us to reduce the instance size, when the number of vertices per side is not equal. Additionally, we can think of the sets A and B of the bi-clique as interchangeable.

For example, consider any set of edges that includes temporal paths from some set $A' \subseteq A$ to all vertices in A . We can extend these paths to reach all vertices in B by including the set of all latest edges, because it is a perfect matching between A and B .

We start with proving that every $b \in B$ is incident to at most one edge in π^- or we can reduce. The same holds for any $a \in A$ with respect to π^+ . For our reduction rules, we adapt the concept of dismantability on temporal cliques, first defined in [6], to temporal bi-cliques.

► **Definition 4.1** (dismountable). *Vertex $a \in A$ is dismantable if there is $a' \in A$ with $a \neq a'$, s.t. $a \preceq_{\pi^-(a')} a'$. Vertex $b \in B$ is dismantable if there is $b' \in B$ with $b \neq b'$, s.t. $b' \preceq_{\pi^+(b')} b$.*

Intuitively, vertex a can delegate its obligation to reach all vertices in B to vertex a' in exchange for including two edges in the bi-spanner. Similarly, vertex b can be reached by all vertices in A that can reach vertex b' if we include the two necessary edges in the bi-spanner.

Note that this corresponds to the definition in [6], but transferred to the case of bi-spanners, where all vertices in A only need to reach vertices in B . Concretely, this means that vertex a can reach vertex a' via the earliest edge of a' on a temporal path of length 2. Therefore, if this configuration occurs in the bi-clique, we can remove a and include the edges $\{a, \pi^-(a')\}$ and $\{a', \pi^-(a')\}$ into our bi-spanner. In the final solution for the whole bi-clique, vertex a then has a temporal path to every vertex $b \in B$ via vertex a' . Similarly, if a vertex $b \in B$ is dismantable, every temporal path that ends in b' can be extended to be a temporal path ending in vertex b , and we can remove b from the bi-clique.

The following lemma is similar to the one in [6]. For completeness, we prove it again.

► **Lemma 4.2** (\star). *Let $D = (A, B, \lambda)$ be a bi-clique. Let $a \in A$, $\pi^-(a) = b$, then $a = \pi^-(b)$ or $\pi^-(b)$ is dismantable. Let $b \in B$, $\pi^+(b) = a$, then $\pi^+(a) = b$ or $\pi^+(a)$ is dismantable.*

We can apply Lemma 4.2 to show that we can always find a dismantable vertex whenever the number of vertices per side differs.

► **Lemma 4.3** (\star). *For $D = (A, B, \lambda)$, if $|A| \neq |B|$ then a dismantable vertex exists.*

Since from now on, both parts of D have the same size after removing dismantable vertices, we write $n := |A| = |B|$ and call n the size of the instance. This justifies also that the definition of \mathcal{D} assumes equal part sizes and gives the following corollary.

► **Corollary 4.4**. *Let $D = (A, B, \lambda)$ and $|A| \leq |B|$. Then, there is a bi-spanner for D of size $\mathcal{D}(|A|) + 2(|B| - |A|)$. Similarly, if $|B| \leq |A|$ there is a bi-spanner of size $\mathcal{D}(|B|) + 2(|A| - |B|)$.*

Finally, combining our lemmas yields that π^- and π^+ are extremally matched.

► **Theorem 4.5** (Extremal Matching), (\star). *Let $D = (A, B, \lambda)$ be a bi-clique. Then D is extremally matched or there is a dismantable vertex.*

After removing dismantable vertices, we now know that the remaining graph must be extremally matched. Since reducing an instance to be extremally matched only requires the inclusion of linearly many edges, we focus from now on extremally matched bi-cliques. This also allows us to switch the roles of A and B , which will be useful later on.

► **Lemma 4.6** (\star). *Let $D = (A, B, \lambda)$ be an extremally matched bi-clique and let S be a bi-spanner for D . If $\pi^-, \pi^+ \subseteq S$, then S is also a bi-spanner for (B, A, λ) .*

With Corollary 4.4 we prove Lemma 4.7. This gives a simpler proof that $\mathcal{C}(n) \in \mathcal{O}(n \log n)$.

