



Blackout-tolerant temporal spanners [☆]

Davide Bilò ^a, Gianlorenzo D'Angelo ^b, Luciano Gualà ^c, Stefano Leucci ^{a,*},
Mirko Rossi ^b

^a Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Via Vetoio s.n.c., L'Aquila, 67100, Italy

^b Gran Sasso Science Institute, Viale Francesco Crispi, 7, L'Aquila, 67100, Italy

^c Department of Enterprise Engineering, University of Rome "Tor Vergata", Via del Politecnico 1, Rome, 00133, Italy

ARTICLE INFO

Article history:

Received 16 December 2022

Received in revised form 20 November 2023

Accepted 27 November 2023

Available online 4 December 2023

Keywords:

Temporal graphs

Temporal spanners

Fault-tolerance

ABSTRACT

We introduce the notions of *blackout-tolerant* temporal α -spanner of a temporal graph G which is a subgraph of G that preserves the distances between pairs of vertices of interest in G up to a multiplicative factor of α , even when the graph edges at a single time-instant become unavailable. In particular, we consider the *single-source*, *single-pair*, and *all-pairs* cases and, for each case we look at three quality requirements: *exact* distances (i.e., $\alpha = 1$), *almost-exact* distances (i.e., $\alpha = 1 + \varepsilon$ for an arbitrarily small constant $\varepsilon > 0$), and *connectivity* (i.e., unbounded α). We provide almost tight bounds on the size of such spanners for *general* temporal graphs and for *temporal cliques*, showing that they are either very sparse (i.e., they have $\tilde{O}(n)$ edges) or they must have size $\Omega(n^2)$ in the worst case, where n is the number of vertices of G . We also investigate multiple blackouts and *k-edge fault-tolerant temporal spanners*.

© 2023 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

In wireless and sensor networks the communication links among nodes frequently change over time due to the fact that hosts might move, be active at different times, or face interference. Recently, this dynamic behavior has been modeled through so-called *temporal graphs*, which are graphs where the edge-set is allowed to change over time. There are multiple definitions of temporal graphs in the literature, with the simplest one being that of Kempe, Kleinberg, and Kumar [17] in which each edge of a graph $G = (V, E)$ has an assigned time-label $\lambda(e) \in \mathbb{N}^+$ representing the instant in which $e \in E$ can be used. A path from a vertex to another in G is said to be a *temporal path* if its traversed edges have non-decreasing time-labels. If there exists a temporal path from u to v , for every two vertices $u, v \in V$, then the graph is *temporally connected*.

One of the main problems in network design is reducing the size (and hence the operational cost) of a network while preserving both its robustness to failures and its communication efficiency.

In static networks this problem has been extensively studied, and has been formalized using the notion of (*fault-tolerant*) *graph spanners* i.e., sparse subgraphs that approximately preserve distances between pairs of vertices of interest, possibly in the presence of edge/vertex failures. While the landscape in static networks is quite well-understood (see, e.g., [1] and the

[☆] This work has been partially funded by the European Union - NextGenerationEU under the Italian Ministry of University and Research (MUR) National Innovation Ecosystem grant ECS00000041 - VITALITY - CUP: D13C21000430001 and by ARS01_00540 - RASTA project, funded by the Italian Ministry of Research PNR 2015-2020; The second author acknowledges the support of the MUR (Italy) Department of Excellence 2023-2027.

* Corresponding author.

E-mail addresses: davide.bilo@univaq.it (D. Bilò), gianlorenzo.dangelo@gssi.it (G. D'Angelo), guala@mat.uniroma2.it (L. Gualà), guala@mat.uniroma2.it (S. Leucci), mirko.rossi@gssi.it (M. Rossi).

<https://doi.org/10.1016/j.jcss.2023.103495>

0022-0000/© 2023 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

	1-BT temporal clique			1-BT	2-BT t. clique	2-BT	3-BT t. clique
	single-pair	single-source	all-pairs	single-pair	single-pair	single-pair	single-pair
temporal spanner	$\uparrow$	$\Omega(n^2)$ Theorem 5	$\rightarrow$	$O(n)$ Theorem 1	$\Rightarrow$ Observation 3	$\Omega(n^2)$ Theorem 6	$\Rightarrow$ Observation 2
temporal $(1+\varepsilon)$ -spanner	$\uparrow$	$\downarrow$	$\downarrow$	$O\left(\frac{n \log^4 n}{\log(1+\varepsilon)}\right)$ Theorem 2	$\Rightarrow$ Observation 3	$\downarrow$	$\downarrow$
temporal preserver	$O(n)$ Observation 1	$\downarrow$	$\downarrow$	$\Omega(n^2)$ Theorem 3	$\Rightarrow$ Observation 2	$\downarrow$	$\downarrow$

Fig. 1. A summary of our results for blackout-tolerant temporal spanners along the two considered dimensions. Each combination that admits a blackout-tolerant temporal spanner of size $\tilde{O}(n)$ is shaded in blue, while those for which a strong lower bound of $\Omega(n^2)$ holds are shaded in red. The actual bounds are shown in the corresponding cells. A single-tailed arrow from a cell A to a neighboring cell B means that the upper/lower bound of A immediately implies the same result for B . A double-tailed arrow from A to B means that the result of A can be shown to also hold for B once it is paired with some additional minor observation. Missing columns (e.g., 1-BT single-source) only contain lower-bound results directly implied from an existing column (e.g., 1-BT single-source temporal-clique). Color figures can be found in the web version of this article.

references therein), the problem is still under active investigation in temporal networks, where it is even more relevant. The problem of computing a sparse *temporal spanner*, i.e., a sparse temporal subgraph of a given temporal graph that preserves temporal connectivity between pairs of vertices, was first introduced in the seminal paper of Kempe, Kleinberg, and Kumar [17] in 2002, and has recently received considerable attention. In particular, [2] has shown that there are temporal graphs with $\Theta(n^2)$ edges that cannot be sparsified. This prompted [10] to focus on *temporal cliques*, i.e., complete temporal graphs. Here the situation drastically improves as any temporal clique admits a temporal spanner with $O(n \log n)$ edges, although the resulting *stretch factor*, i.e., the maximum ratio between the length¹ of the shortest temporal path in the spanner and the corresponding shortest temporal path in the original graph, can be $\Omega(n)$.

This motivated [3] to study temporal α -spanners, i.e., temporal spanners that also guarantee a stretch factor of at most α . In particular, the authors showed that temporal cliques always admit a temporal $(2k-1)$ -spanner with $\tilde{O}(kn^{1+\frac{1}{k}})$ edges,² where $k \geq 1$ is an integer parameter of choice. They also considered *single-source* temporal α -spanners, i.e., temporal α -spanners in which the upper bound of α on the stretch factor only needs to hold for distances from a given source, and showed that any general temporal graph admits a single-source temporal $(1+\varepsilon)$ -spanner with $\tilde{O}(n/\log(1+\varepsilon))$ edges, for any $0 < \varepsilon < n$.

In all the above results it is implicitly assumed that the temporal network remains *fault-free*, i.e., each edge e is available with certainty at the time instant specified by its time-label $\lambda(e)$. Real-world networks, however, are fault-prone and a link could fail to operate properly at the scheduled time (e.g., a connection fails to be established). Even worse, due to cascading failures or catastrophic events, the entire network might become unavailable for a limited amount of time.

In this paper we investigate the problem of computing sparse temporal spanners with low stretch in presence of a failure in all the edges with single unknown time-label τ . We name this kind of failure a *blackout* at time τ , and we say that a temporal subgraph H of a temporal graph G is a *blackout-tolerant* (BT) temporal spanner if, for every blackout τ , the surviving part $H^{-\tau}$ of H is a temporal spanner of the surviving part $G^{-\tau}$ of G . In this setting, the stretch is defined as the largest ratio between the length of a shortest temporal path in $H^{-\tau}$ and the length of the corresponding shortest temporal path in $G^{-\tau}$, among all blackouts τ . We also investigate the case of *b-blackout tolerant* (*b*-BT) spanners, i.e., spanners in which up to b blackouts can occur, and that of *k-edge fault-tolerant* (*k*-EFT) temporal spanners, i.e., spanners that are resilient to the failure of any set of at most k edges.

Our results. Our focus is on the case of a single blackout and we consider different kinds of temporal spanners on both *general temporal graphs* and *temporal cliques*. We classify our temporal spanners along two dimensions. The first dimension pertains the pairs of vertices of interest: all our spanners provide some connectivity or distance guarantee from any vertex in some set $S \subseteq V$ to any vertex in some set $T \subseteq V$. We say that a temporal spanner is *all-pairs* if $S = T = V$; *single-source* if $S = \{s\}$, for some $s \in V$, and $T = V$; and *single-pair* if $S = \{s\}$ and $T = \{t\}$ for some $s, t \in V$. If not otherwise specified we refer to all-pairs spanners.

The second dimension concerns whether the temporal spanner is required to only preserve connectivity from the vertices in S to the vertices in T (in which case we simply talk about temporal spanners) or to guarantee α -approximate distances from the nodes in S to the nodes in T (i.e., it is a temporal α -spanner). The special case $\alpha = 1$ corresponds to preserving exact distances and we refer to the corresponding temporal spanner as a *temporal preserver*.

