

Sharp Thresholds in Random Simple Temporal Graphs*

Arnaud Casteigts[†] Michael Raskin[‡] Malte Renken[§] Viktor Zamaraev[¶]

Abstract

A graph whose edges only appear at certain points in time is called a temporal graph (among other names). Such a graph is temporally connected if each ordered pair of vertices is connected by a path which traverses edges in chronological order (i.e., a temporal path). In this paper, we consider a simple model of random temporal graph, obtained from an Erdős–Rényi random graph $G \sim G_{n,p}$ by considering a random permutation π of the edges and interpreting the ranks in π as presence times.

We give a thorough study of the temporal connectivity of such graphs and derive implications for the existence of several kinds of sparse spanners. It turns out that temporal reachability in this model exhibits a surprisingly regular sequence of thresholds. In particular, we show that, at $p = \log n/n$, any fixed pair of vertices can a.a.s. reach each other; at $2 \log n/n$, at least one vertex (and in fact, any fixed vertex) can a.a.s. reach all others; and at $3 \log n/n$, all the vertices can a.a.s. reach each other, i.e., the graph is temporally connected. Furthermore, the graph admits a temporal spanner of size $2n + o(n)$ as soon as it becomes temporally connected, which is nearly optimal as $2n - 4$ is a lower bound. This result is quite significant because temporal graphs do not admit spanners of size $O(n)$ in general (Kempe, Kleinberg, Kumar, STOC 2000). In fact, they do not even always admit spanners of size $o(n^2)$ (Axiotis, Fotakis, ICALP 2016). Thus, our result implies that the obstructions found in these works, and more generally, any non-negligible obstruction, is statistically insignificant: *nearly optimal spanners always exist in random temporal graphs*.

All the above thresholds are sharp. Carrying the study of temporal spanners a step further, we show that pivotal spanners—i.e., spanners of size $2n - 2$ composed of two spanning trees glued at a single vertex (one descending in time, the other ascending subsequently)—exist a.a.s. at $4 \log n/n$, this threshold being also sharp. Finally, we show that optimal spanners (of size $2n - 4$) also exist a.a.s. at $p = 4 \log n/n$. Whether this value is a sharp threshold is open; we conjecture that it is.

For completeness, we compare the above results to existing results in related areas, including edge-ordered graphs, gossip theory, and population protocols, showing that our results can be interpreted in these settings as well, and that in some cases, they improve known results therein. Finally, we discuss an intriguing connection between our results and Janson’s celebrated results on percolation in weighted graphs.

*An extended abstract of this work appeared in the Proceedings of the 62nd Annual Symposium on Foundations of Computer Science (FOCS) [22]. The work of the first and the second authors was supported by the French ANR, project ANR-22-CE48-0001 (TEMPOGRAL). The work of the second author was also supported by the European Research Council (ERC) under the Horizon 2020 research and innovation program (grant agreement No 787367 (PaVeS)). The work of the third author was supported by the German Research Foundation (DFG), project MATE (NI 369/17).

[†]University of Geneva, Switzerland. Email: arnaud.casteigts@unige.ch. ORCID: 0000-0002-7819-7013. Part of this project was done while the author was at LaBRI, CNRS, University of Bordeaux, France.

[‡]LaBRI, CNRS, University of Bordeaux, France. Email: mraskin@u-bordeaux.fr. ORCID:0000-0002-6660-5673. Part of this project was done while the author was at Department of Informatics, Technical University of Munich, Germany.

[§]Technical University of Berlin, Germany. Email: m.renken@tu-berlin.de. ORCID: 0000-0002-1450-1901.

[¶]Department of Computer Science, University of Liverpool, United Kingdom. Email: Viktor.Zamaraev@liverpool.ac.uk. ORCID: 0000-0001-5755-4141.

1 Introduction

A temporal graph is a graph whose edges are present only at certain times. These graphs can be modeled in various ways, a classical option being as an edge-labeled graph $\mathcal{G} = (G, \lambda)$, where G is a standard graph and λ encodes the presence times of the edges. Temporal graphs have been extensively studied in the past two decades, motivated by the modeling of dynamic networks in areas like social network analysis, communication networks, epidemics, transportation, and biology. As theoretical objects, these graphs pose a number of fundamental questions. In particular, the redefinition of classical graph concepts in a temporal setting often leads to subtle, yet radical, differences. A canonical example is that of a *temporal path*, which is a path running over a non-decreasing sequence of timestamps (or increasing, in the case of strict temporal paths). Clearly, such paths are not symmetric, even with undirected edges, and they are also not transitive—the fact that vertex u can reach vertex v and v can reach w does not imply that u can reach w . Non-transitivity makes temporal paths quite different even from classical paths in directed graphs, with a strong impact on problems and algorithms. Among the early examples, it was shown in [41] that deciding whether a set of k vertex-disjoint temporal paths exist between two given vertices is NP-complete, whereas the analog problem in static graphs is polynomial-time solvable. Likewise, computing a maximum connected component based on temporal paths is essentially as hard as finding a maximum clique in static graphs [10]. Since then, many tractable problems have been shown to have intractable analogs in temporal graphs (see, e.g., [3, 20, 29, 31, 34, 45, 46]).

In this landscape, the status of spanning trees and other spanning structures remains unsettled. In static graphs, a spanning tree is a connected spanning subgraph without cycle, and therefore consists of $n - 1$ edges. It always exists when the graph is connected, and computing one is straightforward. In contrast, Kempe, Kleinberg, and Kumar [41] observe that the size of a minimum temporal spanner varies among (temporally connected) temporal graphs, and in some cases, only spanners of superlinear size exist. For example, hypercubes can be time-labeled in such a way that none of their $\Theta(n \log n)$ edges can be removed without breaking temporal connectivity (i.e., they are minimally connected). More surprisingly, Axiotis and Fotakis [7] construct an infinite family of minimally connected temporal graphs having $\Theta(n^2)$ edges. In other words, even *sparse* analogs of spanning trees do not exist unconditionally in temporal graphs. (Incidentally, finding a minimum spanner is computationally difficult, namely APX-hard [1].)

On the positive side, Casteigts, Peters, and Schoeters [21] showed that temporal graphs whose underlying graph is complete always admit spanners of size $O(n \log n)$. Whether a spanner of size $O(n)$ always exists in complete temporal graph is still open. However, a general result in gossip theory [16] implies that $2n - 4$ time labels are required to achieve temporal connectivity in any underlying graph (provided only that adjacent labels are distinct). Thus, spanners of size $2n - 4$ are optimal, and we call a spanner of size $2n + o(n)$ *nearly optimal*.

All these results were obtained for temporal graphs whose edges have exactly one presence time, i.e., *simple* temporal graphs. Due to their simplicity, these graphs prove to be good prototypes for studying connectivity problems. Moreover, the above results extend to general temporal graphs, either by containment (for negative results), or by the fact that the spanning property is preserved under the addition of time labels (for positive results). In the following, we restrict our attention to simple temporal graphs and refer the reader to [19] and [38] for background on general models.

Property	Shorthand	Sharp threshold	Reference
Point-to-point Reachability	$\forall u \forall v \text{ a.a.s. } u \rightsquigarrow v$	$\log n/n$	Theorem 5.1
First Temporal Source	$\text{a.a.s. } \exists u \forall v u \rightsquigarrow v$	$2 \log n/n$	Theorem 5.3
Temporal Source	$\forall u \text{ a.a.s. } \forall v u \rightsquigarrow v$	$2 \log n/n$	Theorem 5.2
Temporal Connectivity	$\text{a.a.s. } \forall u \forall v u \rightsquigarrow v$	$3 \log n/n$	Theorem 5.4

Table 1: Sharp thresholds for connectivity properties in random temporal graphs. The notation $u \rightsquigarrow v$ denotes the existence of a temporal path between u and v .

1.1 Contributions

The present work investigates temporal reachability and temporal spanners from a *probabilistic* point of view. To this end, we study a simple and natural model of random temporal graphs which can be seen as a temporal analog of Erdős–Rényi random graphs. More precisely, given a number of nodes n and a probability p , a *random simple temporal graph* (RSTG, for short) is obtained by taking an Erdős–Rényi random graph $(V, E) = G(n, p)$ and assigning to each edge $e \in E$ a *unique* presence time $\lambda(e) \in \{1, \dots, |E|\}$ according to a random permutation of all edges. The reachability in this model only depends on the relative order of edge labels; it can thus be equivalently defined by having $\lambda : E \rightarrow [0, 1]$ assign to each edge a presence time chosen uniformly at random (and independently) from the unit interval (with probability 1, all edges receive distinct labels). As will be discussed below, RSTGs are closely related to a number of other models. For example, they are mathematically equivalent to random edge-ordered graphs, although the interpretation of labels as time motivates the study of specific questions. RSTGs are also related to classical processes in gossip theory and population protocols. In a certain sense, they *encode* information pertaining to the ordering of interactions in these models. A significant difference, that also distinguishes RSTGs from stochastic models of dynamic networks like edge-markovian evolving graphs [9, 25], is that there are no repeated interactions among the same nodes. This creates dependencies between past and future events that make RSTGs typically more difficult to handle.

Overview of the results

Inspired by classical results on Erdős–Rényi graphs, we investigate natural analogs of connectivity thresholds in RSTGs. The fact that temporal reachability is neither transitive nor symmetric implies a number of different thresholds pertaining to gradual forms of connectivity. We consider the following properties: (1) a fixed node can reach another fixed node asymptotically almost surely (a.a.s.); (2) a.a.s. at least one node can reach all others; (3) a fixed node can a.a.s. reach all the others; and (4) every node can reach all the others a.a.s. (i.e., the graph is *temporally connected* a.a.s.).

The first set of results in this paper is a complete characterization of these thresholds, which follow a strikingly regular set of incremental values as follows. The first property occurs at $p = \log n/n$ (where \log is the natural logarithm), the second and third occur at $p = 2 \log n/n$, and the fourth at $p = 3 \log n/n$. All these thresholds are sharp; that is, for the threshold factors $c \in \{1, 2, 3\}$, the properties a.a.s. do not hold at $p = (c - o(1)) \log n/n$ and hold a.a.s. at $p = (c + o(1)) \log n/n$. These thresholds are summarized in Table 1.

Subsequently, we consider the densities at which various types of spanners start to exist. In particular, we are interested in the emergence of (5) nearly optimal spanners (of size $2n + o(1)$); (6) pivotal spanners (of size $2n - 2$), composed of two spanning trees glued at a single vertex, one descending in time, the other ascending subsequently; and (7) optimal

Property	Edge count	Sharp threshold	Reference
Nearly optimal spanner	$(2 + o(1))n$	$3 \log n/n$	Theorem 5.12
Optimal spanner	$2n - 4$	$\geq 3 \log n/n$	Theorem 5.12
		$\leq 4 \log n/n$	Theorem 5.13
		$= 4 \log n/n ?$	Conjecture 2
Pivotal spanner	$2n - 2$	$4 \log n/n$	Theorem 5.7

Table 2: Sharp thresholds for spanner properties in random simple temporal graphs.

spanners (of size $2n - 4$). Clearly, none of these structures can exist before $p = 3 \log n/n$, because the corresponding graphs are not even temporally connected. Interestingly, nearly optimal spanners emerge essentially as soon as the graph becomes temporally connected, i.e., the threshold is also $p = 3 \log n/n$. As for pivotal spanners and optimal spanners, both exist a.a.s. at $4 \log n/n$. For pivotal spanners, this constitutes a sharp threshold. Whether this is also the case for optimal spanners is left open—we conjecture that it is. These results are summarized in Table 2.

Significance of the results

In static graphs, no distinction exists between properties (2), (3), and (4). If a node u can reach all the others, then by transitivity and symmetry of the reachability relation, every node can reach all the others. Furthermore, spanning trees of $n - 1$ edges exist unconditionally. In standard Erdős–Rényi graphs, all these properties occur at the same point of the densification process, namely at $p = \log n/n$ [30]. As for Property (1), it occurs gradually before that point, without obeying a sharp threshold (see, e.g., [32, Section 2.2]). In view of these results, the fact that the temporal versions of the four properties occur at three distinct (and sharp) thresholds is quite remarkable. This can be seen as a fine-grained measure of the discrepancy between static and temporal reachability.

Conceptually, our results regarding the existence of spanners also admit a quite remarkable interpretation. Recall that the existence of linear spanners, and even subquadratic spanners, is not deterministically guaranteed in temporal graphs, as witnessed by infinite families of adversarial graphs in [41] and [7] (respectively). These *obstructions* to the existence of linear (and even sparse) spanners establish a fundamental difference between static and temporal reachability. Our results show that this difference does not hold probabilistically. In particular, our result on the existence of nearly optimal spanners,

Theorem 5.12. *There is $\delta = \delta(n) \in o(1)$ such that $p(n) = 3 \log n/n$ is a sharp threshold for an RSTG to admit a spanner of size at most $(2 + \delta)n$,*

combined with the fact that temporal connectivity itself arises at $p = 3 \log n/n$ (Theorem 5.4), establishes that the obstructions in [41] and [7]—and in fact, all conceivable obstructions—are statistically insignificant. In other words, nearly optimal spanners almost surely exist if the graph is temporally connected, and by analogy to static graphs, the universality of sparse spanners can be recovered at least in a probabilistic sense.

Drawing an analogy to *directed* static graphs, our characterization of a threshold at $p = 4 \log n/n$ for the existence of pivotal spanners of size $2n - 2$ is to be contrasted to the unconditional existence of analog constructions in strongly connected static graphs by Kosaraju-Sharir’s principle [50]. Indeed, the fact that this threshold is different from temporal connectivity (at $p = 3 \log n/n$) implies that such constructions are not universal, even in a probabilistic sense. Finally, whether optimal-size spanners (i.e., spanners of size

$2n - 4$) are probabilistically universal remains open. Our upper bound at $p = 4 \log n/n$ proves that this is at least close to being true.

Related models and questions

The RSTG model studied in this paper is a natural temporal analog of the Erdős–Rényi random graph model. Another analog considered in the literature is the model in which a random temporal graph is a sequence of independent Erdős–Rényi random graphs [23, 26]. This model and its generalizations have been studied in the literature from different perspectives [25, 6, 36, 9, 51, 2].

Our results can be related to a number of known results in the fields of gossip theory, population protocols, edge-ordered graphs, and random weighted graphs. Reviewing these connections is another contribution of this paper, and we dedicate Section 6 to it. The main observations can be summarized as follows.

A classical question in gossip theory asks the time it takes for a rumor to spread among a set of agents through random phone calls. Gossip models where interactions occur at random without repetition correspond to our model. However, in many cases, gossip protocols consider repeated interactions. We show that, even in this case, the early evolution of reachability is sufficiently close to RSTGs for the known thresholds to be essentially equivalent. A similar observation holds for some models of population protocols (see, e.g., [5]), where the pattern of interactions is specified by a sequential scheduler. Unless this scheduler is adversarial, it is generally assumed, for analysis, that the interactions are chosen uniformly at random with repetitions. Surprisingly, few connections have been made between gossip theory and population protocols, some analyses being re-discovered from time to time in both areas. Following the notations in [27], we will refer to the setting with repeated interactions as the ANY model and to the one without repetition as the CO (“call once”) model, irrespective of the areas in which they were considered.

Early works in the ANY model appear in a sequence of three papers from the ’70s [49, 13, 37], all under the title “*Random exchanges of information*” but with different authors. These papers are concerned with the time it takes for a fixed agent to receive a potential rumor from every other agent, which by symmetry corresponds to all the agents receiving a fixed single rumor—our third property. They also consider the time it takes for all agents to receive each other’s rumors—our fourth property (temporal connectivity). Asymptotics for the *expected* number of interactions are given in these papers, namely $n \log n$ for the former (which indeed corresponds to the number of edges in an RSTG with $p = 2 \log n/n$), and $1.5n \log n$ for the latter (corresponding to $p = 3 \log n/n$). In more recent works, the former was shown to be concentrated around its expected value (the best known result appears to be by Mocquard et al. [48]). As for temporal connectivity in the ANY model, it is already known that it does not exceed its expected value a.a.s. (see Burman et al. [17] for the currently best bound¹). In this paper, we show how our results can be translated to the ANY model. Doing so, we obtain concentration results (i.e., sharp thresholds) for these two properties, and beyond, for all others. Allowing for repetitions in the ANY model actually makes things easier to handle than in the CO model (the one equivalent to RSTGs), where handling dependencies between past and future calls requires additional efforts. According to [27], not much was known so far in that model, for the same reason. Therefore, our results on RSTGs fill a gap, by showing, among other things, that information propagates at essentially the same speed with or without repetitions.

¹[17] claims a weaker result (see Lemma 2.9 therein), but an intermediate step of the proof directly implies the upper bound of $(1.5 + o(1))n \log n$.

The fact that RSTGs relate to processes in gossip theory and population protocols is not surprising. Less intuitively, the gradual levels of temporal reachability that we observe in this work turn out to be also related to percolation processes in randomly edge-weighted complete graphs, as studied by Janson [39], provided one interprets the edge weights as waiting times. This connection is intriguing, because Janson’s model is not dynamic: the availability of the edges and the value of weights are invariant in time. Thus, in particular, the time required to cross an edge (i.e., the edge weight) is independent of all other factors, whereas in temporal graphs and gossip models two temporal paths, in general, need to wait different amounts of time to cross the same edge. Yet, some of our thresholds for temporal reachability are similar to percolation thresholds in this model; we discuss the reasons in Section 6.2. This agreement indicates that the essential features defining the connectivity properties of random graphs (weighted or temporal) coincide for a very wide range of models, even though the required proof techniques differ between them.

Finally, while the model of edge-ordered graphs is closest to RSTGs, research here has focused more on longest monotone paths and walks than on characterizing the reachability that results from monotone paths. Yet, the known results in randomly edge-ordered graphs translate directly to RSTGs, and vice versa.

1.2 Overview of the techniques

The following content is intended to help the reader navigate through the ideas and techniques used in this paper. In general, in order to prove that an RSTG with edge probability p has (or lacks) some property, we instead consider an RSTG with edge probability 1 and each edge label randomly chosen from $[0, 1]$. Then, we study the properties of the temporal subgraph induced by all the edges whose labels are at most p .

A preliminary result in Section 3 establishes that a.a.s. every vertex can reach every other vertex via a temporal path of length at most 2 when $p \geq \log n / \sqrt{n}$. This result is much weaker than the subsequent ones, but serves as a warm-up to the reader and allows us to subsequently assume that all relevant labels are contained in $[0, t]$ for some $t = o(1)$.

A concept that is quite versatile in the paper is that of a *foremost tree*, i.e., a tree consisting of gluing together a (prefix-stable) set of earliest arrival paths from a given vertex, called a temporal source, to all others. Foremost trees were considered for different purposes in [15] and [43], and a similar concept also appears in [42] (Problem 4.18). Section 4 offers a general study of foremost trees in RSTGs. In particular, we define a suitable martingale for estimating the time by which k vertices have been reached from a given vertex u , which allows us to apply Azuma’s inequality and derive concentration results. Precisely, we consider the sequence of waiting times for the respective next vertex to be reached from the source u and turn this into a martingale by subtracting the respective expected values. As the probability distribution of these waiting times has a one-sided long tail, we are able to achieve improved concentration bounds by analyzing capped versions of these variables (Lemma 4.6). This technique allows us to obtain a sharp threshold for the appearance of the first temporal source. On the other hand, these caps are chosen such that a.a.s. all elements of the sequence agree with their uncapped versions. By this technique, we also obtain sharp thresholds for the event that a fixed vertex reaches a particular number of other vertices (Lemma 4.12), which yields, in turn, sharp thresholds for temporal reachability between two fixed vertices and for a fixed vertex being a temporal source.

In order to characterize the threshold for temporal connectivity, we start by considering the moment when the first temporal sink arises (same as for temporal sources, by symmetry). Naively, one might be tempted to wait for the same amount of time subse-

quently for that vertex to reach all others back as a temporal source, which gives temporal connectivity by the pivoting principle (more on this below). However, this implies waiting for $4 \log n/n$ overall, whereas the actual threshold for temporal connectivity occurs at $3 \log n/n$ and requires a more general approach. Starting with the lower bound, we show that once the first temporal sink has appeared, an additional $\log n/n$ wait is needed for all the other vertices to get at least one more edge, which is necessary for them to become temporal sinks in turn (Lemma 5.5). The fact that temporal connectivity is equivalent to having all the vertices being temporal sinks (or sources, for that matter) implies, in turn, that $3 \log n/n$ is indeed a lower bound. As for the upper bound, we use the fact that the majority of vertices become temporal sinks at $2 \log n/n$ (soon after the first temporal sink appears), after which an additional $\log n/n$ wait guarantees that a.a.s. each remaining vertex becomes reachable from everywhere via at least one of these sinks (Lemma 5.6).