► **Lemma 4.7** (\star). *For all $k, n \in \mathbb{N}^+$, we have $2k(k-1)n + k\mathcal{D}(n) \geq \mathcal{D}(kn)$.*

If we now choose $k > 1$, we get $\mathcal{C}(n) \in \mathcal{O}(n \log n)$ using the master theorem [8]. Note that we can also implement all the necessary reductions to find such a spanner in linear time in the number of edges in the temporal clique. This achieves the same spanner size as in [6] but with a simpler algorithm and proof.

► **Corollary 4.8.** *We have $\mathcal{C}(n) \in \mathcal{O}(n \log n)$ and there is an algorithm that computes a spanner of size $\mathcal{O}(n \log n)$ in time $\mathcal{O}(n^2)$.*

5 Pivotal Bi-Cliques

An important structure that was used extensively in [6, 7] is the notion of a pivot. A pivot is a vertex that can be reached by all vertices until a certain time point t and that can also reach all other vertices starting at time point t . In [7] it is shown that this structure, though intuitively quite restrictive, appears with high probability in random temporal graphs.

However, there are instances without this structure [6]. Thus, we study an even more widely applicable structure, based on what we call a *c-pivot-edge* or *partial pivot-edge* (when omitting the concrete parameter). We exploit partial pivot-edges to apply the divide-and-conquer paradigm to find linear-size temporal spanners. To the best of our knowledge, this paradigm has not yet been applied to this problem. As we see later, the absence of a partial pivot-edge yields strong structural properties and it heavily restricts the class of temporal cliques where no linear spanner can be found using current techniques.

Recall the definition of In and Out in Section 2. We observe that in extremally matched bi-cliques, the sets $\text{In}(e)$ and $\text{Out}(e)$ are always distributed evenly between A and B .

► **Lemma 5.1** (\star). *Let $G = (A, B, \lambda)$ be an extremally matched bi-clique. For all $e \in A \otimes B$, we have $|\text{In}(e) \cap A| = |\text{In}(e) \cap B|$ and $|\text{Out}(e) \cap A| = |\text{Out}(e) \cap B|$.*

Next, we show that we can reach all vertices in $\text{Out}(e)$ from e starting at or after $\lambda(e)$, using exactly $|\text{Out}(e)| - 1$ edges. The analogous statement holds for the vertices in $\text{In}(e)$, in the sense that those can reach e at or after $\lambda(e)$, using only $|\text{In}(e)| - 1$ many edges.

► **Lemma 5.2.** *Let $e = \{a, b\} \in A \otimes B$. All vertices in $\text{Out}(e)$ can be organized in a tree T rooted at a , s.t. T only has edges with label at least $\lambda(e)$ and for each $v \in \text{Out}(e)$ there are temporal paths $a \rightsquigarrow v$ and $b \rightsquigarrow v$ in T . The analogous result holds for vertices in $\text{In}(e)$.*

The proof of Lemma 5.2 is already known [19] and it relies on constructing a foremost tree for $\text{Out}(e)$ (and its inverted version for $\text{In}(e)$). It was considered in [16] and lately has been used in [7], where foremost trees are applied to temporal cliques. With this idea, we can temporally connect all $u \in \text{In}(e)$ to all $v \in \text{Out}(e)$ using a linear number of edges.

► **Corollary 5.3** (\star). *Given an edge $e \in A \otimes B$, we can connect the vertices in $\text{In}(e)$ to those in $\text{Out}(e)$ by using at most $|\text{In}(e)| + |\text{Out}(e)| - 3$ many edges.*

To apply divide-and-conquer, we choose an edge $e \in A \otimes B$ and include the edges according to Corollary 5.3. Then, we create two sub-instances to connect the remaining vertices.

► **Theorem 5.4** (\star). *Let $D = (A, B, \lambda)$ be an extremally matched bi-clique with $n := |A| = |B|$. There exists a bi-spanner in D with size at most*

$$\min_{e \in A \otimes B} \left(\mathcal{D} \left(n - \frac{|\text{In}(e)|}{2} \right) + \mathcal{D} \left(n - \frac{|\text{Out}(e)|}{2} \right) + 2|\text{In}(e)| + 2|\text{Out}(e)| - 3 \right).$$

As $2|\text{In}(e)| + 2|\text{Out}(e)| - 3 \leq 8n - 3$, we see that only a linear number of edges is included to split the instance. With this, we can show that if for some edge $e \in A \otimes B$ the sets $\text{In}(e)$ and $\text{Out}(e)$ overlap, we can reduce the instance significantly. We name the set $\text{In}(e) \cap \text{Out}(e)$ the *pivot-set* of e . If $|\text{In}(e) \cap \text{Out}(e)| \geq 2cn$ for $c \in (0, 1]$, we call e a *c-pivot-edge*. When omitting the concrete constant c , we refer to e as a *partial pivot-edge*.