¹ As in [3], the length of a temporal path is its number of edges.

² The notation $\tilde{O}(f(n))$ is a shorthand for $O(f(n) \log^c f(n))$, for some constant $c > 0$.

	1-EFT temporal clique			1-EFT		2-EFT clique		2-EFT	3-EFT clique
	single-pair	single-source	all-pairs	single-pair	single-source	single-pair	single-source	single-pair	single-pair
temporal spanner	$\uparrow$	$O(n)$ Theorem 9		$O(n)$ Theorem 7	$\Omega(n^2)$ Observation 5	$O(n)$ Observation 9	$\Omega(n^2)$ Observation 8	$\Omega(n^2)$ Observation 7	$\Omega(n^2)$ Observation 8
temporal (1+ ε)-spanner	$\uparrow$			$O\left(\frac{n \log^4 n}{\log(1+\varepsilon)}\right)$ Theorem 8	$\downarrow$	$O\left(\frac{n \log^4 n}{\log(1+\varepsilon)}\right)$ Observation 9	$\downarrow$	$\downarrow$	$\downarrow$
temporal preserver	$O(n)$ Observation 6	$\Omega(n^2)$ Theorem 10	$\rightarrow$	$\Omega(n^2)$ Observation 4	$\downarrow$	$\Omega(n^2)$ Observation 8	$\downarrow$	$\downarrow$	$\downarrow$

Fig. 2. A summary of our results for edge fault-tolerant temporal spanners, using the same graphical notation of Fig. 1. The gray area corresponds to cases left open. Missing columns (e.g., 1-EFT all pairs) only contain lower-bound results directly implied from an existing column. Color figures can be found in the web version of this article.

From a high-level perspective our results show that, for each of the above combinations, the sparsest admissible blackout-tolerant temporal spanner is either very sparse, i.e., it contains $\tilde{O}(n)$ edges, or it must have size $\Omega(n^2)$ for some worst-case family of temporal graphs with n vertices. This implies that all our upper and lower bounds are asymptotically optimal up to polylogarithmic factors.

In more detail, our main results are the following (see also the cells with the bold references in Fig. 1):

- Any temporal graph admits a BT single-pair temporal spanner of linear size. Such a temporal spanner can be computed in polynomial time.
- For any $0 < \varepsilon < n$, we can compute in polynomial time a BT single-pair temporal $(1 + \varepsilon)$ -spanner of size $O\left(\frac{n \log^4 n}{\log(1+\varepsilon)}\right)$.
- There exists a temporal graph G such that any BT single-pair temporal preserver of G has size $\Omega(n^2)$.
- The above result can be extended to show that there exists a *temporal clique* G such that any BT single-source temporal spanner of G has size $\Omega(n^2)$.

For the case of multiple blackouts we adapt our lower bound of $\Omega(n^2)$ on the size of temporal spanners for the following cases: (i) 2-BT single-pair temporal spanners, (ii) 2-BT single-pair temporal preservers of temporal cliques, and (iii) 3-BT single-pair temporal spanners of temporal cliques.

Moreover, we observe that, in the case of two blackouts, temporal cliques admit 2-BT single-pair temporal spanners and 2-BT single-pair temporal $(1 + \varepsilon)$ -spanners of size $O(n)$ and $O(n \log^4 n / \log(1 + \varepsilon))$, respectively.

Thanks to the above results and some additional observations that allow us to import known bounds from the literature, we can achieve a characterization of the landscape of blackout-tolerant temporal spanners along the considered dimensions that is tight up to polylogarithmic factors. We summarize the resulting upper and lower bounds in Fig. 1. All our bounds extend to the case in which each edge can have multiple time-labels (i.e., it can be used in multiple time-instants).

Finally, we also provide some preliminary results for the case of k -edge fault-tolerant temporal spanners, i.e., spanners that provide connectivity or distance guarantees following the failure of any set of at most k edges. As for blackout-tolerant spanners we consider single-source, single-pair, and all-pairs k -EFT spanners on both temporal cliques and general temporal graphs. These results are summarized in Fig. 2, we observe that for almost all cases, the bounds on the size of a k -EFT temporal spanner coincide with those of a k -BT temporal spanner. One case where the bounds differ, is that of single-source 1-EFT temporal spanners on temporal cliques, for which we provide an upper bound of $O(n)$. We leave open the problems of bounding the size of 1-EFT single-source temporal $(1 + \varepsilon)$ -spanners of temporal cliques, and 1-EFT all-pairs temporal spanners and $(1 + \varepsilon)$ -spanners of temporal cliques.

Other related work. In the case of static graphs, fault-tolerant spanners have been widely studied in the presence of edge failures also in the single-pair and single-source cases [1,5–7,14,20,21].

Regarding temporal graphs, besides the aforementioned papers, the problem of understanding the size of a temporal spanner for random temporal graphs has been considered in [11], where the authors study Erdős-Rényi graphs $G_{n,p}$ in which the time-label of an edge is its rank in a random permutation of the graph's edges, and provide sharp thresholds on the probability p for which the graph contains sparse single-pair, single-source, and (connectivity-preserving) temporal spanners.

For temporal graphs multiple distance measures are natural. In addition to the one considered in this work, i.e., the minimum number of edges of a temporal path, other common choices are the *earliest arrival time*, the *latest departure time*, *shortest time*, and *fastest time* (see, e.g., [8,9,18,22]). As far as temporal spanners are considered, [3] shows that these measures result in strong lower bounds on the size of any temporal subgraph preserving all-pair connectivity, even when the considered temporal graph is a clique.

The study of temporal-graph models and algorithms is a wide area of research, and we refer the interested reader to [19] for a survey and to [15] for a discussion on alternative models. Finally, other problems that aim to compute robust paths in temporal graphs that can be subject to disruptions have been studied in [12,13].

Paper organization. Section 2 describes our model of temporal graphs and gives some preliminary definitions. Section 3 and Section 4 present our results for single-pair and single-source temporal spanners in the case of a single blackout, respectively. We discuss the case of multiple blackouts in Section 5 and the case of edge fault-tolerant spanners in Section 6.

2. Model and preliminaries

Let $G = (V, E, \lambda)$ be an undirected graph of n vertices and a labeling function $\lambda : E \rightarrow \mathbb{N}^+$ that assigns a *time-label* $\lambda(e)$ to each edge e . We call such a graph *temporal graph*. If G is complete we will say that it is a *temporal clique*. A temporal path π from vertex u to vertex v is a path in G from u to v such that the sequence $e_1, e_2, \dots, e_k$ of edges traversed by π satisfies $\lambda(e_i) \leq \lambda(e_{i+1})$ for all $i = 1, \dots, k-1$. For a given (non-empty) temporal path π , the departure time of π and the arrival time of π are the time-labels of the first and last edge of π , respectively. For technical convenience we define the departure time of an empty path as $+\infty$ and the arrival time of an empty path as $-\infty$. A temporal path π_1 from u to v is *compatible* with a temporal path π_2 from v to z if the arrival time of π_1 does not exceed the departure time of π_2 , and we denote by $\pi_1 \circ \pi_2$ the temporal path from u to z obtained by concatenating π_1 and π_2 . If a simple temporal path π traverses the vertices u and v , in this order, then we denote by $\pi[u : v]$ the (temporal) subpath of π from u to v .

We say that a vertex u can *reach* a vertex v if there exists a temporal path from u to v in G . The *length* of a (not necessarily temporal) path π is the number of the edges in π and it is denoted by $|\pi|$. We define the *distance* $d_G(u, v)$ from u to v in G as the length of the shortest temporal path from u to v (in G). If u cannot reach v , then $d_G(u, v) = +\infty$.

A temporal graph $H = (V', E', \lambda')$ is a (temporal) subgraph of $G = (V, E, \lambda)$ if (V', E') is a subgraph of (V, E) and $\lambda'(e) = \lambda(e)$, for all $e \in E'$. We denote by $G^{\leq \tau}$, $G^{\geq \tau}$, and $G^{-\tau}$, the subgraphs of G that contain all vertices and exactly the edges with a time-label at most τ , at least τ , and different from τ , respectively. For the sake of readability we slightly abuse the notation for distances by moving the superscripts of the considered graph to the function d itself, e.g., we may write $d_G^{\leq \tau}(u, v)$ to denote the distance from u to v in the graph $G^{\leq \tau}$. When the graph G is clear from the context we may omit it altogether and write, e.g., $d^{\leq \tau}(u, v)$. The lifetime L of G is the maximum time-label assigned to an edge of G . The *time-reversed* version of G is the temporal graph obtained from G by replacing each time-label $\lambda(e)$ with $L - \lambda(e) + 1$. One can observe that reversing a temporal path from u to v with minimum arrival time in G yields a temporal path from v to u with maximum departure time in the time-reversed version of G , and vice-versa.

Given two sets of vertices $S, T \subseteq V$, an (S, T) -temporal spanner of G is a (temporal) subgraph H of G such that for every $u \in S$ and $v \in T$, u can reach v in H if and only if u can reach v in G . We call an (S, T) -temporal spanner: (i) *single-pair* temporal spanner if $S = \{s\}$ and $T = \{t\}$ for some vertices $s, t \in V$; (ii) *single-source* temporal spanner if S contains a single vertex $s \in V$ and $T = V$, (iii) *temporal spanner* if $S = T = V$.