To establish the existence of small temporal spanners in Section 5.4, we give explicit constructions. For pivotal spanners, we require a pivot vertex that is a temporal sink before $2 \log n/n$ and a temporal source after that, but before $4 \log n/n$. The lower bounds are based on the above ones for reachability. Truly optimal spanners are built in a similar way, but instead of a single pivot, they rely on the concept of a pivotal square whose existence is guaranteed by time $4 \log n/n$. The result that $3 \log n/n$ is a sharp threshold for the existence of nearly optimal spanners requires a finer understanding of the height of foremost trees. To this end, we compute inductively the approximate distribution of vertex heights in a foremost tree, which implies a logarithmic upper bound on the number of edges in the foremost path between any two vertices (Section 4.3). Then, we divide the time into three intervals of length $\log n/n$ each. In the first interval, we find a temporal tree through which most vertices can reach a vertex that will act almost as a pivot. Analogously, in the last time interval, we build a temporal tree through which this vertex can reach most vertices. The remaining vertices are individually connected to the pivot by earliest and latest temporal paths whose edges are distinct from the edges of the trees. Finally, we separately connect all pairs of vertices that still cannot reach each other via the pivot. The contribution of these additional paths to the spanners is shown to be asymptotically negligible, which implies an overall size of $(2 + o(1))n$ edges.

1.3 Organization of the document

The main definitions are given in Section 2, including basic concepts and notations in temporal graphs. Section 3 presents an analysis that obtains weaker bounds than our actual results using simpler arguments; while proposed as a warm-up, these bounds are also utilized in subsequent analysis. Section 4 develops our main tools, by describing an algorithm that grows a *foremost tree* (i.e., a tree of time-optimal temporal paths) in a given temporal graph and analyzing the execution of this algorithm on a typical RSTG. Section 5 applies the tools from the previous section to obtain the main results, namely the claimed thresholds in RSTGs. Section 6 describes the adaptation of our analyses to models coming from gossip theory and population protocols. We also explain in detail how our results strengthen known results in these fields. Finally, Section 7 concludes with some remarks and open questions.

2 Preliminaries

2.1 Temporal graphs

A *temporal graph* \mathcal{G} is a pair (G, λ) , where $G = (V, E)$ is a simple undirected graph and λ is a function that assigns to every edge e of G a finite set of elements from some totally ordered set \mathbb{T} . The graph G is called the *underlying graph* of \mathcal{G} , the elements of $\lambda(e)$ are called *time labels* of e , and the pairs (e, t) , where $t \in \lambda(e)$ are called *time-edges*. The temporal graph \mathcal{G} is *simple* if every edge of G is assigned exactly one time label, i.e., $|\lambda(e)| = 1$ for every $e \in E$. We will mostly focus on simple temporal graphs, in which case with a slight abuse of notation we write $\lambda(e)$ to denote the unique time label of an edge $e \in E$ and refer to the time-edge $(e, \lambda(e))$ simply as edge e .

A *temporal (u, v) -path* or a *temporal path* from u to v in \mathcal{G} is a path $u = u_0, u_1, \dots, u_\ell = v$ in G such that $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_\ell$, where $\lambda_i \in \lambda(u_{i-1}u_i)$. The labels λ_1 and λ_ℓ are called the *departure time* and *arrival time* of the temporal path, respectively. A temporal (u, v) -path is called a *foremost (u, v) -path* if it has the earliest arrival time among all temporal (u, v) -paths. Symmetrically, a temporal (u, v) -path is called a *hindmost (u, v) -path* if it has the latest departure time among all temporal (u, v) -paths. A vertex $v \in V$ is called a *temporal source* of \mathcal{G} if every other vertex can be reached from v by a temporal path. Similarly, vertex v is called a *temporal sink* of \mathcal{G} if every other vertex can reach v by a temporal path. The temporal graph \mathcal{G} is *temporally connected* if each vertex can reach every other vertex by a temporal path. A temporal graph $\mathcal{G}' = (G', \lambda')$ is a *temporal subgraph* of \mathcal{G} if G' is a subgraph of G and $\lambda'(e) \subseteq \lambda(e)$ for every edge e of G' . Furthermore, if $V(\mathcal{G}') = V(\mathcal{G})$ and \mathcal{G}' is temporally connected, then \mathcal{G}' is called a (*temporal*) *spanner* of \mathcal{G} . All the temporal graphs we consider in this work are simple. As a result, spanners are uniquely determined by the set of underlying edges that they contain.

It is a straightforward consequence of a classical result in gossiping theory (see, e.g., [12]) that any spanner of an n -vertex simple temporal graph has at least $2n - 4$ edges. (The same holds more generally for time-edges in non-simple temporal graphs.) For this reason, we will refer to spanners with exactly $2n - 4$ edges as *optimal*. Another important type of temporal spanners is *pivotal* temporal spanners. Such a spanner is the union of two disjoint spanning trees rooted at the same vertex v , called *pivot*, such that for some time t all vertices can reach v in one of the trees before time t and v can reach all vertices in the other tree after time t . Clearly, a pivotal spanner contains $2n - 2$ edges. Finally, we define the *restriction of $\mathcal{G} = (G, \lambda)$ to a time interval $[a, b]$* , denoted by $\mathcal{G}_{[a,b]}$, as the temporal graph $\mathcal{G}' = (G', \lambda')$ where $\lambda'(e) = \lambda(e) \cap [a, b]$ and $G' = (V, \{e \in E \mid \lambda'(e) \neq \emptyset\})$.

2.2 Random simple temporal graphs (RSTGs)

The most basic and commonly studied model of random static graphs is the Erdős–Rényi model denoted by $G_{n,p}$, in which there are n vertices and every two of them are connected by an edge independently with probability $p \in [0, 1]$. A possible way of turning $G_{n,p}$ into a model of random *temporal* graphs is by choosing uniformly at random a total order on the set of edges. More specifically, a random temporal graph (G, λ) can be obtained by first drawing the underlying graph G from $G_{n,p}$ and then drawing λ uniformly at random from all bijections $E(G) \rightarrow \{1, 2, \dots, |E(G)|\}$. For technical convenience, however, we work with a slightly different but equivalent model which we denote by $\mathcal{F}_{n,p}$. In this model, the underlying graph is drawn from $G_{n,p}$ as before, but the labeling function λ now maps every edge of $E(G)$ to an independent and uniformly distributed label in the real interval $[0, 1]$. Since with probability 1 no two time labels are equal, this induces a total order on the edges, and by symmetry, all orders are equally likely; thus $\mathcal{F}_{n,p}$ is equivalent to the above model

of random edge orderings of $G_{n,p}$. (This equivalence has been used in some recent works on edge-ordered graphs [4, 43].) We refer to such a graph as a random simple temporal graph (RSTG), or just a random temporal graph, which is simple by default. Throughout the paper we will assume, without loss of generality, that in the graphs from $\mathcal{F}_{n,p}$ that we work with all edges have pairwise different labels.

Instead of working with $\mathcal{F}_{n,p}$ directly, it will often be convenient to first draw a temporal graph $\mathcal{G} = (G, \lambda)$ from $\mathcal{F}_{n,1}$ (resulting in G being a complete graph), and to then consider $\mathcal{G}' = (G', \lambda') = (G, \lambda)_{[0,p]}$ in which all edges with time labels greater than p are deleted. Observe that $G' \sim G_{n,p}$ and that each edge label $\lambda'(e)$ is uniformly distributed on $[0, p]$. In other words, \mathcal{G}' is distributed according to $\mathcal{F}_{n,p}$ up to multiplying all labels by $\frac{1}{p}$. Furthermore, since the temporal properties that we study depend only on the relative order of edge labels and not on their absolute values, without loss of generality, we will omit the multiplicative factor $\frac{1}{p}$ and will work directly with \mathcal{G}' . Similarly, for any $0 \leq a \leq b \leq 1$, up to rescaling $x \mapsto \frac{x-a}{b-a}$, the graph $(G, \lambda)_{[a,b]}$ is distributed according to $\mathcal{F}_{n,p}$, where $p = b - a$.

2.3 Degrees of reachability

A graph property is said to hold asymptotically almost surely (a.a.s.) if the probability that it is satisfied converges to 1 as n goes to infinity. A function $p = p(n)$ is a *sharp threshold* for a temporal graph property \mathcal{P} if for every $\varepsilon > 0$ a random temporal graph from $\mathcal{F}_{n,(1-\varepsilon)p}$ does not have property \mathcal{P} a.a.s., while a random temporal graph from $\mathcal{F}_{n,(1+\varepsilon)p}$ has \mathcal{P} a.a.s.

We study the following fundamental properties related to temporal reachability:

1. **Point-to-point Reachability.** The property that a fixed pair (u, v) of vertices has a temporal path from u to v .
2. **First Temporal Source.** The property that \mathcal{G} contains at least one temporal source.
3. **Temporal Source.** The property that a fixed vertex v is a temporal source, i.e., every vertex in \mathcal{G} can be reached from v by a temporal path.
4. **Temporal Connectivity.** The property that \mathcal{G} is temporally connected, which is equivalent to the property that all vertices in \mathcal{G} are temporal sources.
5. **Nearly Optimal Temporal Spanner.** The property that \mathcal{G} contains a temporal spanner with $(2 + o(1))n$ time-edges.
6. **Optimal Temporal Spanner.** The property that \mathcal{G} contains a temporal spanner with $2n - 4$ time-edges.
7. **Pivotal Temporal Spanner.** The property that \mathcal{G} contains a pivotal temporal spanner.

In Section 5, we establish sharp thresholds for all properties in RSTGs (see Table 1) except for the sixth, for which we prove an upper bound that we conjecture to be a sharp threshold as well.

3 Warm-up: 2-hop approach

This section presents a simple argument to derive upper bounds on temporal source related properties. The main statement says that for any $p \geq 3\sqrt{\log n/n}$ a fixed vertex v is a temporal source in a random temporal graph from $\mathcal{F}_{n,p}$ a.a.s. The simplicity of this approach

comes from the restriction that we consider only temporal paths of length 2. While these bounds are far from optimal, they turn out to be useful for subsequent analyses. Note that here and in the following, we often implicitly assume n to be at least an appropriate constant.

Lemma 3.1. *Let $\alpha \geq 3$ and let $p = \alpha\sqrt{\log n/n}$. Then an arbitrary vertex of $(G, \lambda) \sim \mathcal{F}_{n,p}$ is a temporal source with probability at least $1 - n^{-\alpha^2/4+1}$.*

Proof. Let x be an arbitrary vertex in (G, λ) . For two distinct vertices y, z of (G, λ) that are also different from x , we denote by R_z the event that x reaches z , and by S_{yz} we denote the event that x reaches z in exactly two steps via y . Notice that $\mathbb{P}[R_z]$ is the same for every $z \in V(G) \setminus \{x\}$, and we denote this probability by p_1 . Similarly, $\mathbb{P}[S_{yz}]$ is the same for all pairs $z, y \in V(G) \setminus \{x\}$, $z \neq y$, and is equal to $p^2/2$. Hence, denoting by γ the probability that x is a temporal source in (G, λ) and by using the union bound, we have

$$\gamma = 1 - \mathbb{P}\left(\bigcup_{z \neq x} \overline{R_z}\right) \geq 1 - \sum_{z \neq x} \mathbb{P}[\overline{R_z}] = 1 - (n-1)p_1.$$

Furthermore, for every $n \geq 4$, we have

$$\begin{aligned} p_1 = \mathbb{P}[R_x] &\leq \mathbb{P}\left[\bigcap_{y \neq x, z} \overline{S_{yz}}\right] = \prod_{y \neq x, z} \mathbb{P}[\overline{S_{yz}}] = \mathbb{P}[S_{yz}]^{n-2} \\ &= (1 - p^2/2)^{n-2} \leq e^{-\frac{p^2(n-2)}{2}} = \left(\frac{1}{n}\right)^{\alpha^2 \cdot \frac{n-2}{2n}} \leq \left(\frac{1}{n}\right)^{\frac{\alpha^2}{4}}, \end{aligned}$$

where we used the fact that all S_{yz} are independent. Consequently, we derive the desired conclusion

$$\gamma \geq 1 - (n-1)n^{-\alpha^2/4} \geq 1 - n^{-\alpha^2/4+1}.$$

□

If $\alpha = 3$ in the above lemma, we obtain that for $p = 3\sqrt{\log n/n}$ with probability at least $1 - 1/n^{5/4}$ an arbitrary vertex in $(G, \lambda) \sim \mathcal{F}_{n,p}$ is a temporal source. Similarly, if $\alpha = 4$, by applying the union bound, we obtain that for $p = 4\sqrt{\log n/n}$ with probability at least $1 - 1/n^2$ every vertex in $(G, \lambda) \sim \mathcal{F}_{n,p}$ is a temporal source, i.e., (G, λ) is temporally connected. However, in the subsequent analysis we will use the following corollary of the lemma obtained by taking $\alpha = \sqrt{\log n}$ and applying the union bound.

Corollary 3.2. *Let $p = \frac{\log n}{\sqrt{n}}$. Then, $(G, \lambda) \sim \mathcal{F}_{n,p}$ is temporally connected with probability at least $1 - n^{-\frac{\log n}{4}+2}$.*

4 Foremost tree evolution

Let $\mathcal{G} = (G, \lambda)$ be a random temporal graph from $\mathcal{F}_{n,1}$. Consider the probability that a fixed vertex v reaches another fixed vertex u in $\mathcal{F}_{n,p}$. This is equal to the probability that the temporal subgraph $\mathcal{G}_{[0,p]}$ contains a temporal (v, u) -path P . The latter is equivalent to the fact that the arrival time of P in \mathcal{G} is at most p . Therefore, the estimation of the parameter p for temporal reachability from v to u can be reduced to the estimation of the arrival time of a foremost temporal (v, u) -path in \mathcal{G} . In case of the Temporal Source property, we are interested in the smallest value of p such that a given vertex v is a temporal source in $\mathcal{G}_{[0,p]}$, or equivalently, that any vertex in \mathcal{G} can be reached from v by

time p (notice that since \mathcal{G} is a complete temporal graph, every vertex in \mathcal{G} is reachable from v , e.g., by a single-edge path). A minimal temporal subgraph that preserves foremost reachabilities from v to all other vertices reachable from v is called *foremost tree for v* (formally defined below), and we will be interested in estimating the smallest p such that \mathcal{G} contains a foremost tree for v with time labels of its edges not exceeding p . We proceed by defining formally the relevant notions.

Let (T, λ') be a temporal graph, where T is a tree, and let u be a vertex in T . If u is a temporal source (resp. temporal sink) in (T, λ') , then we say that (T, λ') is an *increasing temporal tree* (resp. *decreasing temporal tree*) rooted at u . A temporal subgraph (T, λ') of \mathcal{G} is called a *partial foremost tree for u in \mathcal{G}* if (T, λ') is an increasing temporal tree rooted at u and for every vertex v of T , the temporal (u, v) -path in (T, λ') is a foremost (u, v) -path in \mathcal{G} . Symmetrically, a temporal subgraph (T, λ') of \mathcal{G} is called a *partial hindmost tree for u in \mathcal{G}* if (T, λ') is a decreasing temporal tree rooted at u and for every vertex v of T , the temporal (v, u) -path in (T, λ') is a hindmost (v, u) -path in \mathcal{G} . A partial foremost tree (resp. partial hindmost tree) (T, λ') for u in \mathcal{G} is called *foremost tree* (resp. *hindmost tree*) for u in \mathcal{G} , if T contains all vertices that are reachable from u (resp. all vertices that can reach u) in \mathcal{G} .

Let (G, λ) be a simple temporal graph with time labels from an interval $[a, b]$. Let $\lambda_{[a,b]}^r : E(G) \rightarrow [a, b]$ be the mapping that mirrors the time labels at the middle point of $[a, b]$, i.e., $\lambda_{[a,b]}^r(e) := b - \lambda(e) + a$. Then $(G, \lambda_{[a,b]}^r)$ is also a simple temporal graph with time labels in $[a, b]$, but the order of time labels is reversed with respect to (G, λ) . In particular, any temporal (u, v) -path in (G, λ) is a temporal (v, u) -path in $(G, \lambda_{[a,b]}^r)$, and an increasing (resp. decreasing) temporal tree rooted at some vertex u in (G, λ) is a decreasing (resp. increasing) temporal tree rooted at u in $(G, \lambda_{[a,b]}^r)$. Also, any foremost (u, v) -path in (G, λ) becomes a hindmost (v, u) -path in $(G, \lambda_{[a,b]}^r)$ and vice versa; hence any (partial) foremost tree for u in (G, λ) is a (partial) hindmost tree for u in $(G, \lambda_{[a,b]}^r)$ and vice versa. Furthermore, since $\lambda_{[a,b]}^r(e)$ is a bijection, if $\mathcal{G} \sim \mathcal{F}_{n,1}$ and $(G', \lambda) = \mathcal{G}_{[a,b]}$, then the graphs (G', λ) and $(G', \lambda_{[a,b]}^r)$ have the same distribution. Due to this relation, in the present section, we will study only foremost paths and trees, but analogous results also hold for hindmost paths and trees. We will state specific results for hindmost paths and trees if they are needed subsequently.

The main purpose of the present section is to analyze temporal and structural properties of a typical foremost tree in $\mathcal{F}_{n,1}$. In Section 4.1 we present an algorithm for constructing a foremost tree for a given temporal source. In Section 4.2 we analyze the execution of the algorithm on a random instance from $\mathcal{F}_{n,1}$ to estimate the speed of growth of the foremost tree. We utilize these results in Section 5 for obtaining sharp thresholds.

4.1 Foremost tree algorithm

We start by presenting an algorithm that, given a temporal graph and a source² vertex v in the graph constructs a foremost tree for v . The algorithm is similar to Prim's algorithm for a minimum spanning tree in static graphs with the only difference being that the next edge to be added to the tree is chosen as the edge with minimum label *among those that extend the current tree to an increasing temporal tree*. Foremost trees have been previously studied by Bui-Xuan, Ferreira, and Jarry [15], as well as by Lavrov and Loh [43]. A similar idea also appears implicitly in Problem 4.18 of the classic algorithms book by Kleinberg and Tardos [42].

²For the sake of simplicity, we assume that the input vertex is a source; however, after trivial adjustments in the pseudocode of the algorithm will also work for any input vertex

Algorithm 1 FOREMOST TREE

Input: Simple temporal graph (G, λ) ; temporal source v in (G, λ) .

Output: Foremost tree for v .

- 1: $T_0 := (\{v\}, \emptyset)$
 - 2: **for** $k := 1$ to $n - 1$ **do**
 - 3: Let S_k be the set of edges in G with one endpoint in $V(T_{k-1})$ and the other endpoint in $V(G) \setminus V(T_{k-1})$
 - 4: $e_k := \arg \min\{\lambda(e) \mid e \in S_k \text{ and } T_{k-1} \cup \{e\} \text{ is an increasing temporal tree rooted at } v\}$
 - 5: $T_k := T_{k-1} \cup \{e_k\}$
- return** (T_{n-1}, λ') , where λ' is the restriction of λ to the edges of T_{n-1}
-

In the next lemma we prove the correctness of the algorithm and the easy but key fact that the time labels of the tree edges monotonically increase in the order in which they are added.

Lemma 4.1. *Let (G, λ) be a simple temporal graph and v be a temporal source in (G, λ) . Then*

- (i) *Algorithm 1 constructs a foremost tree for v in (G, λ) ;*
- (ii) $\lambda(e_1) \leq \lambda(e_2) \leq \dots \leq \lambda(e_{n-1})$.

Proof. To prove the first claim of the lemma, we will show by induction on k that

- (a) T_k is well-defined, i.e., the set from which e_k is chosen is non-empty;
- (b) for every $u \in V(T_k)$ the temporal (v, u) -path in T_k is a foremost (v, u) -path in (G, λ) .

The statement is obvious for $k = 0$. Let $1 \leq k \leq n - 1$ and assume the statement holds for $k - 1$. We will show that it also holds for k .