► **Theorem 5.5** (*c-Pivot-Edge*), (\star) . *Let $c \in (0, 1]$. Let $D = (A, B, \lambda)$ be an extremally matched bi-clique with $n := |A| = |B|$. If there is a c -pivot-edge in D , we can split the instance into two sub-instances with overall number of vertices reduced by $2cn$. Per removed vertex, we include at most $\frac{4}{c}$ edges. If applicable recursively, we get a spanner of size $\frac{8}{c}n$.*

This theorem tells us that if we repeatedly find such a c -pivot-edge $e \in A \otimes B$ with $|\text{In}(e) \cap \text{Out}(e)|$ being at least a (previously fixed) c -fraction of the current size of the instance, we get a linear-size bi-spanner for our graph. Therefore, if we search for counterexamples for linear-size spanners in temporal bi-cliques, we can assume that for any fixed c no such c -pivot-edge is present, that is, for all $e \in A \otimes B$, we have $|\text{In}(e) \cap \text{Out}(e)| \in o(n)$.

This has many interesting consequences, removing large classes of graphs from consideration. To illustrate, let us quickly mention three forbidden structures. First, for $c \in (0, 1]$, no edge e with $|\sigma_v(e) - \sigma_w(e)| \geq cn$ can exist. We call e *c-steep* because, intuitively, it makes a steep jump between the ordering at its two endpoints. Second, in a bi-clique $D = (A, B, \lambda)$ without c -steep edges, for any $i, j \in [n]$ and $a \in A$, we know that for the label spread $S := \{b \in B \mid i \leq \sigma_b(a) \leq j\}$, we have $|S| < j - i + 2cn$. Intuitively, the number of vertices for which vertex a is one of their i -th to j -th neighbors is relatively small. Of course, the same property holds when A and B are swapped. Third, we can now assume that all but sublinearly many vertices $v \in \text{In}(e)$ and all but sublinearly many $w \in \text{Out}(e)$ fulfill the condition that $\pi^+(v) \in \text{In}(e)$ and $\pi^-(w) \in \text{Out}(e)$. Intuitively, the matchings π^- and π^+ don't cross from $\text{In}(e)$ to $\text{Out}(e)$ or vice-versa.

These three consequences show that we have gained strong structural insights. They allow us and future work to focus on a more concrete set of bi-cliques. For a more detailed discussion of these consequences, please see the full version.


6 Bi-Cliques With Reverted Edges

We introduced partial pivot-edges and observed some interesting structural properties in bi-cliques which contain no suitable c -pivot-edge. The first hope would be that every bi-clique admits such a c -pivot-edge. Unfortunately, this is not the case and we provide a graph class, the *shifted matching graph*, in which for every edge, the pivot-set is as small as possible (i.e., $\text{In}(\{a, b\}) \cap \text{Out}(\{a, b\}) = \{a, b\}$) and none of our reduction rules from Section 4 apply.

Still, we present a novel technique, called *e-reverted edges*, to construct small spanners for shifted matching graphs which also provides solutions for a significantly larger class of bi-cliques. We relate the technique to partial pivot-edges from Section 5 and argue why *e-reverted edges* solve a class distinct from partial pivot-edges, beyond the concrete example of shifted matching graphs. Unfortunately, also this is not applicable to all graph classes, which we later demonstrate with a class that we call *product graphs* in Section 7.

We start by giving the definition of the shifted matching graph.