We also define a different kind of *temporal spanner* called *temporal α -spanner*, where instead of preserving the reachability between two sets of vertices, we want to approximate their distances up to a factor $\alpha \geq 1$. More formally, for $\alpha \geq 1$ and two sets of vertices $S, T \subseteq V$, an (S, T) -temporal α -spanner of G is a (temporal) subgraph H of G such that $d_H(u, v) \leq \alpha \cdot d_G(u, v)$, for all $u \in S$ and $v \in T$. If $\alpha = 1$ we refer to an (S, T) -temporal 1-spanner as (S, T) -temporal preserver. The *size* of a temporal spanner (or α -spanner) is the number of its edges.

A *blackout* at time $\tau \in [1, L]$ is the failure of all the edges with time-label τ .

An (S, T) -temporal spanner (resp. (S, T) -temporal α -spanner) H of G is *blackout-tolerant* (BT) if for every $\tau \in [1, L]$, $H^{-\tau}$ is an (S, T) -temporal spanner (resp. (S, T) -temporal α -spanner) of $G^{-\tau}$.

3. Blackout-tolerant single-pair temporal spanners

In this section we provide our upper and lower bounds on the size of single-pair temporal spanners and preservers in the case of a single blackout.

3.1. Upper bounds on BT single-pair temporal spanners

Given a temporal graph G and two vertices $s, t \in V$, we provide a single-pair BT temporal spanner H of G with source s and target t . We define H as the union of two subgraphs A and D of G . In particular, A is a subgraph of G such that for every vertex v , A contains a temporal path from s to v with minimum arrival time. Symmetrically, D is a subgraph of G such that for every vertex v , D contains a temporal path from v to t with maximum departure time. From [8,16], we know that we can compute, in polynomial time, a subgraph of size $O(n)$ containing a path with minimum arrival time from a given source to all other vertices.³ By using the above result on G and the time-reversed version of G , respectively, we have that both A and D have size $O(n)$ and can be found in polynomial time. We have:

³ This is true even for temporal graphs with multiple time-labels.

Theorem 1. A blackout-tolerant single-pair temporal spanner of G of size $O(n)$ can be computed in polynomial time.

Proof. We only need to argue that H , as defined above, is a single-pair temporal spanner of G w.r.t. the two nodes s and t .

To this aim, fix a time-label $\tau \in [1, L]$ such that there exists a temporal path π_τ between s and t in $G^{-\tau}$. Let the vertices of π_τ be $s = v_0, v_1, \dots, v_k = t$, as traversed from s to t and define d as the largest integer in $[1, k]$ such that $\lambda((v_{d-1}, v_d)) < \tau$. If no such integer exists, let $d = 0$.

As a consequence $\pi_\tau[s : v_d]$ has an arrival time of at most $\tau - 1$. This shows that A must contain a temporal path π_A from s to v_d with arrival time at most $\tau - 1$, i.e., π_A is in $A^{-\tau}$. Similarly, $\pi_\tau[v_d : t]$ has a departure time of at least $\tau + 1$. Thus, D must contain a temporal path π_D from v_d to t with departure time at least $\tau + 1$, i.e., π_D is in $D^{-\tau}$. Since π_A and π_D are compatible, their concatenation yields a temporal path from s to t that is entirely contained in $H^{-\tau}$. $\square$

We can modify the above construction to obtain, for any $\varepsilon > 0$ of choice, a blackout-tolerant single-pair temporal $(1 + \varepsilon)$ -spanner H of G w.r.t. s and t . To this aim, we choose A and D as two temporal subgraphs of G that have size $O(\frac{n \log^4 n}{\log(1+\varepsilon)})$ and, satisfy the following for all $v \in V$ and $\tau \in [1, L]$:

$$d_A^{\leq \tau}(s, v) \leq (1 + \varepsilon) \cdot d_G^{\leq \tau}(s, v) \text{ and } d_D^{\geq \tau}(v, t) \leq (1 + \varepsilon) \cdot d_G^{\geq \tau}(v, t).$$

These subgraphs can be found in polynomial time using the results in [3].^{3,4} We can now show that the temporal graph H obtained as the union of A and D contains $(1 + \varepsilon)$ -approximate blackout-tolerant temporal paths from s to t .

Theorem 2. Let $0 < \varepsilon < n$. A blackout-tolerant single-pair temporal $(1 + \varepsilon)$ -spanner of size $O(\frac{n \log^4 n}{\log(1+\varepsilon)})$ can be computed in polynomial time.

Proof. We only need to argue that H , with the updated definitions of A and D , is a single-pair temporal $(1 + \varepsilon)$ -spanner of G w.r.t. s and t .

Fix a time label $\tau \in [1, L]$ such s can reach t in $G^{-\tau}$, and let π_τ be a *shortest* temporal path from s to t in $G^{-\tau}$. Let $s = v_0, v_1, \dots, v_k = t$ be the vertices traversed by π_τ , in order, and d as in the proof of Theorem 1 so that the temporal path $\pi_\tau[s : v_d]$ has an arrival time of at most $\tau - 1$, while the temporal path $\pi_\tau[v_d : t]$ has a departure time of at least $\tau + 1$.

By our choice of A , there must be a temporal path π_A in A such that the arrival time of π_A is at most $\tau - 1$ and $|\pi_A| \leq (1 + \varepsilon) \cdot d_G^{\leq \tau-1}(s, v_d) \leq (1 + \varepsilon) \cdot |\pi_\tau[s : v_d]|$. Similarly, by our choice of D , there must be a temporal path π_D in D such that the departure time of π_D is at least $\tau + 1$ and $|\pi_D| \leq (1 + \varepsilon) \cdot d_G^{\geq \tau+1}(v_d, t) \leq (1 + \varepsilon) \cdot |\pi_\tau[v_d : t]|$. Since π_A and π_D are compatible temporal paths and are both contained in $H^{-\tau}$, we have:

$$\begin{aligned} d_H^{-\tau}(s, t) &\leq |\pi_A \circ \pi_D| = |\pi_A| + |\pi_D| \leq (1 + \varepsilon) \cdot |\pi_\tau[s : v_d]| + (1 + \varepsilon) \cdot |\pi_\tau[v_d : t]| \\ &= (1 + \varepsilon) \cdot |\pi_\tau[s : v_d] \circ \pi_\tau[v_d : t]| = (1 + \varepsilon) |\pi_\tau| = (1 + \varepsilon) d_G^{-\tau}(s, t). \quad \square \end{aligned}$$

3.2. Lower bounds on BT single-pair temporal preservers

In this section, we show that there exists a temporal graph G of $\Theta(n)$ vertices such that, any blackout-tolerant single-pair temporal preserver w.r.t. a fixed pair of vertices s and t of G has size $\Theta(n^2)$.

We will make use of a dense graph construction which we import from [3].⁵

Lemma 1 ([3]). Given n , there exists a temporal graph G' with $\Theta(n)$ vertices and $\Omega(n^2)$ edges that satisfies the following properties:

- G' is the union of n edge-disjoint temporal paths $\pi_1, \dots, \pi_n$ such that, each π_i is a path from a fixed source vertex σ to a distinct target vertex $t_i \neq \sigma$.
- π_i uses only edges with time-label i ;
- π_i is the unique shortest temporal path from σ to t_i ;
- the largest time-label of an edge incident to vertex t_i is i ;
- $|\pi_i| < |\pi_j| - (i - j)$ for all $j < i$;
- any single-source temporal preserver of G' with source σ has size $\Omega(n^2)$.

⁴ A shortest path between x and y with departure time at least τ is a shortest path from y to x with arrival time $L - \tau + 1$ in the time-reversed version of G (where L is the lifetime of G).

⁵ For ease of presentation, we summarize the key properties of the construction in a single lemma. The details can be found in Section 4.3 of the full version [4] of [3].




Fig. 3. Sketch of the construction of the graph G used to prove Theorem 3.

Our lower-bound graph G contains the temporal graph G' as stated in Lemma 1, a temporal path P of $n + 1$ vertices $s_1, \dots, s_{n+1}$ such that each edge (s_i, s_{i+1}) has time-label $n + i$, and an edge (t_i, s_{n-i+1}) with time-label i for each $i = 1, \dots, n$. Finally, let $s = \sigma$ and $t = s_{n+1}$. See Fig. 3.

Lemma 2. Let $1 \leq i < n$ and define $\tau = 2n - i$. The shortest temporal path from s to t in $G^{-\tau}$ is unique and is the concatenation of π_i , the edge (t_i, s_{n-i+1}) , and the subpath from s_{n-i+1} to t in P .

Proof. Notice that any vertex s_{n-j+1} with $j > i$ cannot reach t following a blackout at time $\tau = 2n - i$. Therefore, the shortest temporal path from s to t in $G^{-\tau}$ (this path exists since t is still reachable from s in $G^{-\tau}$) must go through an edge (t_j, s_{n-j+1}) with $j \leq i$ and then it must follow the subpath from s_{n-j+1} to t in P .