Let u be an arbitrary vertex in $V(G) \setminus V(T_{k-1})$, and let P be a foremost (v, u) -path in (G, λ) . Let $e = ab$ be the first edge of P with one endpoint, say a , in T_{k-1} and the other endpoint b not in T_{k-1} . Let also P' be the foremost (v, a) -path in T_{k-1} . Since P' is foremost in (G, λ) , the arrival time of P' is not more than the arrival time of the (v, a) -subpath of P . Consequently, the arrival time of P' is not larger than $\lambda(e)$, and therefore $T_{k-1} \cup \{e\}$ is an increasing temporal tree, which proves part (a) of the statement.

Let now $e_k = ab$ be the edge added to T_{k-1} to form T_k , where $a \in V(T_{k-1})$ and $b \notin V(T_{k-1})$. Taking into account the induction hypothesis, to prove part (b) of the statement, we only need to show that the temporal (v, b) -path in T_k is a foremost (v, b) -path in (G, λ) . Suppose it is not, and let P be a foremost (v, b) -path in (G, λ) . Let e' be the first edge of P with one endpoint in $V(T_{k-1})$ and the other endpoint in $V(T_k)$. By the proof of part (a), we know that $T_{k-1} \cup \{e'\}$ is an increasing temporal tree. Furthermore, since the arrival time of P is less than $\lambda(e_k)$, we have that $\lambda(e') < \lambda(e_k)$. But this contradicts the choice of e_k .

To prove the second claim of the lemma, assume that it does not hold and let $k \geq 2$ be the minimum index such that $\lambda(e_k) < \lambda(e_{k-1})$. Let $e_k = ab$, where $a \in V(T_{k-1})$ and $b \notin V(T_{k-1})$, and let e_i be the last edge of the (v, a) -path in T_k . Clearly $\lambda(e_i) \leq \lambda(e_k)$ because T_k is increasing. Hence, there exists j with $i < j < k$, such that $\lambda(e_{j-1}) \leq \lambda(e_k) < \lambda(e_j)$. Since $i \leq j - 1$, the edge e_i belongs to T_{j-1} , and therefore e_k can extend T_{j-1} to an increasing temporal tree. But this contradicts the choice of e_j at the j -th iteration of the algorithm, as $\lambda(e_k) < \lambda(e_j)$. \square

4.2 Foremost tree growth

The main goal of this section is to estimate the time by which a typical foremost tree in $\mathcal{F}_{n,1}$ acquires a given number of vertices. In this subsection, we always assume n to be at least a sufficiently large constant. To conduct our analysis, we will consider the execution of Algorithm 1 on a complete random temporal graph $(G, \lambda) \sim \mathcal{F}_{n,1}$ from some fixed vertex v as a random process that reveals the edges of the resulting foremost tree for v one by one in the order in which they are added. We define random variables $Y_0^v := 0$ and $Y_k^v := \lambda(e_k^v)$, $k \in [n-1]$, where, following the algorithm, $T_0^v := (\{v\}, \emptyset)$ and for every $k \in [n-1]$

$$\begin{aligned} S_k^v &:= V(T_{k-1}^v) \times (V(G) \setminus V(T_{k-1}^v)), \\ e_k^v &:= \arg \min \{ \lambda(e) \mid e \in S_k^v \text{ and } T_{k-1}^v \cup \{e\} \text{ is an increasing temporal tree} \}, \\ T_k^v &:= T_{k-1}^v \cup \{e_k^v\}. \end{aligned}$$

By definition, Y_k^v is the earliest time when the foremost tree for v contains exactly k edges, or equivalently the earliest time by which v can reach $k+1$ vertices (itself included). We note that since (G, λ) is temporally connected, Lemma 4.1 implies that all Y_k^v , $k \in [n-1]$, are finite. For $k \in [n-1]$, let X_k^v be a random variable equal to $Y_k^v - Y_{k-1}^v$, i.e., to the *waiting time* between the edges e_{k-1}^v and e_k^v . Clearly, we have

$$Y_k^v = \sum_{i=1}^k X_i^v$$

for every $k \in [n-1]$.

The main objects of our analysis are the random variables $X_1^v, X_2^v, \dots, X_{n-1}^v$ and $Y_1^v, Y_2^v, \dots, Y_{n-1}^v$. We will also study the behavior of their *truncated* versions, which are convenient in the applications: for $k \in [n-1]$, let $\hat{X}_k^v := \min\{X_k^v, c_k\}$ and $\hat{Y}_k^v := \sum_{i=1}^k \hat{X}_i^v$, where

$$c_k := \frac{2 \log(\min\{k, n-k\}) + \log \log n}{k(n-k)}.$$

The values of c_k are chosen in such a way that on the one hand they are sufficiently small to ensure the applicability of Azuma's inequality (Theorem 4.5); and on the other hand they are large enough to guarantee that the truncated variables coincide with their original versions a.a.s. The numerical conditions are formalized in the following lemma.

Lemma 4.2 (Properties of c_k). *For any sufficiently large n , we have*

- (i) $\sum_{i=1}^{n-1} c_i^2 \leq \frac{64(\log \log n)^2}{n^2}$;
- (ii) $\sum_{i=1}^{n-1} (1 - c_k)^{k(n-k)} \leq 4/\log n$

Proof. We start by proving the first part of the lemma.

$$\begin{aligned} \sum_{i=1}^{n-1} c_i^2 &= \sum_{i=1}^{n-1} \frac{(2 \log(\min\{i, n-i\}) + \log \log n)^2}{(i(n-i))^2} \leq \sum_{i=1}^{n-1} \frac{(2 \log(\min\{i, n-i\}) \cdot \log \log n)^2}{(i(n-i))^2} \\ &\leq 4(\log \log n)^2 \left(\sum_{i=1}^{\lfloor n/2 \rfloor} \frac{(\log i)^2}{i^2(n/2)^2} + \sum_{i=\lfloor n/2 \rfloor + 1}^{n-1} \frac{(\log(n-i))^2}{(n/2)^2(n-i)^2} \right) \\ &\leq \frac{32(\log \log n)^2}{n^2} \sum_{i=1}^{\infty} \frac{(\log i)^2}{i^2} \leq \frac{64(\log \log n)^2}{n^2}, \end{aligned}$$

where the first inequality holds for all sufficiently large n .

To prove the second part, we observe that

$$(1 - c_k)^{k(n-k)} \leq e^{-k(n-k)c_k} = \frac{1}{(\min\{k, n-k\})^2 \log n}, \quad (1)$$

and therefore

$$\begin{aligned} \sum_{k=1}^{n-1} (1 - c_k)^{k(n-k)} &\leq \sum_{k=1}^{n-1} \frac{1}{(\min\{k, n-k\})^2 \log n} \\ &= \frac{1}{\log n} \left(\sum_{k=1}^{\lfloor n/2 \rfloor} \frac{1}{k^2} + \sum_{k=\lfloor n/2 \rfloor + 1}^{n-1} \frac{1}{(n-k)^2} \right) \\ &\leq \frac{2}{\log n} \sum_{k=1}^{\infty} \frac{1}{k^2} \leq \frac{4}{\log n}. \end{aligned}$$

□

We will now estimate the expected time between the moments of exposing two consecutive edges of the foremost tree. More specifically, we will show that X_k^v and \widehat{X}_k^v coincide for a given vertex a.a.s., and bound the expected values of X_k^v and \widehat{X}_k^v given the information revealed by the process in the first $k-1$ steps. For every $k \in [n-1]$, let \mathcal{A}_k^v be the σ -algebra generated by the information revealed at the first k iterations of the algorithm starting at v , i.e., by the knowledge of the first k edges $e_1^v, e_2^v, \dots, e_k^v$ and their time labels $Y_1^v, Y_2^v, \dots, Y_k^v$. Let also \mathcal{A}_0^v be the trivial σ -algebra. Then, we have the following

Lemma 4.3. *For sufficiently large n , for a vertex v , with probability at least $1 - 4/\log n$ the equality $\widehat{X}_k^v = X_k^v$ holds for every $k \in [n-1]$.*

Further, for a vertex v and every $k \in [n-1]$, we have

$$\begin{aligned} (i) \quad & \frac{1 - Y_{k-1}^v}{k(n-k) + 1} \leq \mathbb{E}[X_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k) + 1}; \\ (ii) \quad & (1 - 1/\log n) \cdot \frac{1 - Y_{k-1}^v}{k(n-k) + 1} \leq \mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k) + 1}. \end{aligned}$$

Proof. For every $k \in [n-1]$, we define the function $w_k^v : S_k^v \rightarrow [0, 1]$ as follows

$$w_k^v(e) = \begin{cases} \lambda(e) - Y_{k-1}^v, & \lambda(e) \geq Y_{k-1}^v \\ \lambda(e) - Y_{k-1}^v + 1, & \lambda(e) < Y_{k-1}^v. \end{cases}$$

Notice that for any two edges $e, f \in S_k^v$ such that $\lambda(f) < Y_{k-1}^v \leq \lambda(e)$ we have $w_k^v(e) < w_k^v(f)$. Furthermore, it follows from Lemma 4.1 that S_k^v contains at least one edge e with $\lambda(e) \geq Y_{k-1}^v$. These two observations imply that e_k^v is exactly the edge on which the minimum of w_k^v is attained, that is,

$$e_k^v = \arg \min\{w_k^v(e) \mid e \in S_k^v\}, \quad (2)$$

and therefore, for every $k \in [n-1]$,

$$X_k^v = \min\{w_k^v(e) \mid e \in S_k^v\}. \quad (3)$$

Observe that upon exposure of edge e_k^v at step k of the algorithm, we reveal some information about the time labels of the other edges in S_k^v . More precisely, we learn that

these time labels are not contained in the interval $[Y_{k-1}^v, Y_k^v]$. Thus, if, for every $k \in [n-1]$ and $e \in S_k^v$, we inductively define the admissible range of $\lambda(e)$, as $I_k^v(e) := [0, 1]$ for $k = 1$, and for every $2 \leq k \leq n-1$

$$I_k^v(e) := \begin{cases} I_{k-1}^v(e) \setminus [Y_{k-2}^v, Y_{k-1}^v], & e \in S_k^v \cap S_{k-1}^v \\ [0, 1], & e \in S_k^v \setminus S_{k-1}^v \end{cases}$$

then $\lambda(e)$ conditioned on \mathcal{A}_{k-1}^v is uniformly distributed on $I_k^v(e)$. Let $\ell = \ell(e)$ be the unique index with $e \in S_\ell^v \setminus S_{\ell-1}^v$. Then, we have $I_k^v(e) = [0, 1] \setminus [Y_{\ell-1}^v, Y_{k-1}^v]$. Thus, by definition, if $\lambda(e) \geq Y_{k-1}^v$, then $w_k^v(e)$ is uniformly distributed on $[0, 1 - Y_{k-1}^v]$, and if $\lambda(e) < Y_{k-1}^v$, then $w_k^v(e)$ is uniformly distributed on $[1 - Y_{k-1}^v, 1 + Y_{\ell-1}^v - Y_{k-1}^v]$. It follows that $w_k^v(e)$ is uniformly distributed on its admissible range

$$J_k^v(e) := [0, 1 + Y_{\ell-1}^v - Y_{k-1}^v]$$

and clearly

$$[0, 1 - Y_{k-1}^v] \subseteq J_k^v(e) \subseteq [0, 1]. \quad (4)$$

Note that X_k^v is a minimum of $k(n-k)$ independent random variables $w_k^v(e)$, $e \in S_k^v$, where for every edge $e \in S_k^v$ the value $w_k^v(e)$ is distributed uniformly on its own admissible range $J_k^v(e)$. Let X_k' be the minimum of $k(n-k)$ independent random variables distributed uniformly in $[0, 1 - Y_{k-1}^v]$, and X_k'' be the minimum of $k(n-k)$ independent random variables distributed uniformly in $[0, 1]$. Then, the inclusion (4) implies

$$\mathbb{P}[X_k' \geq t] \leq \mathbb{P}[X_k^v \geq t \mid A_{k-1}^v] \leq \mathbb{P}[X_k'' \geq t], \quad (5)$$

where A_{k-1}^v is the event that specific edges $e_1^v, e_2^v, \dots, e_{k-1}^v$ with time labels $Y_1^v, Y_2^v, \dots, Y_{k-1}^v$ are revealed in the first $k-1$ steps of the process. In particular, $\mathbb{P}(X_k^v \geq c_k) \leq (1 - c_k)^{k(n-k)}$. Thus, by the union bound and Lemma 4.2, the probability that for some k we have $X_k^v \neq \widehat{X}_k^v$ is at most $4/\log n$. This proves the first part of the lemma.

Now, by integrating inequality (5) and noting that the expected value of the minimum of m independent variables distributed uniformly on $[0, a]$ is equal to $\frac{a}{m+1}$, we obtain

$$\frac{1 - Y_{k-1}^v}{k(n-k) + 1} = \mathbb{E}[X_k'] \leq \mathbb{E}[X_k^v \mid \mathcal{A}_{k-1}^v] \leq \mathbb{E}[X_k''] = \frac{1}{k(n-k) + 1}. \quad (6)$$

To prove item (ii) of the lemma, we first note that by definition $\widehat{X}_k^v \leq X_k^v$, and hence

$$\mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v] \leq \mathbb{E}[X_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k) + 1}.$$

Therefore, it remains to show the lower bound on $\mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v]$. For convenience, let us write M_k for $k(n-k)$. Then, we have

$$\begin{aligned}
\mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v] &= \int_0^\infty \mathbb{P}[\widehat{X}_k^v \geq t \mid \mathcal{A}_{k-1}^v] dt = \int_0^{c_k} \mathbb{P}[\widehat{X}_k^v \geq t \mid \mathcal{A}_{k-1}^v] dt \\
&= \int_0^{c_k} \mathbb{P}[X_k^v \geq t \mid \mathcal{A}_{k-1}^v] dt \geq \int_0^{c_k} \mathbb{P}[X_k' \geq t] dt \\
&= \int_0^{c_k} \left(1 - \frac{t}{1 - Y_{k-1}^v}\right)^{M_k} dt \\
&= \frac{(1 - Y_{k-1}^v) + (c_k - (1 - Y_{k-1}^v)) \left(1 - \frac{c_k}{1 - Y_{k-1}^v}\right)^{M_k}}{M_k + 1} \\
&\geq \frac{(1 - Y_{k-1}^v) - (1 - Y_{k-1}^v)(1 - c_k)^{M_k}}{M_k + 1} \\
&= \frac{(1 - Y_{k-1}^v) (1 - (1 - c_k)^{M_k})}{M_k + 1},
\end{aligned}$$

and the desired bound follows from the fact that $(1 - c_k)^{k(n-k)} \leq \frac{1}{\log n}$ (see (1)). \square

Recall that, by Corollary 3.2, $\mathcal{F}_{n, (\log n)/\sqrt{n}}$ is temporally connected with probability at least $1 - n^{-\frac{\log n}{4} + 2}$, in which case the bound $Y_k^v \leq Y_{n-1}^v \leq \frac{\log n}{\sqrt{n}}$ holds for every v and $k \in [n - 1]$. Thus, we derive the following

Corollary 4.4. *With probability at least $1 - n^{-\frac{\log n}{4} + 2}$, for every vertex v and every $k \in [n - 1]$, we have*

$$\begin{aligned}
(i) \quad & \left(1 - \frac{\log n}{\sqrt{n}}\right) \cdot \frac{1}{k(n-k) + 1} \leq \mathbb{E}[X_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k) + 1}; \\
(ii) \quad & \left(1 - \frac{2}{\log n}\right) \cdot \frac{1}{k(n-k) + 1} \leq \mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k) + 1}.
\end{aligned}$$

Next, we will bound the deviation of the truncated time of the moment when the foremost tree acquires k edges from the expected value of accumulated truncated waiting times between the consecutive edges in the sequence of the first k edges of the tree. For this we require the following standard inequality by Azuma.

Theorem 4.5 (Azuma's inequality [8]). *Let Z_0, Z_1, \dots, Z_n be a martingale with respect to a filtration $\{\emptyset, \Omega\} = \mathcal{A}_0 \subset \mathcal{A}_1 \subset \dots \subset \mathcal{A}_n$. Let also c_1, c_2, \dots, c_n be non-negative numbers such that $\sum_{i=1}^n \mathbb{P}[|Z_i - Z_{i-1}| \geq c_i] = 0$. Then*

$$\mathbb{P}[|Z_n - Z_0| \geq \mu] \leq 2 \exp\left(\frac{-\mu^2}{2 \sum_{i=1}^n c_i^2}\right).$$

Lemma 4.6. *For a fixed vertex v and any sufficiently large n , with probability at least $1 - n^{-\sqrt{\log n} - 1}$ the inequality*

$$\left| \widehat{Y}_k^v - \sum_{i=1}^k \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] \right| < \frac{(\log n)^{0.8}}{n}$$

holds for all $k \in [n - 1]$.

Proof. Let us fix $k \in [n-1]$ and define a martingale $Z_0^v, Z_1^v, \dots, Z_k^v$ with $Z_0^v := 0$ and

$$Z_s^v := \widehat{Y}_s^v - \sum_{i=1}^s \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] = \sum_{i=1}^s \widehat{X}_i^v - \sum_{i=1}^s \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v],$$

for $s \in [k-1]$.

Since $0 \leq \widehat{X}_i^v \leq c_i$, we have $0 \leq \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] \leq c_i$, and therefore

$$\mathbb{P} \left[|Z_i^v - Z_{i-1}^v| > c_i \right] = \mathbb{P} \left[\left| \widehat{X}_i^v - \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] \right| > c_i \right] = 0$$

holds for every $i \in [k]$. Furthermore, by Lemma 4.2 (i), for all sufficiently large n ,

$$\sum_{i=1}^k c_i^2 \leq \sum_{i=1}^{n-1} c_i^2 \leq \frac{64(\log \log n)^2}{n^2},$$

and hence applying Azuma's inequality (Theorem 4.5), we obtain that for sufficiently large n

$$\begin{aligned} \mathbb{P} \left[|Z_k^v| \geq \frac{(\log n)^{0.8}}{n} \right] &\leq 2 \exp \left(\frac{-(\log n)^{1.6}}{n^2} \cdot \frac{1}{2 \sum_{i=1}^k c_i^2} \right) \\ &\leq 2 \exp \left(\frac{-(\log n)^{1.6}}{n^2} \cdot \frac{n^2}{128(\log \log n)^2} \right) \\ &= 2 \exp \left(\frac{-(\log n)^{1.6}}{128(\log \log n)^2} \right) \\ &\leq \exp \left(-(\log n)^{1.55} \right) \\ &\leq n^{-\sqrt{\log n}-2}. \end{aligned}$$

The latter inequality together with the union bound over all $k \in [n-1]$ imply the desired result. \square

The following technical inequality will be useful in the rest of the section.

Lemma 4.7. *For all integers $n \geq 1$ and $0 \leq k \leq n-1$, we have*

$$\begin{aligned} \frac{\log(k+1) + \log n - \log(n-k)}{n} - \frac{3}{n} &\leq \sum_{i=1}^k \frac{1}{i(n-i)+1} \leq \sum_{i=1}^k \frac{1}{i(n-i)} \\ &\leq \frac{\log(k+1) + \log n - \log(n-k)}{n} + \frac{3}{n} \end{aligned}$$

Proof. We have

$$\begin{aligned} \sum_{i=1}^k \frac{1}{i(n-i)+1} &\leq \sum_{i=1}^k \frac{1}{i(n-i)} = \frac{1}{n} \sum_{i=1}^k \left(\frac{1}{i} + \frac{1}{n-i} \right) \\ &\leq \frac{\log(k+1) + \log n - \log(n-k)}{n} + \frac{3}{n} \end{aligned}$$

as well as

$$\begin{aligned}
\sum_{i=1}^k \frac{1}{i(n-i)+1} &\geq \sum_{i=1}^k \frac{1}{i(n+1-i)} = \frac{1}{n+1} \sum_{i=1}^k \left(\frac{1}{i} + \frac{1}{n+1-i} \right) \\
&\geq \frac{\log(k+1) + \log n - \log(n-k)}{n+1} - \frac{1}{n+1} \\
&\geq \frac{\log(k+1) + \log n - \log(n-k)}{n} - \frac{2 \log n}{n^2} - \frac{1}{n} \\
&\geq \frac{\log(k+1) + \log n - \log(n-k)}{n} - \frac{3}{n}.
\end{aligned}$$

□

Now we are ready to prove the main results of this section.