► **Definition 6.1** (*Shifted Matching Graph*). *Let $n \in \mathbb{N}^+$, $A := \{a_0, \dots, a_{n-1}\}$, and $B := \{b_0, \dots, b_{n-1}\}$. We define the shifted matching graph on n vertices per side as $\text{SM}(n) := (A, B, \lambda)$, where for every $i, j \in \{0, \dots, n-1\}$ we label the edge $\lambda(\{a_i, b_j\}) := j - i \bmod n$.*



■ **Figure 2** If at least one of the green or purple paths is temporal, $\{a', b'\}$ is $\{a, b\}$ -reverted.

In shifted matching graphs, the earliest and latest neighbors form a matching, thus we cannot apply the reduction from Theorem 4.5. Note that for all $i \in \{0, \dots, n-1\}$, we have $\pi^-(a_i) = b_i$ and $\pi^+(a_i) = b_{i-1 \bmod n}$. Additionally, we observe that each edge at time t has the form $\{a_i, b_{i+t \bmod n}\}$. Thus, the edges with label t form a perfect matching and a_i sees the b_j in ascending order of indices, circularly shifted by i places compared to a_0 .

► **Observation 6.2.** *Let $t \in \{0, \dots, n-1\}$ be a time-point. Then in $\text{SM}(n)$, the edges with label t form a perfect matching.*

We claim that this graph has no partial pivot-edges that are suitable for reduction in the sense of Theorem 5.5. To prove that this method fails on the shifted matching graph, we first characterize the set of vertices which can be reached from a vertex a_i before and after a given time. Finally, we prove that the intersection of In and Out for all edges at a_i is trivial.

► **Lemma 6.3** (\star). *Let $i, t \in \{0, \dots, n-1\}$ and a_i and $b_{i+t \bmod n}$ be vertices in $\text{SM}(n)$. Then using edges with label at most t , precisely the vertices $a_i, \dots, a_{i+t \bmod n}$ and $b_i, \dots, b_{i+t \bmod n}$ can reach a_i and $b_{i+t \bmod n}$. Using only edges with label at least t , a_i and $b_{i+t \bmod n}$ can precisely reach the vertices $b_{i+t \bmod n}, \dots, b_{i-1 \bmod n}$ and $a_{i+t+1 \bmod n}, \dots, a_i$.*

Using Lemma 6.3, we can see that using partial pivot-edges in $\text{SM}(n)$ leads to no reduction.

► **Lemma 6.4** (\star). *For vertices a_i, b_j in $\text{SM}(n)$, we have $\text{In}(\{a_i, b_j\}) \cap \text{Out}(\{a_i, b_j\}) = \{a_i, b_j\}$.*

However, with a new technique we can still construct linear-size spanners for $\text{SM}(n)$. For this, we need the concept of *e-reverted edges*. This construction sheds light on different ideas, such as making use of π^- and π^+ to switch A and B , that can be used for other graphs than just $\text{SM}(n)$. We start by giving the general technique and then proving that it is applicable to construct a linear-size bi-spanner for the shifted matching graph.

► **Definition 6.5** (*e-reverted*). *For a bi-clique $D = (A, B, \lambda)$, consider an edge $e = \{a, b\}$. We say that an edge $\{a', b'\}$ is *e-reverted*, if $a' \preceq_b \pi^+(b')$ or $\pi^-(a') \preceq_a b'$. Denote the set of edges that are not *e-reverted* as $\text{NotRev}_e := \{\{a', b'\} \in A \otimes B \mid \{a', b'\} \text{ is not } e\text{-reverted}\}$.*

Intuitively, for each edge e there is a linear-size set of edges that connects all pairs of vertices $\{a', b'\}$ that form an *e-reverted* edge. If $a' \preceq_b \pi^+(b')$, the path $a' \rightarrow b \rightarrow \pi^+(b') \rightarrow b'$ is temporal, since b' is also the latest neighbor for $\pi^+(b')$. If $\pi^-(a') \preceq_a b'$, the path $a' \rightarrow \pi^-(a') \rightarrow a \rightarrow b'$ is temporal, since a' is also the earliest neighbor $\pi^-(a')$. See Figure 2. Notice that, no matter which $a' \in A$ and $b' \in B$ we choose, the paths only use edges incident to a or b and edges from π^- or π^+ . We use this to construct a small bi-spanner for D , if NotRev_e is small.