Observe that the shortest temporal path from s to t_j in G is π_j , and that the subpath from s_{n-j+1} to t in P has length j . For $j \leq i$, the concatenation of the above paths via the augmenting edge (t_j, s_{n-j+1}) is a temporal path π_j^* from s to t in $G^{-\tau}$ of length $|\pi_j| + j + 1$ and, in particular, it is the unique shortest temporal path from s to t in $G^{-\tau}$ among those passing through (t_j, s_{n-j+1}) . Hence, to prove the claim we only need to show that the choice of $j \leq i$ that minimizes $|\pi_j^*|$ is $j = i$. This is true since, for any $j < i$, we can invoke Lemma 1 to obtain $|\pi_i^*| = |\pi_i| + i + 1 < |\pi_j| - (i - j) + i + 1 = |\pi_j| + j + 1 = |\pi_j^*|$. $\square$

We are now ready to state the following:

Theorem 3. For any n , there exists a temporal graph G of $\Theta(n)$ vertices, such that any blackout-tolerant single-pair temporal preserver of G (w.r.t. a suitable pair of vertices) has size $\Omega(n^2)$.

Proof. Consider the temporal graph G of this section and observe that any single-pair blackout-tolerant temporal preserver H of G w.r.t. s and t must contain all paths $\pi_1, \dots, \pi_{n-1}$. Indeed, Lemma 2 ensures that, for any $i = 1, \dots, n - 1$, the unique shortest temporal path from s to t following a blackout at time $\tau = 2n - i$ contains π_i as a subpath. By Lemma 1, the union of the edges in all the temporal shortest paths $\pi_1, \dots, \pi_n$ is exactly G' , which has size $\Theta(n^2)$. Hence, since $|\pi_n| \leq n$, we can conclude that H must have size at least $\sum_{i=1}^{n-1} |\pi_i| = \Omega(n^2)$. The claim follows by noticing that G has $\Theta(n) + n + 1$ vertices (see Lemma 1). $\square$

3.3. Upper bounds on BT single-pair temporal preservers on cliques

Given a temporal clique G and two vertices s, t , with $s \neq t$, we have that the shortest temporal path between s and t coincides with the edge (s, t) for any blackout at time $\tau \neq \lambda(u, v)$. We can therefore build a blackout-tolerant single-pair temporal preserver of G w.r.t. s and t by selecting the edge (s, t) along with a shortest temporal path in $G^{-\lambda(s, t)}$.

This is tight as it can be seen by considering the temporal clique in which there is a Hamiltonian path from s to t consisting of edges with time-label 2, and all other edges have time-label 1.

The above discussion is summarized by the following observation:

Observation 1. For any temporal clique $G = (V, E, \lambda)$ of n vertices and a pair of vertices $s, t \in V$, there exists a blackout-tolerant single-pair temporal preserver of G with source s and target t of size $O(n)$. This is tight.

4. Lower bounds on BT single-source temporal spanners

In this section we first show that there is a temporal graph G of $\Theta(n)$ vertices such that any blackout-tolerant single-source temporal spanner of G has size $\Omega(n^2)$. Then, we discuss how such a temporal graph G can be transformed into a clique for which the same lower bound holds.

Our temporal graph G is similar to the one used for the lower bound of the blackout-tolerant single-pair temporal preserver discussed in Section 3.2. We build G as the union of (i) a temporal path P of $n + 1$ vertices $s_1, \dots, s_{n+1}$, such that for $i = 1, \dots, n$, the time-label of the edge (s_i, s_{i+1}) is $n - i + 1$, (ii) a time-shifted copy of the temporal graph G' of Lemma 1 in which each time-label $\lambda(e)$ is replaced by $\lambda(e) + n$, and (iii) an augmenting edge (s_i, σ) with time-label $n + i$ for each $i = 1, \dots, n$. Finally, we choose $s = s_{n+1}$ (see Fig. 4).




Fig. 4. Sketch of the construction of the graph G used to prove Theorem 4.




Fig. 5. Sketch of the construction of the graph $\tilde{G}$ used to prove Theorem 5. Edges with time-label 1 or $M+1$ are dashed. All edges in (the time-shifted copy of) G have a time-label between 3 and $2n+2$, except for the additional edges with time-label 1.

Lemma 3. Let $2 \leq i \leq n$ and define $\tau = n - i + 2$. Vertex s can reach vertex t_i in $G^{-\tau}$. Moreover, all temporal paths from t to t_i in $G^{-\tau}$ contain π_i as a subpath.

Proof. We first suppose that a temporal path π from t to t_i in $G^{-\tau}$ exists and we show that it must contain π_i as a subpath. Let (s_j, σ) be the first augmenting edge traversed by π (this edge exists since the augmenting edges form an s - t_i -cut in G). Since the time-label of (s_j, σ) is $n+j$ and π must have arrival time at most $n+i$ (see again Lemma 1), we must have $j \leq i$. At the same time, we notice that $\tau = n - i + 2$ is the time-label of the edge (s_{i-1}, s_i) . Thus, after the blackout at time τ , only the vertices s_j with $j \geq i$ can be reached by s without traversing an augmenting edge, hence we must have $j \geq i$.

We can then conclude that $j = i$ and that all the edges in $\pi[\sigma : t_i]$ have time-label $n+i$. However, the set of edges with time-label $n+i$ in G (and in $G^{-\tau}$) induces exactly the simple path π_i from σ to t_i . Therefore we have $\pi[\sigma : t_i] = \pi_i$.

To see that s can reach t_i in $G^{-\tau}$, we notice that s can reach s_i in $G^{-\tau}$ with a temporal path π'_i having arrival time $n - i + 1$, and that π'_i and (s_i, σ) (resp. (s_i, σ) and π_i) are compatible. $\square$

We can now prove the following theorem.

Theorem 4. For any n , there exists a temporal graph G of $\Theta(n)$ vertices, such that any blackout-tolerant single-source temporal spanner of G (w.r.t. a suitable source vertex) has size $\Theta(n^2)$.

Proof. Consider the temporal graph G of this section and notice that it has $\Theta(n)$ vertices. By Lemma 3, any blackout-tolerant single-source temporal spanner H of G w.r.t. s must contain all temporal paths π_i for $i = 2, \dots, n$. By Lemma 1 these paths are edge disjoint and their overall number of edges is $\sum_{i=2}^n |\pi_i| = \Omega(n^2)$. Hence the size of H is $\Omega(n^2)$. $\square$

We now show how to strengthen Theorem 4 by transforming G into a temporal clique $\tilde{G}$ for which the same asymptotic lower bound holds. The temporal clique $\tilde{G}$ is obtained by starting from a time-shifted copy of G in which each time-label $\lambda(e)$ is replaced with $\lambda(e) + 2$ and augmenting such graph as follows (see Fig. 5):

- For every vertex t_i , create a new vertex z_i and add the edge (t_i, z_i) with time-label M , with $M = 2n + 3$;
- Add a new vertex $\tilde{s}$ and an edge $(\tilde{s}, z_i)$ with time-label $\tau_i = n - i + 2$ for each $i = 1, \dots, n$. Add the edge $(\tilde{s}, s)$ with time-label 2;
- Add all the remaining edges incident to $\tilde{s}$, and set their time-labels to $M + 1$.
- Finally, add all remaining edges (between any pair of vertices in $\tilde{G}$) and set their time-labels to 1.

Theorem 5. For any n , there exists a temporal clique $\tilde{G}$ of $\Theta(n)$ vertices, such that any blackout-tolerant single-source temporal spanner of $\tilde{G}$ (w.r.t. a suitable source vertex) has size $\Theta(n^2)$.

Proof. Consider the temporal graph $\tilde{G}$ of this section. First of all, notice that all edges incident to $\tilde{s}$ have time-label at least 2 and hence we can ignore all the edges with time-label 1, since they do not belong to any temporal path from $\tilde{s}$.




Fig. 6. Sketch of the construction of the graph G used to prove Theorem 6.

Consider a vertex z_i with $i \geq 2$ and a blackout at time $\tau_i = n - i + 2$. We claim that any temporal path π from $\tilde{s}$ to z_i must use $(\tilde{s}, s)$ as its first edge and must enter in z_i with the edge (t_i, z_i) . Indeed, the only edge incident to z_i in $\tilde{G}^{-\tau_i}$ with a time-label different from 1 is (t_i, z_i) . As a consequence, no edge in π can have time-label $M + 1$. Moreover π cannot contain any vertex z_j with $j \neq i$, since then π would need to traverse (z_j, t_j) which has time-label M , but $\tilde{G}^{-\tau_i}$ contains no path from z_j to t_i that only uses edges with time-label M .

Hence, we know that $\pi[s : t_i]$ is a path from s to t_i in $G^{-\tau_i}$ and, by Lemma 3, any such path includes π_i from the time-shifted copy of G' as a subpath (see Lemma 1). To show that at least one such π exists, notice that there is a temporal path π' from s to t_i in $G^{-\tau_i}$ and hence $(\tilde{s}, s) \circ \pi' \circ (t_i, z_i)$ is a temporal path in $\tilde{G}^{-\tau_i}$.