Theorem 4.8. *With probability at least $1 - 2n^{-\sqrt{\log n}}$, for every vertex v , any large enough n , and $k \in [n-1]$, we have*

$$\left| \widehat{Y}_k^v - \sum_{i=1}^k \frac{1}{i(n-i)+1} \right| < \frac{2(\log n)^{0.8}}{n}.$$

Proof. By Corollary 4.4 (ii), with probability at least $1 - n^{-\frac{\log n}{4}+2}$, for every vertex v and every $k \in [n-1]$, we have

$$-\sum_{i=1}^k \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] - \frac{2}{\log n} \cdot \sum_{i=1}^k \frac{1}{i(n-i)+1} \leq -\sum_{i=1}^k \frac{1}{i(n-i)+1} \leq -\sum_{i=1}^k \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v]. \tag{7}$$

Similarly, by the union bound and Lemma 4.6, with probability at least $1 - n^{-\sqrt{\log n}}$ for every vertex v , any sufficiently large n , and every $k \in [n-1]$, we have

$$\sum_{i=1}^k \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] - \frac{(\log n)^{0.8}}{n} \leq \widehat{Y}_k^v \leq \sum_{i=1}^k \mathbb{E}[\widehat{X}_i^v \mid \mathcal{A}_{i-1}^v] + \frac{(\log n)^{0.8}}{n}. \tag{8}$$

Hence, summing up (7) and (8), we conclude that with probability at least $1 - 2n^{-\sqrt{\log n}}$ for every vertex v and every $k \in [n-1]$

$$-\frac{(\log n)^{0.8}}{n} - \frac{2}{\log n} \cdot \sum_{i=1}^k \frac{1}{i(n-i)+1} \leq \widehat{Y}_k^v - \sum_{i=1}^k \frac{1}{i(n-i)+1} \leq \frac{(\log n)^{0.8}}{n},$$

which implies the result after noticing that by Lemma 4.7 and for large enough n

$$\frac{2}{\log n} \cdot \sum_{i=1}^k \frac{1}{i(n-i)+1} \leq \frac{2}{\log n} \cdot \left(\frac{2 \log n}{n} + \frac{3}{n} \right) \leq \frac{10}{n} \leq \frac{(\log n)^{0.8}}{n}.$$

□

Corollary 4.9. *With probability at least $1 - 2n^{-\sqrt{\log n}}$, for every vertex v and any large enough n , we have*

$$\left| \widehat{Y}_{n-1}^v - \frac{2 \log n}{n} \right| < \frac{3(\log n)^{0.8}}{n}.$$

Proof. By Lemma 4.7 we have

$$\left| \sum_{i=1}^{n-1} \frac{1}{i(n-i)+1} - \frac{2 \log n}{n} \right| \leq \frac{3}{n}.$$

Hence, from Theorem 4.8 we conclude that with probability at least $1 - 2n^{-\sqrt{\log n}}$ the inequality

$$\begin{aligned} \left| \widehat{Y}_{n-1}^v - \frac{2 \log n}{n} \right| &\leq \left| \widehat{Y}_{n-1}^v - \sum_{i=1}^{n-1} \frac{1}{i(n-i)+1} \right| + \left| \sum_{i=1}^{n-1} \frac{1}{i(n-i)+1} - \frac{2 \log n}{n} \right| \\ &\leq \frac{2(\log n)^{0.8}}{n} + \frac{3}{n} < \frac{3(\log n)^{0.8}}{n} \end{aligned}$$

holds for every vertex v and any large enough n . \square

Theorem 4.8 and Lemma 4.3 (introductory claim) imply the following analog of Theorem 4.8 for the non-truncated random variables.

Theorem 4.10. *For a fixed vertex v and any large enough n , with probability at least $1 - 5/\log n$, for every $k \in [n-1]$, we have*

$$\left| Y_k^v - \sum_{i=1}^k \frac{1}{i(n-i)+1} \right| < \frac{2(\log n)^{0.8}}{n}.$$

The latter statement together with Lemma 4.7 immediately implies the following

Corollary 4.11. *For a fixed vertex v and any large enough n , with probability at least $1 - 5/\log n$, for every $0 \leq k \leq n-1$, we have*

$$\left| Y_k^v - \frac{\log n + \log(k+1) - \log(n-k)}{n} \right| < \frac{2(\log n)^{0.8} + 3}{n}.$$

Finally, we use Corollary 4.11 to prove the following lemma that will be utilized repeatedly throughout the paper.

Lemma 4.12. *For every function $z = z(n)$ with $0 \leq z(n) \leq 1$, and any constant $y > 0$ there is $n_0 := n_0(y)$ such that for all $n \geq n_0$ a fixed vertex v in $\mathcal{G} \sim \mathcal{F}_{n,p}$ can reach (resp. be reached by)*

- (i) at least $\left\lceil \frac{n^z}{(\log n)^y} \right\rceil$ vertices with probability at least $1 - \frac{5}{\log n}$, if $p \geq z \frac{\log n}{n} + \frac{3(\log n)^{0.8}}{n}$;
- (ii) fewer than $\left\lceil \frac{n^z}{(\log n)^y} \right\rceil$ vertices with probability at least $1 - \frac{5}{\log n}$, if $p \leq z \frac{\log n}{n} - \frac{3(\log n)^{0.8}}{n}$;
- (iii) at least $n - \left\lceil \frac{n^z}{(\log n)^y} \right\rceil$ vertices with probability at least $1 - \frac{5}{\log n}$, if $p \geq (2-z) \frac{\log n}{n} + \frac{3(\log n)^{0.8}}{n}$;
- (iv) fewer than $n - \left\lceil \frac{n^z}{(\log n)^y} \right\rceil$ vertices with probability at least $1 - \frac{5}{\log n}$, if $p \leq (2-z) \frac{\log n}{n} - \frac{3(\log n)^{0.8}}{n}$.

Proof. Define $d := \left\lceil \frac{n^z}{(\log n)^y} \right\rceil$. By choosing n_0 large enough, we may assume $1 \leq d \leq n-1$. To prove the “can reach” part of (i) and (ii), it suffices to show that for a fixed vertex v , with probability at least $1 - 5/\log n$, we have

$$\left| Y_{d-1}^v - z \frac{\log n}{n} \right| < \frac{3(\log n)^{0.8}}{n}.$$

Observe that

$$\begin{aligned} L &:= \log n + \log d - \log(n-d+1) \\ &= \log n + \log \left(\frac{n^z}{(\log n)^y} + 1 \right) - \log \left(n - \frac{n^z}{(\log n)^y} \right) + O(1) \\ &= z \log n + \log \left(\frac{1}{(\log n)^y} + \frac{1}{n^z} \right) - \log \left(1 - \frac{n^{z-1}}{(\log n)^y} \right) + O(1) \\ &= z \log n + O(\log \log n). \end{aligned}$$

Then we conclude by Corollary 4.11 that with probability at least $1 - 5/\log n$

$$\begin{aligned} \left| Y_{d-1}^v - z \frac{\log n}{n} \right| &\leq \left| Y_{d-1}^v - \frac{L}{n} \right| + \left| \frac{L}{n} - z \frac{\log n}{n} \right| \\ &\leq \frac{2(\log n)^{0.8} + 3}{n} + \frac{O(\log \log n)}{n} < \frac{3(\log n)^{0.8}}{n} \end{aligned}$$

holds, given that n_0 (and thus n) is at least a sufficient constant.

The “can reach” statement of (iii) and (iv) is proven analogously by showing that for $d' := \left\lfloor \frac{n^z}{(\log n)^y} \right\rfloor$ we have

$$\begin{aligned} L' &:= \log n + \log(n-d') - \log(d'+1) \\ &= (2-z) \log n + O(\log \log n). \end{aligned}$$

and thus with probability at least $1 - 5/\log n$

$$\left| Y_{n-d'-1}^v - (2-z) \frac{\log n}{n} \right| \leq \left| Y_{n-d'-1}^v - \frac{L'}{n} \right| + \left| \frac{L'}{n} - (2-z) \frac{\log n}{n} \right| < \frac{3(\log n)^{0.8}}{n}.$$

Finally, the “can be reached by” part of the theorem holds by symmetry under time reversal. \square

4.3 Foremost tree height

In this section we study the height of the foremost trees T_{n-1}^v in a random temporal complete graph $(G, \lambda) \sim \mathcal{F}_{n,1}$. We prove that a.a.s. each foremost tree T_{n-1}^v has at most logarithmic height.

We say that an endpoint u of e_k^v is the *attachment vertex* of e_k^v , if u belongs to T_{k-1}^v . Recall that $e_1^v, e_2^v, \dots, e_{k-1}^v$ denote the edges that are revealed in the first $k-1$ steps of the process and that form T_{k-1}^v ; and $Y_1^v, Y_2^v, \dots, Y_{k-1}^v$ denote the time labels of these edges, respectively. In the next lemma we show that the attachment vertex of e_k^v is distributed almost uniformly in $V(T_{k-1}^v)$. Denote by \mathcal{T}_{k-1}^v the family of all k -vertex trees rooted at v and with vertices from $[n]$.

Lemma 4.13. *There exists a function $\varepsilon = \varepsilon(n) \in O\left(\frac{\log n}{\sqrt{n}}\right)$ such that for all $n \geq 2$, $k \in [n-1]$, and any $T \in \mathcal{T}_{k-1}^v$*

$$\max_{u \in V(T)} \left| \mathbb{P}[u \in e_k^v \mid T_{k-1}^v = T] - \frac{1}{k} \right| \leq \frac{\varepsilon}{k}.$$

Proof. Recall that the event $Y_{n-1}^v \leq \frac{\log n}{\sqrt{n}}$ happens with probability at least $1 - n^{-\frac{\log n}{4} + 2} =: 1 - \varepsilon'$ by Corollary 3.2. Using the notation $w_k^v(e)$ and $J_k^v(e)$ introduced in the proof of Lemma 4.3, recall that for each $k \in [n-1]$, we have

$$e_k^v = \arg \min\{w_k^v(e) \mid e \in S_k^v\},$$

where the variable $w_k^v(e)$ is uniformly distributed on $J_k^v(e)$. As long as $Y_{k-1}^v \leq \frac{\log n}{\sqrt{n}}$, by (4), the interval $J_k^v(e)$ has length between $1 - \frac{\log n}{\sqrt{n}}$ and 1.

Since, for each edge $e \in S_k^v$, the probability of the event $e = e_k^v$ is inversely proportional to the length of $J_k^v(e)$, conditional on $Y_{k-1}^v \leq \frac{\log n}{\sqrt{n}}$ this probability can deviate from $1/|S_k^v|$ at most by a factor of $1 - \frac{\log n}{\sqrt{n}} =: 1 - \varepsilon''$. As each vertex of T_{k-1}^v is incident with the same number of edges in S_k^v , the probability of some vertex u of the tree being incident with e_k^v is thus within a factor of $1 - \varepsilon''$ of $\frac{1}{k}$. Consequently, we have

$$\frac{1 - \varepsilon''}{k} \leq \mathbb{P}\left[u \in e_k^v \mid T_{k-1}^v = T, Y_{k-1}^v \leq \frac{\log n}{\sqrt{n}}\right] \leq \frac{1}{k(1 - \varepsilon'')}.$$

Now, since the event $Y_{k-1}^v \leq \frac{\log n}{\sqrt{n}}$ is implied by $Y_{n-1}^v \leq \frac{\log n}{\sqrt{n}}$, it happens with probability at least $1 - \varepsilon'$, and therefore

$$(1 - \varepsilon') \frac{1 - \varepsilon''}{k} \leq \mathbb{P}[u \in e_k^v \mid T_{k-1}^v = T] \leq \frac{1 - \varepsilon'}{k(1 - \varepsilon'')} + \varepsilon'.$$

When setting $\varepsilon := n\varepsilon' + 4\varepsilon'' \in O\left(\frac{\log n}{\sqrt{n}}\right)$, this yields

$$\left| \mathbb{P}[u \in e_k^v \mid T_{k-1}^v = T] - \frac{1}{k} \right| \leq \frac{\varepsilon}{k},$$

where we used that

$$(1 - \varepsilon') \frac{1 - \varepsilon''}{k} > \frac{1 - \varepsilon' - \varepsilon''}{k} > \frac{1 - \varepsilon}{k}$$

and, since $\varepsilon'' < \frac{3}{4}$,

$$\frac{1 - \varepsilon'}{k(1 - \varepsilon'')} + \varepsilon' < \frac{1}{k(1 - \varepsilon'')} + \varepsilon' < \frac{1 + 4\varepsilon''}{k} + \varepsilon' < \frac{1 + \varepsilon}{k}.$$

□

Let now $N(h, k, n)$ denote the number of vertices of height h in the k -vertex partial foremost tree T_{k-1}^v in $(G, \lambda) \sim \mathcal{F}_{n,1}$, where the height of a vertex is the distance from the vertex to the root v .

Lemma 4.14. *There exists a positive function $\varepsilon = \varepsilon(n) \in O\left(\frac{\log n}{\sqrt{n}}\right)$ such that for all $n \geq 2$, $k \in [n-1]$, and $h \in [k]$,*

$$\frac{1 - \varepsilon}{k} \mathbb{E}[N(h-1, k, n)] \leq \mathbb{E}[N(h, k+1, n)] - \mathbb{E}[N(h, k, n)] \leq \frac{1 + \varepsilon}{k} \mathbb{E}[N(h-1, k, n)].$$

Proof. Let ε be as in Lemma 4.13. The result follows from the observation that $N(h, k + 1, n) - N(h, k, n)$ is either 1 or 0, depending on whether the edge e_k^v was attached to a vertex of height $h - 1$ or not, which, by Lemma 4.13, happens with probability between $(1 - \varepsilon)N(h - 1, k, n)/k$ and $(1 + \varepsilon)N(h - 1, k, n)/k$. \square

Lemma 4.15. *For all n sufficiently large, $k \in [n - 1]$, and $0 \leq h \leq k$, we have the inequality*

$$\mathbb{E}[N(h, k, n)] \leq \frac{(4 \log k)^h}{h!}.$$

Proof. We prove the lemma by double induction on h, k . Clearly, the statement holds whenever $k = 1$ or $h = 0$. Hence, assuming $h \geq 1$, we derive

$$\begin{aligned} \mathbb{E}[N(h, k + 1, n)] &\leq \mathbb{E}[N(h, k, n)] + \frac{2}{k} \mathbb{E}[N(h - 1, k, n)] \quad \text{by Lemma 4.14} \\ &\leq \frac{(4 \log k)^h}{h!} + \frac{2(4 \log k)^{h-1}}{k(h-1)!} \quad \text{by induction hypothesis} \\ &= \frac{(4 \log k)^h}{h!} \left(1 + \frac{2h}{k(4 \log k)} \right) \\ &\leq \frac{(4 \log k)^h}{h!} \left(1 + \frac{2}{k(4 \log k)} \right)^h \\ &= \frac{(4 \log k + 2/k)^h}{h!} \\ &\leq \frac{(4 \log(k + 1))^h}{h!} \end{aligned}$$

where the last inequality holds because

$$4 \log k + \frac{2}{k} = 4 \left(\log k + \frac{1}{2k} \right) \leq 4 \left(\log k + \frac{1}{k+1} \right) \leq 4 \log(k + 1).$$

\square

On a side note, we believe the following more precise asymptotic estimate of $\mathbb{E}[N(h, k, n)]$ to be true.

Conjecture 1.

$$\lim_{k \rightarrow \infty} \sup_{\substack{h, n \\ h \leq k \leq n}} \left| \mathbb{E}[N(h, k, n)] - \frac{(\log k)^h}{h!} \right| = 0.$$

We are now ready to state and prove the main result of the section.

Theorem 4.16. *Let $\mathcal{G} = (G, \lambda) \sim \mathcal{F}_{n,1}$. Then a.a.s. for all vertices v in \mathcal{G} and all $0 \leq \tau \leq \tau' \leq 1$ simultaneously, the foremost tree for v in $\mathcal{G}_{[\tau, \tau']}$ has height at most $14 \log n$.*

Proof. To prove the theorem, we will show that the foremost trees constructed by Algorithm 1 have the desired height. First, without loss of generality we can assume that $\tau, \tau' \in \Lambda := \{\lambda(e) \mid e \in E(G)\}$. Furthermore, it follows from Algorithm 1 that the foremost tree for v in $\mathcal{G}_{[\tau, \tau']}$ is a subtree of the foremost tree for v in $\mathcal{G}_{[\tau, 1]}$. Hence, it suffices to show the claim when $\tau' = 1$. Let $\ell_\tau(v)$ denote the number of vertices of height exactly $\lceil 14 \log n \rceil$ in $T_{v, \tau}$, which denotes the foremost tree for v in $\mathcal{G}_{[\tau, 1]}$. To prove the theorem, it suffices to show that $s := \max\{\ell_\tau(v) \mid v \in V, \tau \in \Lambda\}$ is zero a.a.s.

Since $\mathcal{G}_{[\tau,1]}$ is distributed according to $\mathcal{F}_{1-\tau,n}$ (up to rescaling), the number of vertices of height h in $T_{v,\tau}$ is distributed as the number of vertices of height h in the foremost tree for v in $\mathcal{G}_{[0,1-\tau]}$, which is upper-bounded by the number of vertices of height h in the foremost tree T_{n-1}^v of \mathcal{G} . Therefore we obtain from Lemma 4.15 that

$$\begin{aligned} \mathbb{E}[s] &\leq \mathbb{E} \left[\sum_{v \in V} \sum_{\tau \in \Lambda} \ell_\tau(v) \right] \leq n^3 \cdot \mathbb{E} [N(\lceil 14 \log n \rceil, n, n)] \leq n^3 \frac{(4 \log n)^{\lceil 14 \log n \rceil}}{\lceil 14 \log n \rceil!} \\ &\leq n^3 \left(\frac{4e \log n}{\lceil 14 \log n \rceil} \right)^{\lceil 14 \log n \rceil} \leq n^3 \left(\frac{4e}{14} \right)^{14 \log n} \quad \text{since } \frac{4e}{14} < 1 \\ &= n^{17} \left(\frac{2}{7} \right)^{14 \log n} = n^{17+14 \log(2/7)} \approx n^{-0.54} \in o(1). \end{aligned}$$

Thus, by Markov's inequality, a.a.s. $s = 0$ as required. \square

A theorem similar to Theorem 4.16 holds also for hindmost trees. The following result is an implication of these theorems and the observation that a temporal graph \mathcal{G} contains a foremost (resp. hindmost) (v, u) -path if and only if any foremost tree for v (resp. hindmost tree for u) in \mathcal{G} contains a (v, u) -path.

Theorem 4.17. *Let $\mathcal{G} \sim \mathcal{F}_{n,p}$. Then a.a.s. for every ordered pair of vertices v, u and all $\tau_1, \tau_2 \in [0, p]$, $\tau_1 \leq \tau_2$, the temporal graph $\mathcal{G}_{[\tau_1, \tau_2]}$ either contains a foremost (resp. hindmost) (v, u) -path with fewer than $14 \log n$ edges or contains no temporal (v, u) -path.*

Remark 4.18. *It can also be shown that the foremost tree has a vertex of height at least $\log n/4$ a.a.s., and hence the foremost tree has logarithmic height a.a.s. Indeed, the expected number of vertices below this height is sublinear, thus by Markov's inequality there are vertices above this height a.a.s.*

5 Sharp thresholds for temporal graph properties

In this section, we apply the results obtained in Section 4.2 to establish sharp thresholds for Point-to-point Reachability, First Temporal Source, Temporal Source, and Temporal Connectivity properties. We recall that our general strategy for obtaining a sharp threshold p_0 for a certain property in the model $\mathcal{F}_{n,p}$ is to show that the property *does not* hold in a random temporal complete graph $\mathcal{F}_{n,1}$ before time p_0 , and holds after time p_0 a.a.s. Therefore, in the proofs, if we do not specify explicitly the model a graph under consideration comes from, we assume that it comes from $\mathcal{F}_{n,1}$.

5.1 Point-to-point Reachability

Recall that Point-to-point Reachability is the property that for a fixed pair of vertices there exists a temporal path from the first vertex to the second one. In this section we establish a sharp threshold for this property.

Theorem 5.1. *The function $\frac{\log n}{n}$ is a sharp threshold for Point-to-point Reachability. More specifically, for any sufficiently large n , for two fixed distinct vertices v and u in $\mathcal{G} \sim \mathcal{F}_{n,p}$*

- (i) *there is no temporal (v, u) -path in \mathcal{G} with probability at least $1 - \frac{6}{\log n}$, if $p \leq \frac{\log n}{n} - \varepsilon$;*
- and*

(ii) there is a temporal (v, u) -path in \mathcal{G} with probability at least $1 - \frac{6}{\log n}$, if $p \geq \frac{\log n}{n} + \varepsilon$,

where $\varepsilon := \frac{3(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.