11:12 How to Reduce Temporal Cliques to Find Sparse Spanners

We do not prove Theorem 7.1 right away. Rather, we first consider an application. For this, choose $f(n) := n^{1-\varepsilon}$. Theorem 7.1 tells us that there is a set of graphs $\{D_n\}_{n \in \mathbb{N}}$ such that the size of sides is in $\Theta(n)$, the size of any pivot-set is in $\mathcal{O}(n^{1-\varepsilon})$, and the number of not reverted edges with respect to any edge is in $\Omega(n^{2-\varepsilon})$.

To prove Theorem 7.1, we define a graph class that we later use to define appropriate D_n .

► **Definition 7.2 (Product Graph).** *Consider two bi-cliques $G = (A_G, B_G, \lambda_G)$ and $H = (A_H, B_H, \lambda_H)$. Define the product graph $G \times H = (\mathcal{A}, \mathcal{B}, \Lambda)$ as*

$$\begin{aligned} \mathcal{A} &:= A_G \times A_H, \\ \mathcal{B} &:= B_G \times B_H, \\ \Lambda &:= ((a_G, a_H), (b_G, b_H)) \mapsto (\lambda_G(a_G, b_G), \lambda_H(a_H, b_H)). \end{aligned}$$

Note that the co-domain of Λ is \mathbb{N}^2 . To make it the proper definition of an edge labeling function, take the lexicographic embedding. For the reader's benefit, we continue to write tuples. Intuitively, to construct $G \times H$, we start with the graph G and replace each vertex of A_G with a distinct copy of A_H as well as each vertex of B_G with a distinct copy of B_H , see Figure 3a for an illustration of this idea. We refer to each such expanded vertex as a *bag*. Note that this operation is not commutative. The edge labels Λ are chosen in such a way that any 2-edge path $x \rightarrow y \rightarrow z$ with x and z in the same bag is temporal if the respective path in H is temporal and any 2-edge path $x \rightarrow y \rightarrow z$ with x and z belonging to different bags is temporal if the respective path is temporal in G . We now make this rigorous.

► **Lemma 7.3 (\star).** *Let G and H be bi-cliques. Let $P = (g_1, h_1)(g_2, h_2) \dots (g_\ell, h_\ell)$ be a temporal path in $G \times H$. Then*

1. *the sequence $g_1 g_2 \dots g_\ell$ is a temporal path in G ;*
2. *if $\lambda_G(g_1, g_2) = \lambda_G(g_{\ell-1}, g_\ell)$, the sequence $h_1 h_2 \dots h_\ell$ is a temporal path in H .*

This gives us enough tools to prove how the pivot-set of an edge $e = \{(u_G, u_H), (v_G, v_H)\} \in E(G \times H)$ in $G \times H$ relates to the pivot-set of $\{u_G, v_G\}$ in G . For this, define $\phi: V(G \times H) \rightarrow V(G)$ as $(u_G, u_H) \xrightarrow{\phi} u_G$. This denotes the function that yields for every vertex in $G \times H$ its bag, namely, its original vertex in G . We will also naturally extend ϕ to edges (i.e., $\phi(\{(u_G, u_H), (v_G, v_H)\}) = \{u_G, v_G\}$). Intuitively, for any vertex v in the pivot-set of an edge $e \in E(G \times H)$, the vertex $\phi(v)$ must be in the pivot-set of $\phi(e)$ in G .

► **Lemma 7.4 (\star).** *Let G and H be bi-cliques and let $e \in E(G \times H)$. We have*


1. $\phi(\text{In}(e)) \subseteq \text{In}(\phi(e))$,
2. $\phi(\text{Out}(e)) \subseteq \text{Out}(\phi(e))$, and
3. $\phi(\text{In}(e) \cap \text{Out}(e)) \subseteq \text{In}(\phi(e)) \cap \text{Out}(\phi(e))$.

We now focus on product graphs, where both G and H are shifted matching graphs, as we will use these to prove Theorem 7.1. Therefore, they will serve as the family of graphs which the already introduced techniques do not solve.

► **Definition 7.5.** *Let $m, k \in \mathbb{N}$. Define the graph $\text{SM}(m, k) = \text{SM}(m) \times \text{SM}(k)$ to be the product graph obtained by taking $\text{SM}(m)$ as the outer graph and $\text{SM}(k)$ as the inner graph.*

Observe that each side of $\text{SM}(m, k)$ has m bags each of size k and so, overall, $\text{SM}(m, k)$ has mk vertices per side. We now investigate the size of the pivot-sets in $\text{SM}(m, k)$ as well as the number of the non-reverted edges.