To conclude the proof we notice that $\tilde{G}$ has $\Theta(n)$ vertices and that, by Lemma 1, $\sum_{i=2}^n |\pi_i| = \Omega(n^2)$. $\square$

5. More than one blackout

In this section we consider b -blackout-tolerant temporal spanners, i.e., temporal spanners of temporal graphs that can withstand multiple blackouts. We start by formalizing this notion.

For a given temporal graph G with lifetime L and a subset F of $\{1, \dots, L\}$, we denote by G^{-F} the temporal subgraph of G that contains all vertices and exactly the edges e with $\lambda(e) \notin F$. A b -blackout-tolerant (b -BT) temporal spanner (resp. α -spanner) of a temporal graph G with lifetime L is a temporal subgraph H of G such that H^{-F} is a temporal spanner (resp. α -spanner) of G^{-F} for all $F \subseteq \{1, \dots, L\}$ with $|F| \leq b$.

We now provide a lower bound on the size of any 2-BT single-pair temporal spanner of a temporal graph. Our lower-bound graph G consists of a temporal path P' of $n+1$ vertices $s'_1, \dots, s'_{n+1}$, where $\lambda(s'_i, s'_{i+1}) = n - i + 1$, a time-shifted copy of the graph G' of Lemma 1 in which each time-label $\lambda(e)$ is replaced with $\lambda(e) + n$, a temporal path P of $n+1$ vertices $s_1, \dots, s_{n+1}$ where $\lambda(s_i, s_{i+1}) = 2n + i$, and the following additional edges:

- An edge (s'_i, σ) with time-label $n + i$ for each $i = 1, \dots, n$;
- An edge (t_i, s_{n-i+1}) with time label $n + i$ for each $i = 1, \dots, n$.

Finally, we choose $s = s'_{n+1}$ and $t = s_{n+1}$. See Fig. 6.

Theorem 6. For any n , there exists a temporal graph G of $\Theta(n)$ vertices, such that any 2-blackout-tolerant single-pair temporal spanner of G (w.r.t. a suitable pair of vertices) has size $\Theta(n^2)$.

Proof. Consider the temporal graph G of this section and notice that it has $\Theta(n)$ vertices. Pick any $i = 2, \dots, n-1$ and let $\tau = n - i + 2$ and $\psi = 3n - i$.

Notice that any path from s to t in $G^{-\tau}$ must traverse an edge (s'_j, σ) with $j \geq i$ and, by definition of P' , it has time-label $n + j \geq n + i$.

Symmetrically, by definition of P , any path from s to t in $G^{-\psi}$ must traverse an edge (t_j, s_{n-j+1}) with $j \leq i$, which has time-label $n + j \leq n + i$.

This implies that, if there is a path π from s to t in $G^{-\{\tau, \psi\}}$, π must contain a subpath between σ and s_i in the (time-shifted) graph G' that uses only edges with time-label $n + i$. By Lemma 1, the only such subpath is π_i .

To show that π exists, we notice that we can build a temporal path in $G^{-\{\tau, \psi\}}$ as a concatenation of (i) the unique simple path from s to s'_i in P' , (ii) the edge (s'_i, σ) , (iii) the path π_i in (the time-shifted version of) G' , (iv) the edge (t_i, s_{n-i+1}) , and (v) the subpath path from s_{n-i+1} to t in P .

Using the fact that the paths π_i are edge-disjoint (Lemma 1), we can conclude that any 2-blackout-tolerant single-pair temporal spanner of G with source s and target t has size at least $\sum_{i=2}^{n-1} |\pi_i| = \Omega(n^2)$. $\square$

We now show how to extend some of our upper and lower bound results to a larger number of blackouts by using the following observation: Let G be a temporal graph with lifetime L and $\tau \in \{1, \dots, L\}$. Any b -BT spanner (resp. α -spanner) of G with $b \geq 1$ is a $(b-1)$ -BT spanner (resp. α -spanner) of $G^{-\tau}$.

As a consequence, any lower bound on the size of a $(b - 1)$ -BT spanner (resp. α -spanner) on general (not necessarily complete) temporal graphs also applies to b -BT spanners (resp. α -spanner) on temporal cliques, since we can complete the lower-bound graph with edges having a new time-label τ .

Applying the above argument to our lower bounds of Theorem 3 and Theorem 6 on the size of all-pairs temporal-spanners on general graph, we have the following⁶:

Observation 2. *The following blackout-tolerant temporal spanners admit a lower bound of $\Omega(n^2)$ on their size:*

- 2-BT single-pair temporal preservers of temporal cliques;
- 3-BT single-pair temporal spanners of temporal cliques.

We conclude this section by arguing that our upper bounds for 1-BT single-pair temporal spanners on general graphs, extend to 2 blackouts on temporal cliques. Indeed, when the input graph is a temporal clique, the only blackout that affects the (single edge) shortest path from s to t is the time-label τ of the edge (s, t) . Thus, for a given temporal clique G and pair of vertices s and t , the subgraph H obtained by adding edge (s, t) to a 1-BT single-pair temporal spanner of $G^{-\tau}$ with source s and target t , is a 2-BT single-pair temporal spanner of G with source s and target t .

Applying the above argument to our upper bounds of Theorem 1 and Theorem 2, we have the following:

Observation 3. *Any temporal clique of n vertices admits:*

- A 2-BT single-pair temporal spanner of size $O(n)$; and
- A 2-BT single-pair temporal $(1 + \varepsilon)$ -spanner of size $O(\frac{n \log^4 n}{\log(1 + \varepsilon)})$, for any $0 < \varepsilon < n$.

6. Edge fault-tolerant temporal spanners

In this section, we show several upper and lower bounds of (S, T) -temporal spanners in the case of the failure of a single edge. More formally, given a temporal graph G and an edge e , we denote by G^{-e} the temporal graph G without the edge e , i.e., G after the failure of the edge e .

Given two sets of vertices $S, T \subseteq V$, an (S, T) -temporal spanner H of G is 1-edge fault-tolerant (1-EFT) if for every edge e , H^{-e} is an (S, T) -temporal spanner of G^{-e} . We also extend the same definition to (S, T) -temporal α -spanners for $\alpha \geq 1$.

6.1. 1-EFT temporal spanners on general graphs

Given a temporal graph G and two vertices s and t , we show how to transform G into a temporal graph $\tilde{G}$ such that the set of edges selected in any blackout-tolerant single-pair temporal α -spanner of $\tilde{G}$ w.r.t. s and t also induces a 1-EFT single-pair temporal α -spanner of G w.r.t. s and t , for any $\alpha \geq 1$. This will allow us to use the upper bounds for blackout-tolerant spanners of the previous sections. In our transformation, we will assign multiple time-labels to some edges. We are allowed to do this since our upper bounds on blackout-tolerant spanners also work in the case of multiple time-labels.

For the sake of simplicity in the explanation, here we slightly abuse the notation and extend the definition of edge labels to fractional values, which can be easily converted to positive integer values by replacing each time-label with its rank in the sorted list of distinct time-labels.

Starting with $\tilde{G} = G$, we fix a shortest temporal path $\pi = e_1, \dots, e_k$ from s to t , where $e_1, \dots, e_k$ is the ordered sequence of the edges traversed from s to t . For each time-label τ that appears in π , let e_i be the first edge and e_j be the last edge, traversed in π , having time-label τ . Since π is a temporal path, all the edges between e_i and e_j have time-label τ . Thus in $\tilde{G}$, for each $0 \leq h \leq j - i$, we replace the time-label of the h -th edge e_{i+h} between e_i and e_j with $\tau + h\varepsilon$, where $\varepsilon < \frac{1}{n}$. Finally, for every edge e with time-label τ in $\tilde{G}$ that is not traversed in π , we assign to e the set of time labels $\{\tau, \tau + \varepsilon, \dots, \tau + (j - i + 1)\varepsilon\}$. Since in $\tilde{G}$ the edges have multiple time-labels, we redefine a temporal path π from a vertex u to a vertex v in $\tilde{G}$ as a sequence of triples $\{(v_{i-1}, v_i, t_i)\}_{i=1}^k$, where $v_0 = u$, $v_k = v$, t_i is a time-label of the edge (v_i, v_{i+1}) , and $t_i \leq t_{i+1}$.

Lemma 4. *For each edge e in π , if s can reach t in G^{-e} , there exists a time-label λ such that $d_{\tilde{G}}^{-e}(s, t) = d_{\tilde{G}}^{-\lambda}(s, t)$ and the edges in any temporal path from s to t in $\tilde{G}^{-\lambda}$ induce a temporal path from s to t in G^{-e} .*

Proof. Let $e = (u, v)$ be an edge of π such that s can reach t in G^{-e} . We choose λ as the unique time label of edge e in $\tilde{G}$, and we notice that $\lambda = \tau + h\varepsilon$ for a suitable value of h .

⁶ This also provides an alternative way to show a lower bound of $\Omega(n^2)$ on the size of 1-BT temporal spanners on temporal cliques from the corresponding lower bound of [2] on temporal graphs.




Fig. 7. Qualitative representation of the paths defined in the proof of Lemma 4. In (a), and (b) the edge (u, v) is traversed in opposite directions by π and π' , and the highlighted path is a temporal path that uses less backward edges than π' . In (c), the highlighted path is π_e and is obtained by suitably combining π with the path π^* , which is guaranteed not to use any backward edge. Color figures can be found in the web version of this article.