Proof. Assume n to be sufficiently large. By Lemma 4.12 (ii) (with $z = y = 1$), if $p \leq \frac{\log n}{n} - \varepsilon$, then with probability at least $1 - \frac{5}{\log n}$ vertex v reaches fewer than $\frac{n}{\log n}$ vertices. Analogously, by Lemma 4.12 (iii) (with $z = y = 1$), if $p \geq \frac{\log n}{n} + \varepsilon$, then with probability at least $1 - \frac{5}{\log n}$ vertex v reaches all but $\frac{n}{\log n}$ vertices.

By symmetry, the probability of u to be among the first $\frac{n}{\log n}$ vertices to become reachable from v is equal to $\frac{n}{n \log n} = \frac{1}{\log n}$ and the same holds for the last $\frac{n}{\log n}$ vertices. Hence, by the union bound, with probability at least $1 - \frac{6}{\log n}$, vertex u is not reachable from v if $p \leq \frac{\log n}{n} - \varepsilon$. Analogously and with the same probability, vertex u is reachable from v if $p \geq \frac{\log n}{n} + \varepsilon$. □

5.2 Temporal Source and First Temporal Source

In this section, we first establish a sharp threshold for Temporal Source, i.e., for the property that a fixed vertex is a temporal source. Then, we show that, quite surprisingly, the same function happens to be a sharp threshold for the property of having at least one temporal source in the graph, i.e., for First Temporal Source.

Before proceeding we make the simple but subsequently useful observation that the probability of a vertex v of $\mathcal{F}_{n,p}$ being a temporal source is equal to the probability of v being a temporal sink. This follows from the fact that v is a temporal source in (G, λ) if and only if it is a temporal sink in (G, λ') , where $\lambda'(e) := 1 - \lambda(e)$ for every edge of G . Therefore, the next two theorems are proved for temporal source, but they are stated and valid also for temporal sink.

Theorem 5.2. *The function $\frac{2 \log n}{n}$ is a sharp threshold for Temporal Source (resp. Temporal Sink). More specifically, for any sufficiently large n , a fixed vertex v in $\mathcal{G} \sim \mathcal{F}_{n,p}$*

(i) *is not a temporal source (resp. temporal sink) with probability at least $1 - \frac{5}{\log n}$, if $p \leq \frac{2 \log n}{n} - \varepsilon$;*

(ii) *is a temporal source (resp. temporal sink) with probability at least $1 - \frac{5}{\log n}$, if $p \geq \frac{2 \log n}{n} + \varepsilon$,*

where $\varepsilon := \frac{3(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.

Proof. Follows directly from Lemma 4.12 (iii) and (iv) when setting $z := 0$. □

Theorem 5.3. *The function $\frac{2 \log n}{n}$ is a sharp threshold for First Temporal Source (resp. First Temporal Sink). More specifically, for any sufficiently large n , a random temporal graph $\mathcal{G} \sim \mathcal{F}_{n,p}$*

(i) *does not contain a temporal source (resp. temporal sink) with probability at least $1 - 2n^{-\sqrt{\log n}}$, if $p \leq \frac{2 \log n}{n} - \varepsilon$;*

(ii) *contains a temporal source (resp. temporal sink) with probability at least $1 - \frac{5}{\log n}$, if $p \geq \frac{2 \log n}{n} + \varepsilon$,*

where $\varepsilon := \frac{3(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.

Proof. We observe that the upper bound of $\frac{2\log n}{n} + \varepsilon$ on the threshold follows from Theorem 5.2, and we only need to show the lower bound. Recall that by definition we always have $\widehat{Y}_{n-1}^v \leq Y_{n-1}^v$. Therefore, by Corollary 4.9, with probability at least $1 - 2n^{-\sqrt{\log n}}$ the inequality $Y_{n-1}^v \geq \widehat{Y}_{n-1}^v > \frac{2\log n}{n} - \frac{3(\log n)^{0.8}}{n}$ holds for all vertices v . In other words, with probability at least $1 - 2n^{-\sqrt{\log n}}$ no vertex is a temporal source until time $\frac{2\log n}{n} - \frac{3(\log n)^{0.8}}{n}$. \square

5.3 Temporal Connectivity

In this section, we establish a sharp threshold for Temporal Connectivity.

Theorem 5.4. *The function $\frac{3\log n}{n}$ is a sharp threshold for Temporal Connectivity. More specifically, for any sufficiently large n , a random temporal graph in $\mathcal{F}_{n,p}$*

(i) *is not temporally connected a.a.s., if $p \leq \frac{3\log n}{n} - \varepsilon$, where $\varepsilon := \frac{6(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$;*

(ii) *is temporally connected a.a.s., if $p \geq \frac{3\log n}{n} + \varepsilon$, where $\varepsilon := \frac{3(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.*

We split the proof of Theorem 5.4 into two parts. Note that a temporal graph is temporally connected if and only if each of its vertices is a temporal sink. Our strategy is to show that on the one hand before time $\frac{3\log n}{n}$ at least one vertex in a random temporal complete graph $\mathcal{G} \sim \mathcal{F}_{n,1}$ is not a temporal sink a.a.s. (Lemma 5.5), and on the other hand after time $\frac{3\log n}{n}$ all vertices in \mathcal{G} are temporal sinks a.a.s. (Lemma 5.6).

Lemma 5.5. *Let $p \leq \frac{3\log n}{n} - \varepsilon$, where $\varepsilon := \frac{6(\log n)^{0.8}}{n}$, and let $\mathcal{G} \sim \mathcal{F}_{n,1}$. Then, a.a.s. the temporal graph $\mathcal{G}_{[0,p]}$ contains at least one vertex which is not a temporal sink.*

Proof. Let q be the time at which the first vertex in \mathcal{G} becomes a temporal sink, i.e.,

$$q := \min\{q' \mid \mathcal{G}_{[0,q']} \text{ has a temporal sink}\}.$$

By Theorem 5.3, $q \geq \frac{2\log n}{n} - \frac{3(\log n)^{0.8}}{n} = \frac{2\log n}{n} - \frac{\varepsilon}{2}$ a.a.s. Note that $\mathcal{G}_{[0,q]}$ has at most two temporal sinks, since the addition of a single edge can not turn more than two vertices into temporal sinks. Furthermore, in order for a vertex which is not a temporal sink at time q to become a temporal sink it has to acquire at least one edge incident to it after time q . Thus, to prove the lemma, we will show that the underlying graph H of $\mathcal{G}_{[q,p]}$ has at least 3 isolated vertices, and hence at least one of these vertices is not a temporal sink in $\mathcal{G}_{[0,p]}$.

There to, note that any pair of vertices forms an edge in H with probability $\gamma := \frac{p-q}{1-q}$, unless it is an edge in $\mathcal{G}_{[0,q]}$. Thus, graph H follows the same distribution as drawing a random graph from $G_{n,\gamma}$ and then deleting all edges contained in $\mathcal{G}_{[0,q]}$.

As $\gamma = \frac{p-q}{1-q}$ is maximized when q is minimal, we have for sufficiently large n that

$$\begin{aligned} \gamma &= \frac{p-q}{1-q} \leq \frac{3\frac{\log n}{n} - \varepsilon - 2\frac{\log n}{n} + \varepsilon/2}{1 - 2\frac{\log n}{n} + \varepsilon/2} < \frac{\frac{\log n}{n} - \frac{\varepsilon}{2}}{1 - 2\frac{\log n}{n}} \\ &= \frac{1}{n} \cdot \frac{\log n - 3(\log n)^{0.8}}{1 - 2\frac{\log n}{n}} = \frac{1}{n} \left(\log n - \frac{3(\log n)^{0.8} - 2\frac{(\log n)^2}{n}}{1 - 2\frac{\log n}{n}} \right) \\ &< \frac{1}{n} \left(\log n - \frac{2(\log n)^{0.8}}{1 - 2\frac{\log n}{n}} \right) < \frac{\log n - (\log n)^{0.8}}{n}. \end{aligned}$$

It is known that $G_{n,\gamma}$ contains a.a.s. more than two isolated vertices if $\lim_{n \rightarrow \infty} n(1 - \gamma)^{n-1} = \infty$ [11, Theorem 3.1(ii)]. In order to show the latter, we first evaluate

$$\tau := \log n + n \cdot \log \left(1 - \frac{\log n - (\log n)^{0.8}}{n} \right) = (\log n)^{0.8} - O \left(\frac{(\log n)^2}{n} \right) \xrightarrow{n \rightarrow \infty} \infty$$

by means of Maclaurin expansion. Now, using this, we derive

$$\begin{aligned} n(1 - \gamma)^{n-1} &\geq n \left(1 - \frac{\log n - (\log n)^{0.8}}{n} \right)^{n-1} \\ &= \exp \left(\log n + (n-1) \log \left(1 - \frac{\log n - (\log n)^{0.8}}{n} \right) \right) \\ &\geq e^\tau \xrightarrow{n \rightarrow \infty} \infty. \end{aligned}$$

Thus, $G_{n,\gamma}$, and therefore H , contains at least three isolated vertices a.a.s., as required. \square

Lemma 5.6. *Let $p \geq \frac{3 \log n}{n} + \varepsilon$, where $\varepsilon := \frac{3(\log n)^{0.8}}{n}$, and let $\mathcal{G} \sim \mathcal{F}_{n,1}$. Then, a.a.s. every vertex in $\mathcal{G}_{[0,p]}$ is a temporal sink.*

Proof. Let $q := \frac{2 \log n}{n} + \varepsilon$ and $r := \frac{n}{\log n} \log \log \log n$. It follows from Theorem 5.2 that the probability for an arbitrary vertex *not* to be temporal sink in $\mathcal{G}_{[0,q]}$ is at most $\frac{5}{\log n}$. Hence, the expected number of vertices that are not temporal sinks in $\mathcal{G}_{[0,q]}$ is at most $\frac{5n}{\log n}$, and therefore, by Markov's inequality, the probability of $\mathcal{G}_{[0,q]}$ having more than r vertices that are not temporal sinks is at most $\frac{5n}{\log n} \cdot \frac{\log n}{n \cdot \log \log \log n} = \frac{5}{\log \log \log n} \in o(1)$. We will show that a.a.s. no more than $\frac{\log n}{n}$ extra time is required for these at most r remaining vertices to become temporal sinks.

Let $S \subseteq V$ be the set of vertices that are temporal sinks in $\mathcal{G}_{[0,q]}$. Observe that if at some time $t > q$ a vertex $w \in V \setminus S$ acquires an edge that connects w with a temporal sink in S , then w becomes a temporal sink no later than time t (notice, that w can become a temporal sink before time t). Therefore, to prove the lemma, we will show that a.a.s. by time $p = q + \frac{\log n}{n}$ every vertex in $V \setminus S$ acquires such an edge.

First, note that the underlying graph $G_{[0,q]}$ of $\mathcal{G}_{[0,q]}$ is distributed according to $G_{n,q}$. Hence, we may apply [11, Corollary 3.13] which states that the maximum degree $\Delta(G_{[0,q]})$ of $G_{[0,q]}$ a.a.s. satisfies

$$\begin{aligned} \left| \Delta(G_{[0,q]}) - qn - \sqrt{2q(1-q)n \cdot \log n} + \log \log n \sqrt{\frac{q(1-q)n}{8 \log n}} + \log(2\sqrt{\pi}) \sqrt{\frac{q(1-q)n}{2 \log n}} \right| \\ \leq \frac{\log n}{2} \sqrt{\frac{q(1-q)n}{\log n}}. \end{aligned}$$

Thus, a.a.s. we have

$$\begin{aligned} \Delta(G_{[0,q]}) &\leq qn + \sqrt{2q(1-q)n \cdot \log n} + \frac{\log n}{2} \sqrt{\frac{q(1-q)n}{\log n}} \\ &< 3 \log n + \sqrt{6} \log n + \frac{\sqrt{3}}{2} \log n < 7 \log n \end{aligned}$$

where we used that $\varepsilon < \frac{\log n}{n}$ for sufficiently large n . For every vertex $w \in V \setminus S$ define

$$C_w := \{vw \mid v \in S \text{ and } \lambda(vw) > q\}$$

as the set of edges connecting w to S after time q . By the above, a.a.s. for every w , we have that $|C_w| \geq n - 1 - r - 7 \log(n) \geq n - 2r =: d$. Let T_w be the *waiting time* for the first of these edges to appear, i.e., $T_w := \min\{\lambda(vw) - q \mid vw \in C_w\}$. Note that the time labels of the edges in C_w are independently and uniformly distributed on the interval $I = (q, 1]$. Thus, we have $\mathbb{P}[T_w > x] \leq (1 - x)^d$ for any $0 \leq x \leq 1 - q$. In particular,

$$\begin{aligned} \tau := \mathbb{P}\left[T_w > \frac{\log n}{n}\right] &\leq \left(1 - \frac{\log n}{n}\right)^d \\ &= \left(1 - \frac{\log n}{n}\right)^{\frac{n}{\log n}(\log n - 2 \log \log \log n)} \\ &\leq e^{-(\log n - 2 \log \log \log n)} \\ &= \frac{(\log \log n)^2}{n} < 1. \end{aligned}$$

Now let T be the waiting time of the last vertex in $V \setminus S$ to become adjacent to at least one of the temporal sinks in S , i.e., $T = \max\{T_w \mid w \in V \setminus S\}$. Because the waiting times $T_w, w \in V \setminus S$ are all independent, we have

$$\mathbb{P}[T > x] = 1 - \prod_w \mathbb{P}[T_w < x] \leq 1 - (1 - (1 - x)^d)^r.$$

In the rest of the proof we show that $\lim_{n \rightarrow \infty} \mathbb{P}[T > \frac{\log n}{n}] = 0$. For this purpose, we derive

$$\begin{aligned} 0 &\geq r \cdot \log(1 - \tau) \geq r \cdot \log\left(1 - \frac{(\log \log n)^2}{n}\right) \\ &\geq -2r \frac{(\log \log n)^2}{n} \quad [\text{since } (\log \log n)^2/n < 0.7, \text{ assuming } n \geq 2] \\ &= -2 \frac{(\log \log n)^2 \cdot \log \log \log n}{\log n} \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

Thus, we have $\lim_{n \rightarrow \infty} r \cdot \log(1 - \tau) = 0$ and therefore

$$\mathbb{P}\left[T > \frac{\log n}{n}\right] \leq 1 - (1 - \tau)^r = 1 - \exp(r \cdot \log(1 - \tau)) \xrightarrow{n \rightarrow \infty} 1 - e^0 = 0.$$

□

5.4 Temporal Spanners

The results from Section 5.2 can be directly used to show that $\frac{4 \log n}{n}$ is a sharp threshold for the existence of a pivotal spanner. Indeed, let $\mathcal{G} \sim \mathcal{F}_{n,p}$ for $p \geq \frac{4 \log n}{n} + \frac{6(\log n)^{0.8}}{n}$, and let v be an arbitrary vertex in \mathcal{G} . Then, by Theorem 5.2, with probability at least $1 - 5/\log n$ vertex v is a temporal sink in $\mathcal{G}_{[0,p_1]}$, and with at least the same probability v is a temporal source in $\mathcal{G}_{[p_1,p]}$, where $p_1 = \frac{2 \log n}{n} + \frac{3(\log n)^{0.8}}{n}$. Hence, by the union bound, with probability at least $1 - 10/\log n$ temporal graph \mathcal{G} contains a pivotal temporal spanner formed as the union of the hindmost tree for v in $\mathcal{G}_{[0,p_1]}$ and the foremost tree for v in $\mathcal{G}_{[p_1,p]}$. On the other hand, if $p \leq \frac{4 \log n}{n} - \frac{6(\log n)^{0.8}}{n}$, then by Theorem 5.3 for any $p_1 \in [0, p]$ with probability at least $1 - \frac{5}{\log n}$ either $\mathcal{G}_{[0,p_1]}$ has no temporal sink, or $\mathcal{G}_{[p_1,p]}$ has no temporal source (or both), and therefore \mathcal{G} contains no pivotal temporal spanner. Consequently, $\frac{4 \log n}{n}$ is a sharp threshold for the emergence of a Pivotal Temporal Spanner.

Theorem 5.7. *The function $\frac{4\log n}{n}$ is a sharp threshold for the emergence of a Pivotal Temporal Spanner. More specifically, for any sufficiently large n , a random temporal graph in $\mathcal{F}_{n,p}$*

(i) *has no pivotal temporal spanner with probability at least $1 - \frac{5}{\log n}$, if $p \leq \frac{4\log n}{n} - \varepsilon$;*
and

(ii) *has a pivotal temporal spanner with probability at least $1 - \frac{10}{\log n}$, if $p \geq \frac{4\log n}{n} + 2\varepsilon$,*

where $\varepsilon := \frac{3(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.

While a pivotal temporal spanner contains only 2 more edges than an optimal spanner, the sharp threshold for the existence of a pivotal temporal spanner is quite far from the temporal connectivity threshold. This naturally raises a number of questions. For example, when does an optimal spanner appear? Or, can we have a temporal spanner with a linear number of edges a.a.s. when p is close to the connectivity threshold? We partially address the former question in Section 5.4.3 by showing that an optimal spanner exists a.a.s. at $p = \frac{4\log n}{n}$; furthermore, we conjecture that $\frac{4\log n}{n}$ is a sharp threshold for the Optimal Temporal Spanner property. In Section 5.4.2, we provide a strong answer to the latter question, by proving that $\frac{3\log n}{n}$ is a sharp threshold for the existence of a nearly optimal temporal spanner, i.e., a spanner with $(2 + o(1))n$ time-edges.

The proofs of both results are different elaborations of the above simple construction of a pivotal temporal spanner. The construction of a nearly optimal spanner is more involved among the two, and we develop it iteratively, by first presenting in the next section a simpler construction of a temporal spanner with at most $5n$ edges at $p = 3.5\frac{\log n}{n}$.

5.4.1 Warm-up: a spanner with $5n$ edges at $3.5\frac{\log n}{n}$

In this section we show that a random temporal graph from $\mathcal{F}_{n,p}$ a.a.s. contains a temporal spanner with at most $5n$ edges if $p \geq (3.5 + o(1))\frac{\log n}{n}$. The proof relies upon a construction of a pair of sparse subgraphs. Each of these subgraphs consists of two edge-disjoint trees with a common root, such that every vertex in one of the trees can reach the root before some time t , while every vertex in the other tree can be reached from the root after t . Thus all vertices of the first tree can reach all vertices of the second tree through the root. The two subgraphs together provide temporal connectivity for all but a sublinear number of pairs of vertices. Each of the remaining pairs is connected directly by a foremost path.

Theorem 5.8. *A random temporal graph from $\mathcal{F}_{n,p}$ a.a.s. has a spanner with at most $(4 + o(1))n$ edges if $p \geq 3.5\frac{\log n}{n} + \frac{6(\log n)^{0.8}}{n}$.*

Proof. Let $\mathcal{G} \sim \mathcal{F}_{n,1}$. Set $\varepsilon := \frac{3(\log n)^{0.8}}{n}$, and

$$r := 1.5\frac{\log n}{n} + \varepsilon, \quad q := 2\frac{\log n}{n} + \varepsilon.$$

Let u and v be two arbitrary vertices. By Theorem 5.2 and Lemma 4.12 (iii) (with $z = 0.5$, $y = 1$), we have that each of the following events happens with probability at least $1 - 5/\log n$

1. u is a temporal sink in $\mathcal{G}_{[0,q]}$; we denote by T_u the hindmost tree for u in $\mathcal{G}_{[0,q]}$.
2. v is a temporal source in $\mathcal{G}_{[r,p]}$; we denote by T_v the foremost tree for v in $\mathcal{G}_{[r,p]}$.

3. u reaches at least $n - \frac{\sqrt{n}}{\log n}$ vertices in $\mathcal{G}_{[q,p]}$; we denote by T'_u the foremost tree for u in $\mathcal{G}_{[q,p]}$.
4. v is reached by at least $n - \frac{\sqrt{n}}{\log n}$ vertices in $\mathcal{G}_{[0,r]}$; we denote by T'_v the hindmost tree for v in $\mathcal{G}_{[0,r]}$.