► **Lemma 7.6 (\star).** *Let $m, k \in \mathbb{N}$ and consider $\text{SM}(m, k)$. For any $e \in E(\text{SM}(m, k))$, we have that the size of the pivot-set $\text{In}(e) \cap \text{Out}(e)$ is at most $2k$.*



■ **Figure 4** Overview of techniques. Previously proposed techniques are blue, new ones are green. Example families of graphs are purple. The red area contains instances that are not solved by any known techniques.

► **Lemma 7.7** (\star). *Let $k, m \in \mathbb{N}$ and consider $SM(m, k)$. For any $e \in E(SM(m, k))$, the size of the set of not e -reverted edges NotRev_e is at least $(m-1)\binom{k}{2}$.*

The last two lemmas show that our tools are not strong enough to solve this graph class yet. With our previous techniques, we are not able to give a linear-size spanner for every product graph. By Theorem 7.1, for $f(n) \in \Theta(\sqrt{n})$, Theorem 5.4 and Theorem 6.6 only give us $\Theta(n^{3/2})$ spanners. Still, product graphs exhibit a lot of structure, in the sense that all edges between two bags have roughly the same time labels. We use this to construct bi-spanners for the product graph based on bi-spanners for the two underlying graphs.

► **Lemma 7.8** (\star). *Let $G = (A_G, B_G, \lambda_G)$ and $H = (A_H, B_H, \lambda_H)$ be temporal bi-cliques of sizes n_G and n_H and let S_G, S_H be temporal bi-spanners respectively. Then, there is a bi-spanner in $G \times H$ of size at most $|S_G| \cdot |S_H|$.*

With more careful analysis, this construction can be improved to a bound of $|S_G|n_H + |S_H|n_G$.

If G and H each have linear-size bi-spanners, this proves the existence of a linear-size bi-spanner for $G \times H$. Since we know from Section 6 how to construct a linear-size bi-spanner for $SM(m)$, we can now construct a linear-size bi-spanner for $SM(m, k)$.

Additionally, we can augment Lemma 7.8 to also deal with different graphs H_e per edge e in G . To do this, we have to change the definition of our representatives from a_H^* and b_H^* to functions r_A, r_B , such that for $b_G \in B_G$ the label of $\{(\pi_{S_G}^+(b_G), r_A(b_G)), (b_G, r_B(b_G))\}$ is minimal. This enables us not only to construct bi-spanners for $G \times H$ but any bi-cliques that can be decomposed into bags of vertices.

8 Conclusion

We present a significant step forwards towards answering whether temporal cliques admit linear-size temporal spanners, see Figure 4 for an illustration. The obvious problem that stems from this paper, is whether classes of graphs exist that are not solved by our techniques. We conjecture that this is the case but the classes of graphs remaining should be quite artificial. Another possible avenue is to generalize some of our techniques to include a wider variety of graphs. For example, we could consider investigating the temporal ordering of edges and try to generalize the concept of reverted edges. Also, product graphs and their solution suggest a way in which temporal bi-cliques may be composed and decomposed. Future research could generalize this idea and more closely analyze the cases in which a temporal spanner can be obtained by combining temporal spanners on subinstances. Towards the opposite direction, it is now easier to find a $\Omega(n \log n)$ counterexample, if it exists at all, since we have provided many properties that any counterexample must avoid.