To prove the claim we will show that (i) there is a temporal path from s to t in $\bar{G}^{-\lambda}$ of length $d_G^{-e}(s, t)$, and (ii) the edges of a temporal path from s to t in $\bar{G}^{-\lambda}$ induce a temporal path from s to t in G^{-e} of the same length.

We start by showing (i). Given a path π' from s to t in G , we say that π' uses a backward edge if there is some edge (u, v) that belongs to both π' and π , u is traversed before v in π' , and v is traversed before u in π . We now argue that there exists a shortest temporal path π_e from s to t in G^{-e} such that $\pi_e = \pi[s : x] \circ \tilde{\pi}[x : y] \circ \pi[x : y]$ for suitable vertices x, y , where $\tilde{\pi}[x : y]$ is a temporal path that does not traverse any edge of π (in either direction), and x (resp. y) might coincide with s (resp. t).

Let π^* be any shortest temporal path from s to t in G^{-e} . We can always choose a path π^* that does not use any backward edge. To see this suppose towards a contradiction that all shortest temporal paths from s to t in G^{-e} use a backward edge, and let π' be a shortest temporal path from u to v minimizing the number of backward edges, and let (u, v) be a backward edge of π' .

We can decompose π as $\pi[s : v] \circ (\pi[v, u] \circ \pi[u : t])$ and π' as $\pi'[s : u] \circ (u, v) \circ \pi'[v : t]$. There are two cases:

- Edge (v, u) is traversed before e in π (see Fig. 7 (a)). Since $\pi[s : v]$ is a shortest temporal path from s to v among those having arrival time at most $\lambda(v, u) = \lambda(u, v)$, we have that $|\pi[s : v]| \leq |\pi'[s : u] \circ (u, v)|$. Moreover, the arrival time of $\pi[s : v]$ is at most $\lambda(u, v)$, i.e., it is at most the departure time of $\pi'[v : t]$. Thus, $\pi[s : v] \circ \pi'[v : t]$ is a temporal path from s to t in G^{-e} such that $|\pi[s : v] \circ \pi'[v : t]| \leq |\pi'|$ and uses less backward edges than π' ;
- Edge (v, u) is traversed after e in π (see Fig. 7 (b)). Since π is a shortest temporal path from u to t among those having departure time at least $\lambda(v, u)$, we have that $|\pi[u : t]| \leq |(u, v) \circ \pi'[v : t]|$. Moreover, the arrival time of $\pi'[s : u]$ is at most $\lambda(u, v)$, i.e., it is at most the departure time of $\pi[u : t]$. Thus, $\pi'[s : u] \circ \pi[u : t]$ is a temporal path from s to t in G^{-e} such that $|\pi'[s : u] \circ \pi[u : t]| \leq |\pi'|$ and uses less backward edges than π' .

Thus from now on, we assume that the path π^* does not contain any backward edge and we show that a suitable modification of π^* yields π_e . Let (x', x) be the last edge of π^* that is also in π and precedes e (in π). If no such edge exists, let $x = s$. Similarly, let (y, y') be the first edge of π^* that is also in π and follows e (in π). If no such edge exists, let $y = t$.

The path $\pi_e = \pi[s : x] \circ \tilde{\pi}[x : y] \circ \pi[y : t]$ is obtained by choosing $\tilde{\pi}[x : y] = \pi^*[x : y]$. See Fig. 7 (c).

We now argue that the edges of π_e induce a corresponding temporal path in $\bar{G}^{-\lambda}$ by specifying, for each edge f in π_e the time label at which edge f is traversed in $\bar{G}^{-\lambda}$. If f is in π , then f is used at time of its sole time-label in $\bar{G}^{-\lambda}$. Otherwise, f is in $\pi_e[x : y]$ and we handle it as follows: if $\lambda(f) \neq \lambda(e)$, then f is traversed at time $\lambda(f)$ in $\bar{G}$. Otherwise $\lambda(f) = \lambda(e)$, and f is traversed at time $\lambda(e) + (h + 1)\varepsilon$ in $\bar{G}$. Observe that the time-labels of any pair of consecutive edges are non-decreasing, hence the path π_e with the chosen time-labels is a temporal path in $\bar{G}^{-\lambda}$.

To prove (ii), consider any edge f in $\bar{G}^{-\lambda}$ with a corresponding time label τ and notice that f is also an edge in G^{-e} with time-label $\lfloor \tau \rfloor$. Hence, the edge of any temporal path in $\bar{G}^{-\lambda}$ induce a temporal path in G^{-e} (of the same length). $\square$

Combining Lemma 4 with Theorem 1 and Theorem 2, we immediately obtain the following results:

Theorem 7. For any temporal graph G there exists a single-pair 1-edge fault-tolerant temporal spanner of G of size $O(n)$.

Theorem 8. For any temporal graph G and $\varepsilon > 0$ there exists a single-pair 1-edge fault-tolerant temporal $(1 + \varepsilon)$ -spanner of G of size $O\left(\frac{n \log^4 n}{\log(1 + \varepsilon)}\right)$.

We conclude this section by observing that, our lower-bound of $\Omega(n^2)$ on the size of 1-BT single-pair temporal preserver extends to 1-EFT single-pair temporal preserver. In particular, consider the lower-bound graph G defined in Section 3.2, for which any 1-BT single-pair temporal preserver of G (w.r.t. a suitable pair of vertices) has size of $\Omega(n^2)$. For G , we showed

by Lemma 2 that for any time $\tau = 2n - i$ with $1 \leq i < n$, the shortest temporal path from s to t in $G^{-\tau}$ is unique and contains the temporal path π_i of the temporal graph G' of Lemma 1. However, we can notice that in G there is exactly one edge e with time-label τ . Thus, a blackout at time τ coincides with the failure of the edge e , i.e. $G^{-\tau} = G^{-e}$. This allows us to conclude that:

Observation 4. Any 1-EFT single-pair temporal preserver of G (w.r.t. a suitable pair of vertices) has size $\Omega(n^2)$.

Moreover, a similar argument can be used to extend the lower bound of $\Omega(n^2)$ on the size of 1-BT single-source temporal spanner for general graphs of Theorem 4 to 1-EFT single-source temporal spanners:

Observation 5. Any 1-EFT single-source temporal spanner of G (w.r.t. a suitable source vertex) has size $\Omega(n^2)$.

6.2. 1-EFT single-source temporal spanners on temporal cliques

In this section, we show that any temporal clique has a 1-EFT single-source temporal spanner of size $O(n)$.

In order to build our spanner, we select a suitable set Π that contains a temporal path π_v in $G^{-(s,v)}$ for every $v \in V$ that is reachable from s in $G^{-(s,v)}$ via a temporal path. Notice that any such set Π immediately induces a single-source 1-EFT temporal spanner H_Π obtained as the union of all the edges that are incident to s or belong to at least one path in Π . Indeed, for each v , we have two edge-disjoint temporal paths in H_Π , i.e., the temporal path containing only the edge (s, v) and π_v (if π_v does not exist then all temporal paths from s to v use (s, v)).

In the rest of the section, we focus on computing a set Π for which the size of H_Π is $O(n)$.

We say that a temporal path $\pi \in \Pi$ that contains a vertex $v \neq s$ enters v with the edge (u, v) if (u, v) is the unique edge of $\pi[s, v]$ that is incident to v . We define $E_\Pi(v)$ as the set of edges (u, v) with $u \neq s$ such that there is a path $\pi \in \Pi$ that visits v and enters v with the edge (u, v) . Our strategy to find Π is as follows: we start from a tentative set Π that contains an arbitrary path π_v from s to v in $G^{-(s,v)}$, for each vertex v that is reachable from s . Then, we iteratively update Π by rerouting some paths. In particular, as long as there is some vertex v with $|E_\Pi(v)| \geq 3$, we modify Π to guarantee that $|E_\Pi(v)| \leq 2$ and that the values $|E_\Pi(u)|$ for $u \neq v$ do not increase. We stop as soon as $|E_\Pi(v)| \leq 2$ for all vertices $v \in V$, as this immediately implies that H_Π has size at most $3(n-1)$. Notice that this procedure must stop after at most $n-1$ iterations.

We now describe how Π is updated. Let v be a vertex with $|E_\Pi(v)| \geq 3$. Our update procedure works in two phases. The first phase modifies Π to ensure that it satisfies a useful structural property (see Lemma 5). This phase cannot increase any value $|E_\Pi(u)|$ and in particular, it might already decrease $|E_\Pi(v)|$ below 3, in which case we are already done. Otherwise, we proceed with the second phase which exploits the above structural property and modifies Π further to reduce the value $|E_\Pi(v)|$ to 2 (without increasing any of the other values $|E_\Pi(u)|$, with $u \neq v$).