Clearly, these four trees together have at most $4(n-1)$ edges. Let A be the set of vertices *not* reaching v through T'_v , and analogously let B be the set of vertices *not* reached by u through T'_u . Since all edges of T_u appear before time q while all edges of T'_u appear later, $T_u \cup T'_u$ provides reachability for all ordered vertex pairs in $V \times (V \setminus B)$. Analogously, $T'_v \cup T_v$ provides reachability for all pairs in $(V \setminus A) \times V$. Thus it remains to provide reachability only for the pairs in $A \times B$.

By Theorem 5.4, $\mathcal{G}_{[0,p]}$ is temporally connected a.a.s. Hence, by Theorem 4.17, a.a.s. for every ordered pair $(a, b) \in A \times B$ there exists a foremost (a, b) -path in $\mathcal{G}_{[0,p]}$ with fewer than $14 \log n$ edges. Therefore, in total these paths have at most $|A \times B| \cdot 14 \log n \leq \frac{n}{(\log n)^2} \cdot 14 \log n = \frac{14n}{\log n}$ edges. This collection of paths together with the four trees form a temporal spanner with at most $4(n-1) + \frac{14n}{\log n} = (4 + o(1))n$ edges in $\mathcal{G}_{[0,p]}$ a.a.s. \square

5.4.2 Sharp threshold for nearly optimal temporal spanners

In this section we show that at time $3 \frac{\log n}{n}$, which is when temporal connectivity emerges (Theorem 5.4), a graph contains a temporal spanner of nearly optimal size $(2 + o(1))n$ a.a.s. To achieve this, we extend and refine the approach used in the previous section. In particular, we give a more subtle construction of the two sparse subgraphs (Lemma 5.9 and Lemma 5.10) and conduct a more involved accompanying analysis (Lemma 5.11) to show that the two subgraphs provide temporal connectivity for all but a sublinear number of pairs by time $(3 + o(1)) \frac{\log n}{n}$. As before, any leftover pairs are connected via direct foremost paths.

Let \mathcal{G} be a temporal graph and let u and v be two vertices in \mathcal{G} . The *earliest arrival time* to v from u in \mathcal{G} , denoted by $\mu_{\mathcal{G}}(u, v)$, is the minimum time t such that there exists a temporal (u, v) -path in \mathcal{G} with arrival time t . Similarly, the *latest departure time* from u to v in \mathcal{G} , denoted by $\sigma_{\mathcal{G}}(u, v)$, is the maximum time t such that there exists a temporal (u, v) -path in \mathcal{G} with departure time t . If \mathcal{G} has no temporal (u, v) -path, we set $\mu_{\mathcal{G}}(u, v) = \infty$ and $\sigma_{\mathcal{G}}(u, v) = -\infty$. For notational convenience we define $\infty - \infty := 0$. We proceed with the lemma showing that if p is large enough, then for any fixed vertex u in $\mathcal{G} \sim \mathcal{F}_{n,p}$ there is a.a.s. a sparse temporal subgraph in which the earliest arrival time from any vertex v to u is almost the same as in \mathcal{G} .

Lemma 5.9. *Let $\delta_0 := \frac{28}{\log n}$ and $\varepsilon_0 := \frac{6}{(\log n)^{0.2}}$. Then for any $p \geq (1 + \varepsilon_0) \frac{\log n}{n}$ and any fixed vertex u of $\mathcal{G} \sim \mathcal{F}_{n,p}$, a.a.s. there is a temporal subgraph \mathcal{G}' of \mathcal{G} with at most $(1 + \delta_0)n$ edges so that*

$$|\mu_{\mathcal{G}}(v, u) - \mu_{\mathcal{G}'}(v, u)| \leq \varepsilon_0 \frac{\log n}{n}$$

holds for all vertices $v \in V(\mathcal{G})$.

Proof. Let u be an arbitrary vertex in \mathcal{G} . Set $q_+ := (1 + \varepsilon_0/2) \frac{\log n}{n}$ and $q_- := (1 - \varepsilon_0/2) \frac{\log n}{n}$. By Lemma 4.12 (iii) and (ii) (with $z = 1$, $y = 2$), vertex u is reachable by all but at most $\frac{n}{(\log n)^2}$ vertices in \mathcal{F}_{n,q_+} a.a.s., and vertex u is reachable by at most $\frac{n}{(\log n)^2}$ vertices in \mathcal{F}_{n,q_-} a.a.s. We construct \mathcal{G}' as the union of

- (1) the hindmost tree for u in $\mathcal{G}_{[0, q_+]}$, and
- (2) an earliest-arrival hindmost (v, u) -path in \mathcal{G} of minimal length for each vertex v with $\mu_{\mathcal{G}}(v, u) \notin [q_-, q_+]$.

Observe that $\mu_{\mathcal{G}}(v, u) \leq \mu_{\mathcal{G}'}(v, u) \leq q_+$ for all vertices v with $q_- \leq \mu_{\mathcal{G}}(v, u) \leq q_+$ and that $\mu_{\mathcal{G}}(v, u) = \mu_{\mathcal{G}'}(v, u)$ for all other vertices v . Thus $|\mu_{\mathcal{G}}(v, u) - \mu_{\mathcal{G}'}(v, u)| \leq \varepsilon_0 \frac{\log n}{n}$ holds for all vertices of \mathcal{G} .

To estimate the size of \mathcal{G}' , we observe that the hindmost tree in (1) contains at most $n - 1$ edges, and by the choice of q_- and q_+ , the number of paths added in (2) is a.a.s. at most $\frac{2n}{(\log n)^2}$. Furthermore, by Theorem 4.17, a.a.s. each of these paths has length at most $14 \log n$. Thus, a.a.s. \mathcal{G}' has at most $n - 1 + \frac{2n}{(\log n)^2} \cdot 14 \log n \leq (1 + \delta_0)n$ edges, as required. \square

The following counterpart of Lemma 5.9 holds by symmetry under time reversal and shows that, for any fixed vertex u , a.a.s. \mathcal{G} contains a sparse temporal subgraph in which the latest departure times from u to any vertex v is almost the same as in \mathcal{G} .

Lemma 5.10. *Let $\delta_0 := \frac{28}{\log n}$ and $\varepsilon_0 := \frac{6}{(\log n)^{0.2}}$. Then for any $p \geq (1 + \varepsilon_0) \frac{\log n}{n}$ and any fixed vertex u of $\mathcal{G} \sim \mathcal{F}_{n,p}$, a.a.s. there is a temporal subgraph \mathcal{G}' of \mathcal{G} with at most $(1 + \delta_0)n$ edges so that*

$$|\sigma_{\mathcal{G}}(u, v) - \sigma_{\mathcal{G}'}(u, v)| \leq \varepsilon_0 \frac{\log n}{n}$$

holds for all vertices $v \in V(\mathcal{G})$.

In the next lemma we combine Lemmas 5.9 and 5.10 to construct a sparse spanner in $\mathcal{G} \sim \mathcal{F}_{n,p}$, where $p = (3 + o(1)) \frac{\log n}{n}$; the spanner provides temporal connectivity for all but at most a sublinear number of pairs of vertices.

Lemma 5.11. *Let $\delta := \frac{56}{\log n}$. There exists a function $\varepsilon = \varepsilon(n) \in o(1)$ such that for $p = (3 + \varepsilon) \frac{\log n}{n}$ a.a.s. a temporal graph $\mathcal{G} \sim \mathcal{F}_{n,p}$ contains a temporal subgraph \mathcal{G}' with at most $(2 + \delta)n$ edges and there are at most $\frac{n \log \log n}{(\log n)^2}$ ordered pairs of vertices (v, u) for which there is no temporal (v, u) -path in \mathcal{G}' .*

Proof. Let $\varepsilon_0 := \frac{6}{(\log n)^{0.2}}$. We assume that n is large enough so that $\varepsilon_0 < 1$. By Theorem 5.2 there is $\varepsilon_1 = \varepsilon_1(n) \in o(1)$ such that an arbitrary vertex of $\mathcal{F}_{n, (2 + \varepsilon_1) \frac{\log n}{n}}$ is a temporal source a.a.s. If necessary, we increase ε_1 to ensure $(\varepsilon_1 - \varepsilon_0) \log n > 3(\log n)^{0.8}$. Set $r := \lfloor \log \log n \rfloor$, $\varepsilon := 2\varepsilon_1 + 2/r$, and

$$p_1 := (2 + \varepsilon_1 + 1/r) \frac{\log n}{n}$$

$$p_2 := p - p_1 = (1 + \varepsilon - \varepsilon_1 - 1/r) \frac{\log n}{n} = (1 + \varepsilon_1 + 1/r) \frac{\log n}{n} = p_1 - \frac{\log n}{n}$$

Fix an arbitrary vertex $u \in \mathcal{G}$. Note that, by choice of ε_1 and as $p_1 > (2 + \varepsilon_1) \frac{\log n}{n}$, vertex u is a.a.s. temporal sink in $\mathcal{G}_{[0, p_1]}$ and temporal source in $\mathcal{G}_{[p_2, p]}$.

Define $\delta_0 := \frac{28}{\log n}$. By Lemma 5.9, there exists a.a.s. a temporal subgraph \mathcal{G}_1 of $\mathcal{G}_{[0, p_1]}$ of at most $(1 + \delta_0)n$ edges such that for all vertices $v \in V(\mathcal{G})$

$$|\mu_{\mathcal{G}}(v, u) - \mu_{\mathcal{G}_1}(v, u)| \leq \varepsilon_0 \frac{\log n}{n}.$$

In particular, u is a temporal sink in \mathcal{G}_1 and thus every vertex v has $\mu_{\mathcal{G}_1}(v, u) \leq p_1$. Symmetrically, by Lemma 5.10, there exists a.a.s. a temporal subgraph \mathcal{G}_2 of $\mathcal{G}_{[p_2, p]}$ of at most $(1 + \delta_0)n$ edges such that for all vertices $w \in V(\mathcal{G})$

$$|\sigma_{\mathcal{G}}(u, w) - \sigma_{\mathcal{G}_2}(u, w)| \leq \varepsilon_0 \frac{\log n}{n}$$

and in particular $\sigma_{\mathcal{G}_2}(u, w) \geq p_2$ for every vertex w , because u is a temporal source in $\mathcal{G}_{[p_2, p]}$.

We define \mathcal{G}' as the union $\mathcal{G}_1 \cup \mathcal{G}_2$ and observe that \mathcal{G}' has at most $(2 + 2\delta_0)n = (2 + \delta)n$ edges. Note that there exists a temporal (v, w) -path in $\mathcal{G}_1 \cup \mathcal{G}_2$ if $\mu_{\mathcal{G}_1}(v, u) < \sigma_{\mathcal{G}_2}(u, w)$. Otherwise, i.e., if $\mu_{\mathcal{G}_1}(v, u) \geq \sigma_{\mathcal{G}_2}(u, w)$, we call the ordered pair (v, w) *mismatched*. We claim that the total number of mismatched pairs, and therefore the number of pairs of vertices that are not temporally connected in \mathcal{G}' , is at most $r \cdot \frac{n}{(\log n)^2} \leq \frac{n \log \log n}{(\log n)^2}$. To this end, we introduce, for $0 \leq i \leq r$,

$$\begin{aligned} q_i &:= p_2 + i \cdot \frac{p_1 - p_2}{r} = p_2 + \frac{i}{r} \cdot \frac{\log n}{n} \\ A_i &:= \{v \in V(\mathcal{G}) \mid \mu_{\mathcal{G}_1}(v, u) \geq q_i\} \\ B_i &:= \{w \in V(\mathcal{G}) \mid \sigma_{\mathcal{G}_2}(u, w) \leq q_i\} \end{aligned}$$

and, for $0 \leq i < r$,

$$C_i := A_i \times B_{i+1}.$$

Note that each mismatched pair is contained in some C_i . Indeed, a mismatched pair (v, w) belongs to C_i for the maximum $i \in \{0, 1, \dots, r-1\}$ such that $q_i \leq \mu_{\mathcal{G}_1}(v, u)$. For every vertex $v \in A_i$, $0 \leq i < r$, we have

$$\begin{aligned} \mu_{\mathcal{G}}(v, u) &\geq \mu_{\mathcal{G}_1}(v, u) - \varepsilon_0 \frac{\log n}{n} \\ &\geq q_i - \varepsilon_0 \frac{\log n}{n} \\ &= p_2 + \left(\frac{i}{r} - \varepsilon_0\right) \frac{\log n}{n} \\ &= \left(1 + \varepsilon_1 + \frac{1}{r} - \varepsilon_0 + \frac{i}{r}\right) \cdot \frac{\log n}{n} \\ &= \left(1 + \varepsilon_1 - \varepsilon_0 + \frac{i+1}{r}\right) \cdot \frac{\log n}{n} \\ &> \left(1 + \frac{i+1}{r}\right) \cdot \frac{\log n}{n} + \frac{3(\log n)^{0.8}}{n} =: a_i. \end{aligned}$$

On the other hand, by Lemma 4.12 (iii) (with $z = 1 - (i+1)/r$ and $y = 1$), with probability at least $1 - 5/\log n$, at least $n - \left\lfloor \frac{n^{1-(i+1)/r}}{\log n} \right\rfloor$ vertices in $\mathcal{G}_{[0, a_i]}$ can reach u , and therefore $\mu_{\mathcal{G}}(v, u) \leq a_i$ holds for all those vertices. Consequently, by the union bound, with probability at least $1 - 5r/\log n$, we have $|A_i| \leq \frac{n^{1-(i+1)/r}}{\log n}$ for every $0 \leq i < r$.

Similarly, for every vertex $w \in B_i$, $0 < i \leq r$, we have

$$\begin{aligned}
\sigma_{\mathcal{G}}(u, w) &\leq \sigma_{\mathcal{G}_2}(u, w) + \varepsilon_0 \frac{\log n}{n} \\
&\leq q_i + \varepsilon_0 \frac{\log n}{n} \\
&= p_2 + \left(\frac{i}{r} + \varepsilon_0\right) \frac{\log n}{n} \\
&= p - \left(2 + \varepsilon_1 + \frac{1}{r} - \varepsilon_0 - \frac{i}{r}\right) \frac{\log n}{n} \\
&< p - \left(\left(2 - \frac{i-1}{r}\right) \frac{\log n}{n} + \frac{3(\log n)^{0.8}}{n}\right) =: p - b_i.
\end{aligned}$$

On the other hand, by Lemma 4.12 (iii) (with $z = (i-1)/r$ and $y = 1$), with probability at least $1 - 5/\log n$, at least $n - \lfloor \frac{n^{(i-1)/r}}{\log n} \rfloor$ vertices in $\mathcal{G}_{[p-b_i, p]}$ can be reached from u , and therefore $\sigma_{\mathcal{G}}(u, w) \geq p - b_i$ holds for all those vertices. Consequently, by the union bound, with probability at least $1 - 5r/\log n$, we have $|B_i| \leq \frac{n^{(i-1)/r}}{\log n}$ for every $0 < i \leq r$.

Thus a.a.s. for all $0 \leq i < r$

$$|C_i| = |A_i| \times |B_{i+1}| \leq \frac{n^{1-(i+1)/r}}{\log n} \cdot \frac{n^{i/r}}{\log n} = \frac{n^{1-1/r}}{(\log n)^2} < \frac{n}{(\log n)^2},$$

and hence, as claimed, overall there are at most $\sum_{i=0}^{r-1} |C_i| \leq r \cdot \frac{n}{(\log n)^2}$ mismatched pairs. \square

We are now ready to state and prove our main result.

Theorem 5.12. *Let $\delta := \frac{56}{\log n} + \frac{14 \log \log n}{\log n} \in o(1)$. Then $p = 3 \frac{\log n}{n}$ is a sharp threshold for a random temporal graph $\mathcal{G} \sim \mathcal{F}_{n,p}$ to have a spanner of size at most $(2 + \delta)n$.*

Proof. Let $\varepsilon'(n)$ be given by Lemma 5.11, and $\varepsilon''(n)$ be given by Theorem 5.4. Set $\varepsilon := \varepsilon'(n) + \varepsilon''(n) \in o(1)$. If $p \leq (3 - \varepsilon) \frac{\log n}{n}$, then by Theorem 5.4 a.a.s. a random temporal graph $\mathcal{G} \sim \mathcal{F}_{n,p}$ is not temporally connected, and therefore it has no spanner at all. Hence, to prove the theorem, we show in the rest of the proof that if $p \geq (3 + \varepsilon) \frac{\log n}{n}$, then a.a.s. $\mathcal{G} \sim \mathcal{F}_{n,p}$ has a spanner of size at most $(2 + \delta)n$.

Let \mathcal{G}' be a subgraph of \mathcal{G} with at most $\left(2 + \frac{56}{\log n}\right)n$ edges as guaranteed by Lemma 5.11. Then a.a.s. \mathcal{G}' contains at most $\frac{n \log \log n}{(\log n)^2}$ ordered pairs (v, w) with no temporal (v, w) -path. Since \mathcal{G} is a.a.s. temporally connected by Theorem 5.4, it a.a.s. contains a temporal (v, w) -path for each of these pairs (v, w) , and by Theorem 4.17 a.a.s. each of these temporal paths can be chosen to have length at most $14 \log n$. By adding these paths to \mathcal{G}' , we obtain a temporal spanner of size at most

$$\left(2 + \frac{56}{\log \log n}\right)n + \frac{n \log \log n}{(\log n)^2} \cdot 14 \log n = (2 + \delta)n.$$

\square

5.4.3 Optimal Temporal Spanners

In this section we show that a random temporal graph from $\mathcal{F}_{n,p}$ contains an optimal spanner (i.e., a spanner with exactly $2n - 4$ edges) a.a.s. whenever $p \geq (4 + o(1)) \frac{\log n}{n}$. The main idea of the construction is to replace the pivot vertex in the pivotal temporal spanner construction by a pivot 4-cycle.

Theorem 5.13. *A random temporal graph in $\mathcal{F}_{n,p}$ a.a.s. has an optimal spanner if $p \geq 4\frac{\log n}{n} + \varepsilon$ where $\varepsilon := \frac{16(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.*

Proof. Assume $n \geq 8$ and set $\varepsilon_0 := \frac{\varepsilon}{4} = \frac{4(\log n)^{0.8}}{n} > \frac{3(\log(n-3))^{0.8}}{n-3} > \frac{\log \log n}{n}$. Let also $\mathcal{G} \sim \mathcal{F}_{n,1}$. Our goal is to show that $\mathcal{G}_{[0,p]}$ has an optimal spanner a.a.s. For this purpose, set

$$\begin{aligned} p_1 &:= 2\frac{\log n}{n} + \varepsilon_0, \\ p_2 &:= 2\frac{\log n}{n} + 2\varepsilon_0, \\ p_3 &:= 2\frac{\log n}{n} + 3\varepsilon_0. \end{aligned}$$

Define $V^{(4)}$ to be the set of all 4-tuples of pairwise distinct vertices of \mathcal{G} . We say that such a tuple $(w, x, y, z) \in V^{(4)}$ of \mathcal{G} forms a *square* if $\{\lambda(wx), \lambda(yz)\} \subset [p_1, p_2]$ and $\{\lambda(xy), \lambda(wz)\} \subset [p_2, p_3]$. Note that in a square each vertex can reach any other by a temporal path with labels from $[p_1, p_3]$. Our first goal is to show that a.a.s. there exist sufficiently many squares. This will guarantee that at least one of them can be chosen as a pivot 4-cycle.

Let S be the number of squares in \mathcal{G} . Clearly $\mathbb{E}[S] = n^{(4)}\varepsilon_0^4$, where $n^{(i)} := n \cdot (n-1) \cdot (n-2) \cdots (n-i+1)$ denotes the falling factorial. To bound the variance of S , we need to investigate

$$\mathbb{E}[S^2] = \sum_{A \in V^{(4)}} \sum_{B \in V^{(4)}} \mathbb{P}[A, B \text{ are both squares}] = \sum_{i=4}^8 \sum_{\substack{A, B \in V^{(4)} \\ |A \cup B|=i}} \mathbb{P}[A, B \text{ are both squares}].$$

For $i = 8$ (i.e., A, B being vertex-disjoint) we clearly have $n^{(8)}$ summands, whereas for $i \leq 7$ the number of pairs $A, B \in V^{(4)}$ with $|A \cup B| = i$ can be upper-bounded by $(8n)^i \leq 8^7 n^i$. The time label of every edge in $A \cup B$ belongs to an interval of length ε_0 and thus $\mathbb{P}[A, B \text{ are both squares}] \leq \varepsilon_0^i$. Therefore,

$$\mathbb{E}[S^2] \leq n^{(8)}\varepsilon_0^8 + \sum_{i=4}^7 8^7 n^i \varepsilon_0^i \leq \mathbb{E}[S]^2 + 8^7 \cdot \sum_{i=4}^7 4^i (\log n)^{0.8i} \leq \mathbb{E}[S]^2 + 8^7 \cdot 4 \cdot 4^7 (\log n)^{5.6}$$

and thus the variance $\mathbb{D}[S]$ is no more than $8^7 \cdot 4^8 (\log n)^{5.6} \in o(\mathbb{E}[S]^2)$. By Chebyshev's inequality, we obtain that $S \geq \mathbb{E}[S]/2$ a.a.s.