References

- 1 Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis, and Viktor Zamaraev. Temporal vertex cover with a sliding time window. *Journal of Computer and System Sciences*, 107:108–123, 2020. doi:10.1016/j.jcss.2019.08.002.
- 2 Kyriakos Axiotis and Dimitris Fotakis. On the Size and the Approximability of Minimum Temporally Connected Subgraphs. In *International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, pages 149:1–149:14, 2016. doi:10.4230/LIPIcs.ICALP.2016.149.
- 3 Davide Bilò, Gianlorenzo D’Angelo, Luciano Gualà, Stefano Leucci, and Mirko Rossi. Blackout-tolerant temporal spanners. In *Algorithmics of Wireless Networks (ALGOSENSORS 2022)*, pages 31–44, 2022.
- 4 Davide Bilò, Gianlorenzo D’Angelo, Luciano Gualà, Stefano Leucci, and Mirko Rossi. Sparse Temporal Spanners with Low Stretch. In *European Symposium on Algorithms (ESA 2022)*, pages 19:1–19:16, 2022. doi:10.4230/LIPIcs.ESA.2022.19.
- 5 Davide Bilò, Sarel Cohen, Tobias Friedrich, Hans Gawendowicz, Nicolas Klodt, Pascal Lenzner, and George Skretas. Temporal network creation games. In *International Joint Conference on Artificial Intelligence (IJCAI 2023)*, pages 2511–2519, 2023. doi:10.24963/ijcai.2023/279.
- 6 Arnaud Casteigts, Joseph G Peters, and Jason Schoeters. Temporal cliques admit sparse spanners. In *International Colloquium on Automata, Languages and Programming (ICALP 2019)*, pages 134–1, 2019. doi:10.4230/LIPIcs.ICALP.2019.134.
- 7 Arnaud Casteigts, Michael Raskin, Malte Renken, and Viktor Zamaraev. Sharp thresholds in random simple temporal graphs. In *Symposium on Foundations of Computer Science (FOCS 2022)*, pages 319–326, 2022. doi:10.1109/FOCS52979.2021.00040.
- 8 Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2022.
- 9 Argyrios Deligkas, Eduard Eiben, Tiger-Lily Goldsmith, and George Skretas. Being an influencer is hard: The complexity of influence maximization in temporal graphs with a fixed source. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, pages 2222–2230, 2023. doi:10.5555/3545946.3598898.
- 10 Argyrios Deligkas, Eduard Eiben, and George Skretas. Minimizing reachability times on temporal graphs via shifting labels. In *International Joint Conference on Artificial Intelligence (IJCAI 2023)*, pages 5333–5340, 2023. doi:10.24963/ijcai.2023/592.
- 11 Riccardo Dondi and Alexandru Popa. Timeline cover in temporal graphs: Exact and approximation algorithms. In *Combinatorial Algorithms (IWOCA 2023)*, pages 173–184, 2023. doi:10.1007/978-3-031-34347-6_15.
- 12 Jessica A. Enright, Kitty Meeks, George B. Mertzios, and Viktor Zamaraev. Deleting edges to restrict the size of an epidemic in temporal networks. *Journal of Computer and System Sciences*, 119:60–77, 2021. doi:10.1016/j.jcss.2021.01.007.
- 13 Eugen Fùchsle, Hendrik Molter, Rolf Niedermeier, and Malte Renken. Delay-robust routes in temporal graphs. In *International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*, pages 30:1–30:15, 2022. doi:10.4230/LIPIcs.STACS.2022.30.
- 14 Thekla Hamm, Nina Klobas, George B. Mertzios, and Paul G. Spirakis. The complexity of temporal vertex cover in small-degree graphs. In *AAAI Conference on Artificial Intelligence (AAAI 2022)*, pages 10193–10201, 2022. doi:10.1609/aaai.v36i9.21259.
- 15 David Kempe, Jon Kleinberg, and Amit Kumar. Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences*, 64(4):820–842, 2002. doi:10.1006/jcss.2002.1829.
- 16 Mikhail Lavrov and Po-Shen Loh. Increasing hamiltonian paths in random edge orderings. *Random Structures & Algorithms*, 48(3):588–611, 2016. doi:10.1002/rsa.20592.
- 17 George B. Mertzios, Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Temporal network optimization subject to connectivity constraints. In *International Colloquium on*

- Automata, Languages and Programming (ICALP 2013)*, pages 657–668, 2013. doi:10.1007/978-3-642-39212-2_57.
- 18 Othon Michail. An introduction to temporal graphs: An algorithmic perspective. *Internet Mathematics*, 12(4):239–280, 2016. doi:10.1080/15427951.2016.1177801.
 - 19 B Bui Xuan, Afonso Ferreira, and Aubin Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *International Journal of Foundations of Computer Science*, 14(02):267–285, 2003. doi:10.1142/S0129054103001728.