Phase 1. We start by ordering the edges in $E_\Pi(v)$ in non-decreasing order of their time label, breaking ties arbitrarily. For $e \in E_\Pi(v)$, let $r(e)$ be the rank of edge e in the chosen order. Then, we iteratively look for two paths $\pi_x, \pi_y \in \Pi$ that satisfy the following conditions: i) π_x and π_y both contain v , ii) the edge e_x of $\pi_x[s, v]$ that enters v differs from the edge e_y of $\pi_y[s, v]$ that enters v , iii) the edge leaving s of π_x coincides with the edge leaving s of π_y . Without loss of generality assume that $r(e_x) < r(e_y)$. If $\pi_x[s, v]$ traverses y , we update π_y in Π by replacing it with the path $\pi_x[s, y]$ (notice that $\pi_x[s, y]$ cannot contain (s, y)). Otherwise we replace π_y with $\pi_x[s, v] \circ \pi_y[v, y]$ (notice that $\pi_x[s, v]$ cannot contain (s, y) and that $\pi_x[s, v]$ is compatible with $\pi_y[v, y]$ since the arrival time of $\pi_x[s, v]$ is at most $\lambda(e_y)$). We repeat the above procedure (with the same order of the edges in $E_\Pi(v)$) as long as there are paths π_x, π_y that satisfy properties i)–iii). To prove that this procedure terminates after at most a polynomial number of updates, consider the set Π_v of all paths $\pi_u \in \Pi$ that contain v , let e_u be the edge entering v in $\pi_u[s, v]$, and observe that each update to Π decreases the quantity $\sum_{\pi_u \in \Pi_v} r(e_u)$ by at least 1 (and, initially, $\sum_{\pi_u \in \Pi_v} r(e_u) < n^2$).

At the end of phase 1, the set Π satisfies the following structural property:

Lemma 5. For any two paths $\pi, \pi' \in \Pi$ that traverse v such that the last edge of $\pi[s, x]$ differs from the last edge of $\pi'[s, x]$, we have that the first edge of $\pi[s, x]$ differs from the first edge of $\pi'[s, x]$.

Phase 2. If we still have $|E_\Pi(v)| \geq 3$ after phase 1, we modify Π further to ensure that $|E_\Pi(v)| = 2$ and that the quantity $|E_\Pi(u)|$ of all other vertices u does not increase. We modify Π as follows: we select e_1 as the edge with the minimum time-label in $E_\Pi(v)$ and e_2 as the edge with the minimum time-label in $E_\Pi(v) \setminus \{e_1\}$ (ties are broken arbitrarily). For $i \in \{1, 2\}$, let π_i be a path in Π such that $\pi_i[s, v]$ contains e_i , and let (s, u_i) be the first edge of π_i . From Lemma 5, we have $u_1 \neq u_2$. For any edge $e \in E_\Pi(v) \setminus \{e_1, e_2\}$, and for every vertex u such that $\pi_u[s, v]$ contains e , pick the smallest $i \in \{1, 2\}$ such that $u_i \neq u$ (such an i always exists), and update Π by replacing π_u with $\pi_i[s, v] \circ \pi_u[v, u]$. Notice that $\pi_i[s, v] \circ \pi_u[v, u]$ is a temporal path from s to u in $G^{-(s,u)}$ since $\lambda(e) \geq \lambda(e_2) \geq \lambda(e_1)$.

The following theorem summarizes the result of this section:

Theorem 9. *Given a complete temporal graph G and a source s , it is possible to compute, in polynomial time, a 1-EFT single-source temporal spanner w.r.t. G and s having size at most $3(n-1)$.*

6.3. 1-EFT temporal preservers on temporal cliques

In this section we consider the problem of designing temporal preservers of temporal cliques.

In particular, it is easy to see that if one is interested in preserving distances from a single source node s to a single target node t , then the same arguments of Section 3.3 apply. Indeed, it suffices to build H as the union of the edge (s, t) with the edges in a shortest temporal path from s to t in $G^{-(s,t)}$.

Observation 6. *For any temporal clique $G = (V, E, \lambda)$ of n vertices and a pair of vertices $s, t \in V$, there exists a 1-EFT single-pair temporal preserver of G with source s and target t of size $O(n)$. This is asymptotically tight.*

As consequence of the above discussion we focus on 1-EFT single-source and all-pair preservers, and we show that there are graphs for which any such preserver must be dense. We do so by constructing a temporal clique G such that any 1-EFT single-source temporal preserver of G w.r.t. a suitable source vertex s has size $\Omega(n^2)$, which immediately implies the same result also for the all-pairs case.

The temporal graph G is obtained by starting from the temporal graph G' of Lemma 1 where each time-label λ is replaced with $\lambda + 1$, and augmenting the resulting graph as follows:

- For each edge (u, v) that is not in G' , add (u, v) to G with time-label 1;
- Add new vertex s , and all the edges (s, v) from s to the vertices v in G' . We set the time-label of (s, σ) to 2, and the time-labels of all other edges (s, v) to $M + 2$, where M is the maximum time-label in G' .

Theorem 10. *For any n , there exists a graph with $\Theta(n)$ vertices such that any 1-EFT single-source temporal preserver of G w.r.t. s has size $\Omega(n^2)$.*

Proof. Consider the temporal graph G of this section and notice that it has $\Theta(n)$ vertices.

Since all the edges incident to s in G have time-label greater than 1, all temporal paths from s contain no edges with time-label 1. Consider now a target vertex t_i in G' . Since G is a clique, the shortest temporal path from s to t_i in G is the edge (s, t_i) . If the edge (s, t_i) fails, the unique temporal path from s to t_i in $G^{-(s,t_i)}$ is $(s, \sigma) \circ \pi_i$, where π_i is the unique shortest temporal path from s to t_i in (the time-shifted version of) G' .

Thus, any 1-EFT single-source temporal preserver H of G w.r.t. s must contain all temporal paths π_i , with $1 \leq i \leq n$. By Lemma 1, $|H| \geq \sum_{i=1}^n |\pi_i| = \Omega(n^2)$. $\square$

6.4. More than one edge failure

In this section, we consider k -edge fault-tolerant (f -EFT) temporal spanners, i.e., temporal spanners in the case of the failure of k edges. Given a temporal graph $G = (V, E, \lambda)$ and a set of edges $F = \{e_1, \dots, e_k\}$, we denote with G^{-F} the temporal graph G without the set of edges F , i.e. G after the failure of all the edges in F . A k -EFT temporal spanner of a temporal graph G is a subgraph H of G such that H^{-F} is a temporal spanner of G^{-F} for all $F \subseteq E$, with $|F| \leq k$.

We now observe how the lower-bound of $\Omega(n^2)$ on 2-BT single-pair temporal spanner on general graphs (Section 5) extends to 2-EFT single-pair temporal spanner on general graphs. Indeed, consider the lower-bound graph G defined in Section 5, for which any 2-BT single-pair temporal spanner of G (w.r.t. a suitable pair of vertices) has size of $\Omega(n^2)$.

For G , we showed in the proof Theorem 6 that, for any $1 < i < n$, we can define two time-labels τ_i and ϕ_i such that, the shortest temporal path from s to t in $G^{-\tau_i, \phi_i}$ is unique and contains the temporal path π_i of the temporal graph G' of Lemma 1. However, we can notice that in G there are exactly two edges e and f with time label τ_i and ϕ_i , respectively. Thus, two blackouts at time τ_i and ϕ_i coincide with the failure of the edge e and f , i.e. $G^{-\tau_i, \phi_i} = G^{-\{e, f\}}$. This allows us to conclude that:

Observation 7. *Any 2-EFT single-pair temporal spanner of G (w.r.t. a suitable pair of vertices) has size of $\Omega(n^2)$.*

In the last part of this section, we show how to extend some of our lower and upper bound results in the case of a single edge failure to the case of multiple-edge failures.

Let G be a temporal graph. Notice that any f -EFT temporal spanner of G , with $f \geq 1$ is an $(f-1)$ -EFT temporal spanner of G^{-e} , for any $e \in E$.

We can use this observation to provide a reduction scheme that allows us to show that, any lower-bound on the size of $(f-1)$ -EFT (s, T) -temporal (α) -spanner on general temporal graphs also extends to f -EFT (s, T) -temporal (α) -spanner on temporal cliques.

Indeed, consider any lower-bound graph G for (s, T) -temporal (α) -spanners, where s is a suitable source vertex of G and $T \subseteq V \setminus \{s\}$. Let M be the maximum time-label of an edge in G , we complete G into a clique in the following way:

1. For any edge e in G , we replace the time-label $\lambda(e)$ by $\lambda(e) + 1$, i.e., we shift all the time-labels in G by 1;
2. We add all the missing edges incident to s and set their time-label to $M + 2$;
3. We add all the remaining missing edges in G and we set their time-labels to 1.

Lemma 6. Let $\tilde{H}$ be an f -EFT (s, T) -temporal (α) -spanner of $\tilde{G}$, and let H be the subgraph of G that contains an edge (u, v) if and only if (u, v) is in both G and $\tilde{H}$. H is a $(f - 1)$ -EFT (s, T) -temporal (α) -spanner of G .

Proof. It suffices to show the claim for a generic value of $\alpha \geq 1$, since a temporal spanner is a temporal n -spanners and a temporal preserver is a temporal 1-spanner.

Observe that no temporal path from s uses any edge with time-label 1 in $\tilde{G}$.