Assume now a 4-tuple (w, x, y, z) of vertices to be given (which might or might not form a square). Let \mathcal{G}' be the temporal subgraph of \mathcal{G} obtained by deleting the vertices x, y, z . Observe that \mathcal{G}' is distributed as an element of $\mathcal{F}_{n-3,1}$, independently of whether (w, x, y, z) form a square in \mathcal{G} . Note that the intervals $[0, p_1]$ and $[p_3, p]$ both have length $\frac{2\log n}{n} + \varepsilon_0$. Thus, by Theorem 5.2 and the union bound, with probability at least $1 - \frac{10}{\log(n-3)}$, vertex w is a temporal sink of $\mathcal{G}'_{[0,p_1]}$ and also a temporal source in $\mathcal{G}'_{[p_3,p]}$. If (w, x, y, z) additionally form a square in \mathcal{G} , then we call (w, x, y, z) a *good square*. Then, by taking the good square (w, x, y, z) together with the foremost tree for w in $\mathcal{G}'_{[p_3,p]}$ and the hindmost tree for w in $\mathcal{G}'_{[0,p_1]}$, we obtain a spanner of $\mathcal{G}_{[0,p]}$ that has $4 + 2(n-4) = 2n - 4$ edges.

It only remains to show that there exists a good square. According to the above, the expected number of bad (i.e., non-good) squares S_{bad} satisfies $\mathbb{E}[S_{\text{bad}}] \leq \mathbb{E}[S] \frac{10}{\log(n-3)}$. By

Markov's inequality and our bound on S , the probability for every square of \mathcal{G} to be bad is at most

$$\mathbb{P}[S_{\text{bad}} = S] \leq \mathbb{P}\left[S < \frac{\mathbb{E}[S]}{2}\right] + \frac{\mathbb{E}[S_{\text{bad}}]}{(\mathbb{E}S)/2} \leq o(1) + \frac{20}{\log(n-3)} \in o(1).$$

Thus, \mathcal{G} has a good square a.a.s. □

We conjecture that $4\frac{\log n}{n}$ is in fact a sharp threshold for the existence of an optimal spanner.

Conjecture 2. *The sharp threshold for the existence of an optimal spanner in a random simple temporal graph exists and is equal to $4\frac{\log n}{n}$.*

6 Connections to other models

In this section, we explore some connections to other models in the literature. First and most importantly, we show how our results translate to gossiping and population protocols, with or without repetition of interactions between pairs of nodes. Afterwards, we look into quite surprising parallels to known results about randomly weighted static graphs. Finally, we mention existing results in random edge-ordered graphs, where the model is equivalent to RSTGs but the questions that have been studied are different.

6.1 Gossiping and population protocols

In the classical gossiping setting, there are n agents, each of whom knows a single secret. The agents can communicate via telephone calls. Whenever an agent calls another agent, the two exchange all the secrets they know. An agent who learns all n secrets becomes an *expert*. The standard goal is to reach a configuration in which all agents are experts. A sequence of calls leading to this state is called a *gossiping protocol*³.

In the theory of population protocols, a similar process appears in the analysis of time complexity. Here the order of interactions is governed by a sequential *scheduler*. Unless this scheduler is adversarial, it is generally assumed, for analysis, that the interactions between pairs of nodes are chosen uniformly at random, with repetitions. The complexity of a protocol is then measured in terms of the number of interactions required to reach a certain configuration, which depends on the problem at hand.

If we consider the population protocol's interactions to be symmetric (i.e., not oriented), then the above two models are equivalent in terms of interactions. For the sake of concreteness, in what follows, we will rely on the terminology of gossiping protocols (in particular [27]). Two natural models of random interactions can be defined, depending on whether repetitions are allowed between the same pair of agents:

ANY: No restrictions apply; every call is performed uniformly at random, without dependency on previous calls;

CO: Stands for *call-once*, i.e., no call can be repeated; there is at most one call between any pair of agents.

³There are related models of *parallel rumor spreading* in which multiple pairwise calls can occur in a single round. In those models, the target measure is usually the number of rounds needed for secrets spreading (see, e.g., [33, 40]).

According to a classical result from the '70s, gossiping protocols require at least $2n - 4$ calls even deterministically (see [12] for a historical note on the corresponding combinatorial problem). On the probabilistic side, the duration of gossiping protocols has been studied mostly in the ANY model. In a sequence of three papers [49, 13, 37] with the same title “*Random exchanges of information*”, asymptotics for the *expected* number of calls until a fixed agent becomes an expert and until all agents become experts were obtained for the randomized ANY model of communication, i.e., for the model where every next call happens between a pair of distinct agents chosen uniformly at random among all pairs of distinct agents.

Theorem 6.1 ([49, 13, 37]). *In the randomized ANY model of communication,*

1. *the expected number of calls until a fixed agent becomes an expert is $n \log n + O(n)$;*
2. *the expected number of calls until all agents become experts is $\frac{3}{2}n \log n + O(n)$.*

Unsurprisingly, the number of calls until a fixed agent becomes an expert also turns out to be concentrated around its expected value. To our knowledge, the best concentration to date was obtained in [48]. We extract only the following:

Theorem 6.2 ([48, 47]). *In the randomized ANY model of communication, a.a.s. the number of calls until a fixed agent becomes an expert is $(1 + o(1))(n \log n)$.*

As for the number of calls until all agents become experts, it is known that it does not exceed its expected value a.a.s.. The best bound to date appears to be from Burman et al. [17]. We remark that the actual lemma in [17] states a weaker bound on time (to achieve a stronger bound on the probability), but an intermediate step of the proof directly implies the following:

Theorem 6.3 ([17], proof of Lemma 2.9). *In the randomized ANY model of communication, a.a.s. the number of calls until all agents become experts is at most $(1 + o(1))(\frac{3}{2}n \log n)$ a.a.s.*

According to Ditmarsch, Kokkinis, and Stockmarr [27], there seem to be no similar estimations known for the randomized CO model, which may arguably be due to the fact that the no-repetition property of this model creates dependencies between past and future events which are more difficult to handle. In Section 6.1.1 (below), we show how our results readily translate to the CO model. Quite significantly, we provide asymptotic estimates for the actual number of calls, not only the expected value. In Section 6.1.2, we demonstrate the flexibility of our results by showing that they apply to the ANY model as well, strengthening the known results in gossiping theory and in the analysis of population protocols. Indirectly, the transferability of our results also implies that information propagates at essentially the same speed with or without repetitions in the model.

Finally, we do not discuss our results on temporal spanners in this section, although they could be translated in a similar way, because the corresponding problems have not been considered in the areas of gossiping protocols and population protocols. The remainder of this section provides technical details on the claimed equivalences.

6.1.1 Random information exchanges in the CO model

Recall that in our model $\mathcal{F}_{n,1}$, the time labels induce edge orderings of the complete graph, each of which is equiprobable. The same distribution of edge orderings is obtained if we construct a random ordering by choosing edges one by one uniformly at random among all

edges that are not yet chosen, i.e., in the same way as calls are originated in the randomized **CO** model. Therefore, by interpreting edges as calls and time labels as ranks of the calls, the number of edges in $\mathcal{G}_{[0,p]}$, where $\mathcal{G} \sim \mathcal{F}_{n,1}$, is the number of calls that happened no later than time p . Furthermore, since the underlying graph of $\mathcal{G}_{[0,p]}$ is distributed according to $G_{n,p}$ and the number of edges in $G_{n,p}$ is concentrated around $\binom{n}{2}p$ when $p = \Theta(\log n/n)$, our results about random temporal graphs from Section 5 translate straightforwardly to the following results in the randomized **CO** model.

Theorem 6.4. *In the randomized **CO** model, a.a.s.*

1. *the number of calls until two fixed agents exchange their secrets is $\frac{1}{2}n \log n \cdot (1 + o(1))$; (Theorem 5.1)*
2. *the number of calls until at least one agent becomes an expert is $n \log n \cdot (1 + o(1))$; (Theorem 5.3)*
3. *the number of calls until a fixed agent becomes an expert is $n \log n \cdot (1 + o(1))$; (Theorem 5.2)*
4. *the number of calls until all agents become experts is $\frac{3}{2}n \log n \cdot (1 + o(1))$. (Theorem 5.4)*

6.1.2 Random information exchanges in the **ANY** model

By replacing the uniform distribution of edge labels with another distribution, we can obtain random temporal graphs with different temporal dynamics. This flexibility, in particular, allows us to simulate the randomized **ANY** model. In order to achieve this, we replace the uniform distribution for edge labels in $\mathcal{F}_{n,1}$ with a suitable Poisson point process.

More specifically, we introduce a random temporal graph model \mathcal{H}_n , in which the labels of each of the $\binom{n}{2}$ potential edges appear independently according to a Poisson point process with rate 1 starting at time 0 and running infinitely long. In this model each edge gets a countably infinite set of labels with probability 1. We also define a finite random temporal graph model $\mathcal{H}_{n,p} := (\mathcal{H}_n)_{[0,p]}$ where the process is stopped at time p . Note that, in $\mathcal{H}_{n,p}$, an edge may appear an arbitrary number of times, including zero in which case the edge is not present in the underlying graph. The expected number of appearances for each edge is exactly p . Furthermore, as the Poisson distribution has variance equal to the expectation, by Chebyshev's inequality the number of appearances of a fixed edge is concentrated around its expected value. Hence, the total number of edge appearances in $\mathcal{H}_{n,p}$ is concentrated around $\binom{n}{2}p$.

Since all edges appear according to independent identical Poisson point processes, at any fixed point in time, the next edge to appear is distributed uniformly at random among all $\binom{n}{2}$ possible edges, i.e., in the same way as calls are scheduled in the randomized **ANY** model. Hence, if p_0 is a threshold probability for a temporal property in $\mathcal{H}_{n,p}$, then it translates to a $\binom{n}{2}p_0$ -calls threshold in the randomized **ANY** model for the corresponding property.

In this section, we will show that all our main results about $\mathcal{F}_{n,p}$ can be transferred to $\mathcal{H}_{n,p}$ with only minor changes in the proofs. Precisely, the corresponding results from Section 5 translate to the following results in the randomized **ANY** model.

Theorem 6.5. *In the randomized **ANY** model, a.a.s.*

1. *the number of calls until two fixed agents exchange their secrets is $\frac{1}{2}n \log n \cdot (1 + o(1))$;*

2. the number of calls until at least one agent becomes an expert is $n \log n \cdot (1 + o(1))$;
3. the number of calls until a fixed agent becomes an expert is $n \log n \cdot (1 + o(1))$;
4. the number of calls until all agents become experts is $\frac{3}{2}n \log n \cdot (1 + o(1))$.

In the remainder of this section, we explain which changes (if any) need to be made to the statements of Sections 3 to 5 in order to apply to the $\mathcal{H}_{n,p}$ model.

Results based on 2-hop approach (Section 3)

Lemma 6.6 (cf. Lemma 3.1). *Let $\alpha \geq 3$ and let $p = \alpha\sqrt{\log n/n}$. Then, for all $n \geq 4$ and $p \leq 1$, an arbitrary vertex of $(G, \lambda) \sim \mathcal{H}_{n,p}$ is a temporal source with probability at least $1 - n^{-\alpha^2/12+1}$.*

Proof. The construction with intermediate vertices stays the same. Note that the number of edges of a fixed path x, y, z that one can traverse in \mathcal{H}_n before some point in time can be modeled as a Poisson process stopped after two occurrences. Thus, the probability $\mathbb{P}[S_{yz}]$ that vertex x can reach z along said path before time p becomes $1 - e^{-p}(1+p) > 1 - \exp(-p + p - p^2/2 + p^3/3) = 1 - \exp(-p^2/2 + p^3/3) > 1 - \exp(-p^2/6)$. Here the first inequality uses the Taylor series of the natural logarithm, and the second uses $p \leq 1$. The probability p_1 of all $n - 2$ intermediate vertices being unsuitable is at most $(1 - (1 - \exp(-p^2/6)))^{n-2} = \exp(-(p^2/6)(n - 2)) = \exp(-\alpha^2(\log n)(n - 2)/(6n)) \leq n^{-\alpha^2/12}$ by independence. The rest of the proof is the same as before with $\alpha^2/12$ instead of $\alpha^2/4$. \square

Just like Corollary 3.2, we obtain the following corollary.

Corollary 6.7 (cf. Corollary 3.2). *Let $p = \frac{\log n}{\sqrt{n}}$. Then, $(G, \lambda) \sim \mathcal{H}_{n,p}$ is temporally connected with probability at least $1 - n^{-\frac{\log n}{12}+2}$.*

Results on foremost tree evolution (Section 4)

A change in the foremost tree algorithm (Algorithm 1) is necessary to account for the multiple labels on some edges. The algorithm is now applied to the graphs distributed according to \mathcal{H}_n . The change is similar to treating each edge label as a separate edge. More specifically, for each edge e in S_k , i.e., an edge connecting two vertices, exactly one of which is included in the tree T_{k-1} , we consider the earliest label $\lambda_{e,k} \in \lambda(e)$ such that $T_{k-1} \cup \{(e, \lambda_{e,k})\}$ is an increasing temporal tree. With probability 1, such a label exists. Just like before, we select the edge e with the minimal possible $\lambda_{e,k}$.

Apart from the fact that (G, λ) is now sampled from \mathcal{H}_n and notation changes to use $\lambda_{e,k}$ instead of $\lambda(e)$ where appropriate, the definitions of T_k^v , Y_k^v , c_k , \widehat{Y}_k^v and \mathcal{A}_k^v apply without changes, as do Lemma 4.1 and Lemma 4.2:

Lemma 6.8 (cf. Lemma 4.1). *Let (G, λ) be a temporal graph (not necessarily simple) and v be a temporal source in (G, λ) . Then*

- (i) the modified Algorithm 1 constructs a foremost tree for v in (G, λ) ;
- (ii) $\lambda_{e_{1,1}} \leq \lambda_{e_{2,2}} \leq \dots \leq \lambda_{e_{n-1,n-1}}$.

Proof. Apart from notation changes, this lemma is proven exactly as before. \square

Instead of Lemma 4.3, we may state a slightly stronger result here which also replaces Corollary 4.4:

Lemma 6.9 (cf. Lemma 4.3). *For a vertex v , with probability at least $1 - 4/\log n$ the equality $\widehat{X}_k^v = X_k^v$ holds for every $k \in [n - 1]$.*

Further, for a vertex v and every $k \in [n - 1]$, we have

$$(i) \quad \mathbb{E}[X_k^v \mid \mathcal{A}_{k-1}^v] = \frac{1}{k(n-k)};$$

$$(ii) \quad \frac{1 - 1/\log n}{k(n-k)} \leq \mathbb{E}[\widehat{X}_k^v \mid \mathcal{A}_{k-1}^v] \leq \frac{1}{k(n-k)}.$$

Proof. The proof of the first part of the lemma is exactly the same as the proof of corresponding claim in Lemma 4.3 with the only difference that $e^{-k(n-k)c_k}$ (from the computation in Lemma 4.2) is now the exact value of $\mathbb{P}(X_k^v \geq c_k)$ rather than an upper bound on it.

(i): Due to the memorylessness of the exponential distribution, X_k^v conditioned on \mathcal{A}_{k-1}^v is distributed as the minimum of $k(n-k)$ independently $\text{Exp}(1)$ -distributed random variables. Thus, the distribution of X_k^v is exactly $\text{Exp}(k(n-k))$ and as such it has expected value $\frac{1}{k(n-k)}$.

(ii): The only difference to the proof of Lemma 4.3 is in the evaluation of the integral

$$\begin{aligned} \int_0^{c_k} \mathbb{P}[X_k^v \geq t \mid \mathcal{A}_{k-1}^v] dt &= \int_0^{c_k} \mathbb{P}[X_k^v \geq t] dt \\ &= \int_0^{c_k} \exp(-k(n-k)t) dt \\ &= \frac{1 - \exp(-k(n-k)c_k)}{k(n-k)}. \end{aligned}$$

Then, the desired bound follows from the fact that $\exp(-k(n-k)c_k) \leq \frac{1}{\log n}$ (see (1)). \square

The remaining results in Section 4.2 (Lemma 4.6 to Lemma 4.12) and their proofs carry over to $\mathcal{H}_{n,p}$ with only trivial modifications.

Regarding Section 4.3, we observe that Lemma 4.13 holds trivially for \mathcal{H}_n , as the attachment point is uniformly distributed for the Poisson point process case. Consequently, Lemma 4.14 and Lemma 4.15 also hold, and so do Theorem 4.16 and Theorem 4.17 as all the properties of $\mathcal{F}_{n,p}$ used in the proofs are shared by $\mathcal{H}_{n,p}$.

Results on sharp thresholds for temporal graph properties (Section 5)

Theorems 5.1 to 5.3 (sharp thresholds for point-to-point reachability, temporal source, and first temporal source) apply without modifications as they follow solely from the concentration results obtained in Section 4.2. Lemma 5.5 only requires a minor change to the proof:

Lemma 6.10 (cf. Lemma 5.5). *Let $p \leq \frac{3 \log n}{n} - \varepsilon$, where $\varepsilon := \frac{6(\log n)^{0.8}}{n}$, and let $\mathcal{G} \sim \mathcal{H}_n$. Then, a.a.s. the temporal graph $\mathcal{G}_{[0,p]}$ contains at least one vertex which is not a temporal sink.*

Proof. The proof works the same, except that the probability for a pair of vertices to form an edge in the underlying graph H of $\mathcal{G}_{[q,p]}$ is always

$$\gamma = 1 - e^{q-p} \leq p - q \leq \frac{3 \log n}{n} - \varepsilon - \frac{2 \log n}{n} + \frac{\varepsilon}{2} < \frac{\log n - (\log n)^{0.8}}{n}.$$

In particular, H is distributed as an element of $G_{n,\gamma}$ and the remainder of the proof may be copied as is. \square

The proof of Lemma 5.6 is simplified due to the fact that we do not need to bound $\Delta(\mathcal{G}_{[0,q]})$ anymore.

Lemma 6.11 (cf. Lemma 5.6). *Let $p \geq \frac{3 \log n}{n} + \varepsilon$, where $\varepsilon := \frac{3(\log n)^{0.8}}{n}$, and let $\mathcal{G} \sim \mathcal{H}_n$. Then, a.a.s. every vertex in $\mathcal{G}_{[0,p]}$ is a temporal sink.*

Proof. We define q and r and S as in Lemma 5.6, and observe that $|V \setminus S| \leq r$ a.a.s. The set C_w of potential edges connecting $w \in V \setminus S$ to an element of S is simplified to

$$C_w = \{vw \mid v \in S\}$$

which now trivially implies $|C_w| = n - 1 - r \geq n - 2r = d$. As the waiting time of each edge is exponentially distributed, we now get $\mathbb{P}[T_w > x] \leq e^{-xd}$. In particular, the estimation

$$\tau := \mathbb{P}\left[T_w > \frac{\log n}{n}\right] \leq e^{-(\log n - 2 \log \log \log n)}$$

still applies, so the remainder of the proof may be copied without modifications. \square

Thus, we again obtain Theorem 5.4, i.e., the sharp threshold on temporal connectivity in $\mathcal{H}_{n,p}$.

Of our results in Section 5.4, all except Theorem 5.13 are based solely on results which we have already proven to also hold for $\mathcal{H}_{n,p}$. Finally, Theorem 5.13 only requires a minor modification to the proof.