Let $t \in T$, let F be a set of at most $(f - 1)$ edges in G such that t is reachable from s in G^{-F} , and define π^* as a shortest temporal path from s to t in G^{-F} . We show that H^{-F} contains a temporal path π from s to t such that $|\pi| \leq \alpha \cdot |\pi^*|$.

If (s, t) is an edge in G , then we choose π as the shortest temporal path from s to t in $\tilde{H}^{-F}$. Notice that π exists since the (time-shifted version of) π^* belongs to $\tilde{G}^{-F}$. Moreover, π cannot use any edge with time label $M + 2$ in $\tilde{G}$, and hence π is also path in H^{-F} . Finally, $|\pi| = d_{\tilde{H}}^{-F}(s, t) \leq \alpha \cdot d_{\tilde{G}}^{-F}(s, t) \leq \alpha \cdot d_G^{-F}(s, t) = \alpha \cdot |\pi^*|$.

If (s, t) is not an edge of G , we chose π as the shortest temporal path from s to t in $\tilde{H}^{-F \cup \{(s, t)\}}$, which exists since π^* belongs to $\tilde{G}^{-F}$ and $\tilde{H}$ is a f -EFT (s, T) -temporal (α) -spanner of $\tilde{G}$. By using the same arguments we can show that π cannot use edges of time-label $M + 2$, and hence again is a path in H^{-F} such that $|\pi| = d_{\tilde{H}}^{-F \cup \{(s, t)\}}(s, t) \leq \alpha \cdot d_{\tilde{G}}^{-F \cup \{(s, t)\}}(s, t) \leq \alpha \cdot d_G^{-F}(s, t) = \alpha \cdot |\pi^*|$. $\square$

As a consequence any lower-bound on the size of $(f - 1)$ -EFT (s, T) -temporal (α) -spanner of G is also a lower-bound on f -EFT (s, T) -temporal (α) -spanner of $\tilde{G}$. Then, combining the above lemma to our lower bounds of Observation 4, Observation 7, Observation 5, we have the following:

Observation 8. The following edge-fault tolerant temporal spanners admit a lower bound of $\Omega(n^2)$ on their size:

- 2-EFT single-pair temporal preservers on temporal cliques;
- 3-EFT single-pair temporal spanners on temporal cliques;
- 2-EFT single-source temporal spanners on temporal cliques.

Finally, we observe that our upper bound for 1-EFT single-pair temporal spanners and α -spanners on general graphs extends to 2-EFT single-pair temporal spanners and α -spanners on temporal cliques. Indeed, when the input graph is a temporal clique, the only edge failure that affects the shortest paths from s to t is the single edge (s, t) . Thus, for a given temporal clique G and a pair of vertices s and t , the subgraph H obtained by adding (s, t) to a 1-EFT single-pair temporal (α) -spanner of $G^{-\{(s, t)\}}$, with source s and target t , is a 2-EFT single pair temporal (α) -spanner of G w.r.t. s and t . By Theorem 7 and Theorem 8, we have that:

Observation 9. Any temporal clique of n vertices admits:

- A 2-EFT single-pair temporal spanner of size $O(n)$;
- A 2-EFT single-pair temporal $(1 + \varepsilon)$ -spanner of size $O\left(\frac{n \log^4 n}{\log(1 + \varepsilon)}\right)$, for any $\varepsilon > 0$.

CRedit authorship contribution statement

Davide Bilò: Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Gianlorenzo D'Angelo:** Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Luciano Gualà:** Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Stefano Leucci:** Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Mirko Rossi:** Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] A.R. Ahmed, G. Bodwin, F.D. Sahneh, K. Hamm, M.J.L. Jebelli, S.G. Kobourov, R. Spence, Graph spanners: a tutorial review, *Comput. Sci. Rev.* 37 (2020) 100253, <https://doi.org/10.1016/j.cosrev.2020.100253>.
- [2] K. Axiotis, D. Fotakis, On the size and the approximability of minimum temporally connected subgraphs, in: I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, D. Sangiorgi (Eds.), 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11–15, 2016, Rome, Italy, in: *LIPIcs*, vol. 55, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 149.
- [3] D. Bilò, G. D'Angelo, L. Gualà, S. Leucci, M. Rossi, Sparse temporal spanners with low stretch, in: S. Chechik, G. Navarro, E. Rotenberg, G. Herman (Eds.), 30th Annual European Symposium on Algorithms, ESA 2022, September 5–9, 2022, Berlin/Potsdam, Germany, in: *LIPIcs*, vol. 244, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 19.
- [4] D. Bilò, G. D'Angelo, L. Gualà, S. Leucci, M. Rossi, Sparse temporal spanners with low stretch, *CoRR*, arXiv:2206.11113 [abs], 2022, <https://doi.org/10.48550/arXiv.2206.11113>.
- [5] D. Bilò, L. Gualà, S. Leucci, G. Proietti, Fault-tolerant approximate shortest-path trees, *Algorithmica* 80 (12) (2018) 3437–3460, <https://doi.org/10.1007/s00453-017-0396-z>.
- [6] D. Bilò, L. Gualà, S. Leucci, G. Proietti, Multiple-edge-fault-tolerant approximate shortest-path trees, *Algorithmica* 84 (1) (2022) 37–59, <https://doi.org/10.1007/s00453-021-00879-8>.
- [7] G. Bodwin, F. Grandoni, M. Parter, V.V. Williams, Preserving distances in very faulty graphs, *CoRR*, arXiv:1703.10293 [abs], 2017, <http://arxiv.org/abs/1703.10293>.
- [8] B.M. Bui-Xuan, A. Ferreira, A. Jarry, Computing shortest, fastest, and foremost journeys in dynamic networks, *Int. J. Found. Comput. Sci.* 14 (2003) 267–285, <https://api.semanticscholar.org/CorpusID:2271039>.
- [9] M. Calamai, P. Crescenzi, A. Marino, On computing the diameter of (weighted) link streams, in: D. Coudert, E. Natale (Eds.), 19th International Symposium on Experimental Algorithms, SEA 2021, June 7–9, 2021, Nice, France, in: *LIPIcs*, vol. 190, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 11.
- [10] A. Casteigts, J.G. Peters, J. Schoeters, Temporal cliques admit sparse spanners, *J. Comput. Syst. Sci.* 121 (2021) 1–17, <https://doi.org/10.1016/j.jcss.2021.04.004>.
- [11] A. Casteigts, M. Raskin, M. Renken, V. Zamaraev, Sharp thresholds in random simple temporal graphs, in: 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7–10, 2022, IEEE, 2021, pp. 319–326.
- [12] E. Füchsle, H. Molter, R. Niedermeier, M. Renken, Delay-robust routes in temporal graphs, in: P. Berenbrink, B. Monmege (Eds.), 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15–18, 2022, Marseille, France (Virtual Conference), in: *LIPIcs*, vol. 219, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 30.
- [13] E. Füchsle, H. Molter, R. Niedermeier, M. Renken, Temporal connectivity: coping with foreseen and unforeseen delays, in: J. Aspnes, O. Michail (Eds.), 1st Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2022, March 28–30, 2022, Virtual Conference, in: *LIPIcs*, vol. 221, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 17.
- [14] M. Gupta, S. Khan, Multiple source dual fault tolerant BFS trees, in: I. Chatzigiannakis, P. Indyk, F. Kuhn, A. Muscholl (Eds.), 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10–14, 2017, Warsaw, Poland, in: *LIPIcs*, vol. 80, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 127.
- [15] P. Holme, Temporal networks, in: R. Alhaji, J.G. Rokne (Eds.), *Encyclopedia of Social Network Analysis and Mining*, 2nd edition, Springer, 2018.
- [16] S. Huang, A.W. Fu, R. Liu, Minimum spanning trees in temporal graphs, in: T.K. Sellis, S.B. Davidson, Z.G. Ives (Eds.), Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 – June 4, 2015, ACM, 2015, pp. 419–430.
- [17] D. Kempe, J.M. Kleinberg, A. Kumar, Connectivity and inference problems for temporal networks, *J. Comput. Syst. Sci.* 64 (4) (2002) 820–842, <https://doi.org/10.1006/jcss.2002.1829>.
- [18] G.B. Mertzios, O. Michail, P.G. Spirakis, Temporal network optimization subject to connectivity constraints, *Algorithmica* 81 (4) (2019) 1416–1449, <https://doi.org/10.1007/s00453-018-0478-6>.
- [19] O. Michail, An introduction to temporal graphs: an algorithmic perspective, *Internet Math.* 12 (4) (2016) 239–280, <https://doi.org/10.1080/15427951.2016.1177801>.
- [20] M. Parter, D. Peleg, Fault-tolerant approximate BFS structures, *ACM Trans. Algorithms* 14 (1) (2018) 10:1–10:15, <https://doi.org/10.1145/3022730>.
- [21] M. Parter, D. Peleg, Fault tolerant approximate BFS structures with additive stretch, *Algorithmica* 82 (12) (2020) 3458–3491, <https://doi.org/10.1007/s00453-020-00734-2>.
- [22] H. Wu, J. Cheng, S. Huang, Y. Ke, Y. Lu, Y. Xu, Path problems in temporal graphs, *Proc. VLDB Endow.* 7 (9) (2014) 721–732, <https://doi.org/10.14778/2732939.2732945>.