Theorem 6.12 (cf. Theorem 5.13). *$\mathcal{H}_{n,p}$ a.a.s. has an optimal spanner if $p \geq 4 \frac{\log n}{n} + \varepsilon$ where $\varepsilon := \frac{16(\log n)^{0.8}}{n} \in o\left(\frac{\log n}{n}\right)$.*

Proof. In comparison to the proof of Theorem 5.13, only the probability of an edge appearing in an interval of length ε_0 changes from ε_0 to $1 - e^{-\varepsilon_0}$. Thus, $\mathbb{E}[S] = n^{(4)}(1 - e^{-\varepsilon_0})^4$ and

$$\mathbb{E}[S^2] \leq n^{(8)}(1 - e^{-\varepsilon_0})^8 + \sum_{i=4}^7 8^i n^i (1 - e^{-\varepsilon_0})^i \leq \mathbb{E}[S]^2 + 8^7 \cdot \sum_{i=4}^7 n^i \varepsilon_0^i.$$

Therefore, the bound on the variance $\mathbb{D}[S]$ and the remainder of the proof remain valid. \square

6.2 Complete graphs with random weights

Another model related to RSTGs (technically, to $\mathcal{F}_{n,p}$), although less closely than gossiping and population protocols scheduling, is the randomly edge-weighted complete graph model studied by Janson [39] and related to percolation theory. Janson's model considers a (static) complete graph with edge weights drawn independently and uniformly⁴ from $[0, 1]$. Instead of arrival times, this model simply measures path lengths as sums of edge weights. Intriguingly, some of our sharp thresholds have analogs with the same numerical values in that model. Part of the underlying reason is that, as shown in our analysis in Lemma 4.3, if one reaches either endpoint of any particular edge at any point in time which is sufficiently close to 0, then the waiting time until that edge appears is distributed close to uniformly on $[0, 1]$. The latter condition matches the assumption made by Janson about the distribution of the edge weights in his model and causes the growth of single-source trees to

⁴The results also hold for any other distribution that closely matches the uniform distribution around 0.

behave identically in both models. This explains why the time thresholds for Point-to-Point Reachability and Temporal Source in our model agree with the weight thresholds for One-to-One Shortest Path and One-to-All Shortest Paths in Janson’s model, respectively.

However, the above argument is not sufficient to account for the coincidence of the thresholds involving multiple sources, i.e., the time threshold for Temporal Connectivity and the weight threshold for All-to-All Shortest Paths. Indeed, the strong independence assumption in the randomly edge-weighted complete graph model does not hold even in the gossiping model with repeated interactions: In the former, the time required to cross an edge (i.e., the edge weight) is independent of all other factors, whereas in temporal graphs and gossiping two temporal paths emanating from different sources, in general, need to wait different amounts of time to cross the same edge, as those depend on the times at which the paths reach an endpoint of that edge. The rather surprising observation that the above two thresholds are nevertheless the same is ultimately a consequence of the fact that the determining factor for pairwise connectivity in each of the two models turns out to be the waiting time until at least one edge appears for every vertex. This time is identical in both models as all vertices start at time 0.

In essence, one might say that our results prove the temporal dependencies present in random temporal graphs to be insignificant for reachability times, as witnessed by the fact that the resulting waiting time thresholds coincide with analogous thresholds in the simple weighted complete graph model, where such dependencies are absent. Note, however, that our results aim in a different direction than those of Janson [39]. Instead of characterizing the shape of the limit probability distribution, we focus on the speed of convergence and on the properties of partial foremost trees. Also, it is not clear that the aspects related to spanners have natural interpretations in Janson’s model, as these graphs admit classical spanning trees unconditionally, which is not true in temporal graphs, whatever the density.

6.3 More dependencies

Our techniques can be applied to models in which edge appearances are more dependent, and their distribution deviates even further from uniformity. For instance, with some technical work we can show that the results obtained in this paper also hold in a model where an adversary blocks a small fraction of the potential edges at each vertex before any random choice is made. Additionally, we could also allow an adversary to pick different time label distributions for the edges as long as each of them is asymptotically close to the uniform distribution around time 0.

6.4 Edge-ordered graphs

If only the relative order of time labels but not their absolute values is of interest, simple temporal graphs with pairwise different time labels are clearly equivalent to *edge-ordered graphs*, which are graphs with a given total order of their edges.

In combinatorics, edge-ordered graphs have been studied from various perspectives. Perhaps closest to our work is the study of the lengths of longest monotone paths and walks, i.e., paths or walks whose edge sequences are strictly increasing. Note that in the terminology of temporal graphs, these are exactly the notions of temporal paths or walks.

Interest in this topic dates back to the ’70s, with Chvátal and Komlós [24] asking for worst-case bounds on the length of a longest monotone path or walk in any edge ordering of the complete graph K_n . Graham and Kleitman [35] gave an answer for the case of walks and found a lower bound of $\Omega(\sqrt{n})$ and an upper bound of $3n/4$ for the case of paths. The upper bound has subsequently been improved to roughly $n/2$ [18]. The lower bound stayed

current for a long time, but recently Bucić et al. [14] raised it to $n^{1-o(1)}$, nearly closing the gap.

More closely related to our paper is the work of Lavrov and Loh [43], who considered random edge orderings of K_n and found that even monotone Hamiltonian paths exist with probability at least $1/e$ and conjectured them to exist a.a.s. This conjecture was recently proven true by Martinsson [44]. Even more recently, Angel et al. [4] established the longest monotone walk in a random edge ordering of $G_{n,p}$ to a.a.s. have length about enp . These results can be interpreted also for temporal graphs. Specifically the result of Martinsson shows that any $\mathcal{G} \sim \mathcal{F}_{n,1}$ a.a.s. contains a temporal path visiting all vertices, while that of Angel et al. implies that the length of a longest temporal walk in $\mathcal{F}_{n,p}$ is concentrated around enp .

7 Concluding remarks and open questions

In this paper, we studied a natural model of random temporal graphs, in which every edge of an Erdős–Rényi graph $G_{n,p}$ is assigned a single presence time chosen uniformly at random in the unit interval $[0, 1]$. The study of various degrees of temporal reachability and existence of small temporal spanners in these graphs revealed a rich diversity of thresholds, in stark contrast with static graphs. Put together, these thresholds offer a measure of the discrepancy between static and temporal connectivity in random graphs.

Despite the simplicity of RSTGs, we have shown that they capture, in a scale-preserving way, several classical phenomena observed in gossip theory and population protocols. Furthermore, some of these results were shown to strengthen and/or to complete existing results in these fields. In stark contrast with *deterministic* temporal graphs, where significant obstructions exist which prevent the existence of linear-size spanners even in some dense graphs, we have shown that all possible obstructions to the existence of nearly optimal temporal spanners turn out to be statistically insignificant in random temporal graphs. Thus, the concept of a nearly optimal spanner recovers some form of universality in the spirit of spanning trees in static graphs. Observe that the spanners we construct are highly centralized, with almost all communication passing through a single pivot vertex. It would be interesting to determine whether this is inherent or incidental.

All but one of our characterizations are sharp thresholds. The remaining one concerns the existence of optimal temporal spanners (i.e., spanners of size exactly $2n - 4$) in RSTGs, for which we prove an upper bound of $p = 4 \log n/n$. Whether this bound is actually a sharp threshold is left open; we conjecture that it is (Conjecture 2).

Now that the most basic phase transitions in temporal reachability are characterized in RSTGs, it would be quite natural to start looking at more complex properties, e.g., motivated by networking applications. On the one hand, one might want to obtain connectivity via at least 2 (or k) disjoint temporal paths, where disjoint can mean edge-disjoint or vertex-disjoint. On the other hand, one could extend the task of sending a message to the task of sending a message and receiving a reply, i.e., roundtrip communication. We conjecture that there are similar sharp thresholds for roundtrip communication. In particular, our preliminary investigation suggests the following:

Conjecture 3 (Roundtrip connectivity).

1. For two fixed vertices, the sharp threshold for possibility of roundtrip communication is $2 \log n/n$;
2. For one fixed vertex u , the sharp threshold for roundtrip communication from u to all other vertices and back, as well as the sharp threshold for roundtrip communication

from all the other vertices to u and back, is $3 \log n/n$;

3. The sharp threshold for roundtrip communication between all pairs of vertices (i.e., roundtrip connectivity) is $4 \log n/n$.

A classical concept in Erdős–Rényi random graphs is the one of a *giant component*, i.e., a set of $\Theta(n)$ mutually reachable vertices. In this paper, we showed that $\log n/n$ is a sharp threshold for any fixed pair of vertices to reach each other via temporal paths. However, in a subtle way, this property does not imply that a subset of $\Theta(n)$ vertices can *all* reach each other (due to the intransitivity of temporal paths). Nevertheless, we conjecture that such components also appear at $\log n/n$ and that the threshold is sharp, in contrast to static graphs where the giant component emerges progressively.

Conjecture 4. $\log n/n$ is a sharp threshold for the existence of a temporal component of size $\Theta(n)$. In fact, the same holds even for components of size $n - o(n)$.

Finally, we have shown in the paper that in an RSTG, every pair of temporally connected vertices is connected by a temporal path containing at most a logarithmic number of edges. We conjecture that for sufficiently sparse RSTGs this also constitutes a lower bound. Precisely:

Conjecture 5. There exists a positive-valued function f such that for any constant C , if $p = C \log n/n$, then the graph a.a.s. contains some pairs of vertices with no connecting temporal path having fewer than $f(C) \log n$ edges.

In a dedicated section, we have shown that our analyses can be applied essentially in the same way for multiple types of processes beyond RSTGs. We hope that this versatility can be extended further and characterized by some simple conditions. This was the case with *parallel* rumor spreading time [28], and we hope that sequential rumor propagation could enjoy the same.

In conclusion, the notion of reachability in temporal graphs is quite different from reachability in static graphs, and the present paper illustrates the fact that reachability in *random* temporal graphs is also quite different from reachability in *deterministic* temporal graphs. We hope that our results will help paving the way for further (and more complex) investigations of temporal reachability.

Acknowledgements. We are very grateful to the anonymous reviewers for their careful reading and many useful and detailed comments that improved the presentation of the paper.

References

- [1] Eleni C. Akrida, Leszek Gasieniec, George B. Mertzios, and Paul G. Spirakis. “The complexity of optimal design of temporally connected graphs”. In: *Theory of Computing Systems* 61.3 (2017), pp. 907–944. DOI: [10.1007/s00224-017-9757-x](https://doi.org/10.1007/s00224-017-9757-x).
- [2] Eleni C. Akrida, George B. Mertzios, Sotiris E. Nikolettseas, Christoforos L. Raptopoulos, Paul G. Spirakis, and Viktor Zamaraev. “How fast can we reach a target vertex in stochastic temporal graphs?”. In: *Journal of Computer and System Sciences* 114 (2020), pp. 65–83. DOI: [10.1016/j.jcss.2020.05.005](https://doi.org/10.1016/j.jcss.2020.05.005).
- [3] Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis, and Viktor Zamaraev. “Temporal vertex cover with a sliding time window”. In: *Journal of Computer and System Sciences* 107 (2020), pp. 108–123. DOI: [10.1016/j.jcss.2019.08.002](https://doi.org/10.1016/j.jcss.2019.08.002).

- [4] Omer Angel, Asaf Ferber, Benny Sudakov, and Vincent Tassion. “Long Monotone Trails in Random Edge-Labelings of Random Graphs”. In: *Combinatorics, Probability and Computing* 29.1 (2020), pp. 22–30. DOI: [10.1017/S096354831900018X](https://doi.org/10.1017/S096354831900018X).
- [5] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. “Computation in networks of passively mobile finite-state sensors”. In: *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*. 2004, pp. 290–299. DOI: [10.1145/1011767.1011810](https://doi.org/10.1145/1011767.1011810).
- [6] Chen Avin, Michal Koucký, and Zvi Lotker. “How to explore a fast-changing world (cover time of a simple random walk on evolving graphs)”. In: *Proceedings of the 35th International Colloquium on Automata, Languages, and Programming (ICALP)*. 2008, pp. 121–132. DOI: [10.1007/978-3-540-70575-8_11](https://doi.org/10.1007/978-3-540-70575-8_11).
- [7] Kyriakos Axiotis and Dimitris Fotakis. “On the Size and the Approximability of Minimum Temporally Connected Subgraphs”. In: *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*. 2016, 149:1–149:14. DOI: [10.4230/LIPIcs.ICALP.2016.149](https://doi.org/10.4230/LIPIcs.ICALP.2016.149).
- [8] Kazuoki Azuma. “Weighted sums of certain dependent random variables”. In: *Tohoku Mathematical Journal* 19.3 (1967), pp. 357–367. DOI: [10.2748/tmj/1178243286](https://doi.org/10.2748/tmj/1178243286).
- [9] Hervé Baumann, Pierluigi Crescenzi, and Pierre Fraigniaud. “Parsimonious flooding in dynamic graphs”. In: *Distributed Computing* 24.1 (2011), pp. 31–44. DOI: [10.1007/s00446-011-0133-9](https://doi.org/10.1007/s00446-011-0133-9).
- [10] Sandeep Bhadra and Afonso Ferreira. “Complexity of connected components in evolving graphs and the computation of multicast trees in dynamic networks”. In: *Proceedings of the 2nd International Conference on Ad-Hoc Networks and Wireless (ADHOC-NOW)*. 2003, pp. 259–270. DOI: [10.1007/978-3-540-39611-6_23](https://doi.org/10.1007/978-3-540-39611-6_23).
- [11] Béla Bollobás. *Random Graphs*. 2nd ed. Cambridge University Press, 2001. DOI: [10.1017/CB09780511814068](https://doi.org/10.1017/CB09780511814068).
- [12] Béla Bollobás. *The art of mathematics: Coffee time in Memphis*. Cambridge University Press, 2006.
- [13] David W. Boyd and J. Michael Steele. “Random exchanges of information”. In: *Journal of Applied Probability* 16.3 (1979), pp. 657–661. DOI: [10.2307/3213094](https://doi.org/10.2307/3213094).
- [14] Matija Bucić, Matthew Kwan, Alexey Pokrovskiy, Benny Sudakov, Tuan Tran, and Adam Zsolt Wagner. “Nearly-linear monotone paths in edge-ordered graphs”. In: *Israel Journal of Mathematics* 238.2 (2020), pp. 663–685. DOI: [10.1007/s11856-020-2035-7](https://doi.org/10.1007/s11856-020-2035-7).
- [15] Binh-Minh Bui-Xuan, Afonso Ferreira, and Aubin Jarry. “Computing shortest, fastest, and foremost journeys in dynamic networks”. In: *International Journal of Foundations of Computer Science* 14.02 (2003), pp. 267–285. DOI: [10.1142/S0129054103001728](https://doi.org/10.1142/S0129054103001728).
- [16] Richard T. Bumby. “A Problem with Telephones”. In: *SIAM Journal on Algebraic Discrete Methods* 2.1 (1979). DOI: [10.1137/0602002](https://doi.org/10.1137/0602002).
- [17] Janna Burman, Ho-Lin Chen, Hsueh-Ping Chen, David Doty, Thomas Nowak, Eric Severson, and Chuan Xu. “Time-optimal self-stabilizing leader election in population protocols”. In: *Proceedings of the 40th Annual ACM Symposium on Principles of Distributed Computing (PODC)*. 2021, pp. 33–44. DOI: [10.1145/3465084.3467898](https://doi.org/10.1145/3465084.3467898).
- [18] A. Robert Calderbank, Fan R. K. Chung, and Dean G. Sturtevant. “Increasing sequences with nonzero block sums and increasing paths in edge-ordered graphs”. In: *Discrete Mathematics* 50 (1984), pp. 15–28. DOI: [10.1016/0012-365X\(84\)90031-1](https://doi.org/10.1016/0012-365X(84)90031-1).

- [19] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. “Time-varying graphs and dynamic networks”. In: *International Journal of Parallel, Emergent and Distributed Systems* 27.5 (2012), pp. 387–408. DOI: [10.1080/17445760.2012.668546](https://doi.org/10.1080/17445760.2012.668546).
- [20] Arnaud Casteigts, Anne-Sophie Himmel, Hendrik Molter, and Philipp Zschoche. “Finding Temporal Paths Under Waiting Time Constraints”. In: *Algorithmica* (2021). DOI: [10.1007/s00453-021-00831-w](https://doi.org/10.1007/s00453-021-00831-w).
- [21] Arnaud Casteigts, Joseph G. Peters, and Jason Schoeters. “Temporal cliques admit sparse spanners”. In: *Journal of Computer and System Sciences* 121 (2021), pp. 1–17. DOI: [10.1016/j.jcss.2021.04.004](https://doi.org/10.1016/j.jcss.2021.04.004).
- [22] Arnaud Casteigts, Michael Raskin, Malte Renken, and Viktor Zamaraev. “Sharp thresholds in random simple temporal graphs”. In: *Proceedings of the 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 319–326. DOI: [10.1109/FOCS52979.2021.00040](https://doi.org/10.1109/FOCS52979.2021.00040).
- [23] Augustin Chaintreau, Abderrahmen Mtibaa, Laurent Massoulie, and Christophe Diot. “The diameter of opportunistic mobile networks”. In: *Proceedings of the 3rd International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. 2007, pp. 1–12. DOI: [10.1145/1364654.1364670](https://doi.org/10.1145/1364654.1364670).
- [24] Václav Chvátal and János Komlós. “Some combinatorial theorems on monotonicity”. In: *Canadian Mathematical Bulletin* 14.2 (1971), pp. 151–157. DOI: [10.4153/CMB-1971-028-8](https://doi.org/10.4153/CMB-1971-028-8).
- [25] Andrea Clementi, Claudio Macci, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. “Flooding time in edge-markovian dynamic graphs”. In: *Proceedings of the 27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*. 2008, pp. 213–222. DOI: [10.1145/1400751.1400781](https://doi.org/10.1145/1400751.1400781).
- [26] Andrea Clementi, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. “Broadcasting in dynamic radio networks”. In: *Journal of Computer and System Sciences* 75.4 (2009), pp. 213–230. DOI: [10.1016/J.JCSS.2008.10.004](https://doi.org/10.1016/J.JCSS.2008.10.004).
- [27] Hans van Ditmarsch, Ioannis Kokkinis, and Anders Stockmarr. “Reachability and expectation in gossiping”. In: *Proceedings of PRIMA 2017: Principles and Practice of Multi-Agent Systems*. 2017, pp. 93–109. DOI: [10.1007/978-3-319-69131-2_6](https://doi.org/10.1007/978-3-319-69131-2_6).
- [28] Benjamin Doerr and Anatolii Kostrygin. “Randomized Rumor Spreading Revisited”. In: *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP)*. Vol. 80. 2017, 138:1–138:14. DOI: [10.4230/LIPIcs.ICALP.2017.138](https://doi.org/10.4230/LIPIcs.ICALP.2017.138).
- [29] Jessica Enright, Kitty Meeks, and Fiona Skerman. “Assigning times to minimise reachability in temporal graphs”. In: *Journal of Computer and System Sciences* 115 (2021), pp. 169–186. DOI: [10.1016/j.jcss.2020.08.001](https://doi.org/10.1016/j.jcss.2020.08.001).
- [30] Paul Erdős and Alfréd Rényi. “On the evolution of random graphs”. In: *A Magyar Tudományos Akadémia. Matematikai Kutató Intézetének Közleményei* 5 (1960), pp. 17–61.
- [31] Till Fluschnik, Hendrik Molter, Rolf Niedermeier, Malte Renken, and Philipp Zschoche. “Temporal graph classes: A view through temporal separators”. In: *Theoretical Computer Science* 806 (2020), pp. 197–218. DOI: [10.1016/j.tcs.2019.03.031](https://doi.org/10.1016/j.tcs.2019.03.031).
- [32] Alan Frieze and Michał Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015. DOI: [10.1017/CB09781316339831](https://doi.org/10.1017/CB09781316339831).

- [33] Alan M. Frieze and Geoffrey R. Grimmett. “The shortest-path problem for graphs with random arc-lengths”. In: *Discrete Applied Mathematics* 10.1 (1985), pp. 57–77. DOI: [10.1016/0166-218X\(85\)90059-9](https://doi.org/10.1016/0166-218X(85)90059-9).
- [34] Emmanuel Godard and Dorian Mazauric. “Computing the dynamic diameter of non-deterministic dynamic networks is hard”. In: *Proceedings of the 10th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics (ALGOSENSORS)*. 2014, pp. 88–102. DOI: [10.1007/978-3-662-46018-4_6](https://doi.org/10.1007/978-3-662-46018-4_6).
- [35] Ronald L. Graham and Daniel J. Kleitman. “Increasing paths in edge ordered graphs”. In: *Periodica Mathematica Hungarica* 3.1-2 (1973), pp. 141–148. DOI: [10.1007/BF02018469](https://doi.org/10.1007/BF02018469).
- [36] Peter Grindrod and Desmond J. Higham. “Evolving graphs: dynamical models, inverse problems and propagation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 466.2115 (2010), pp. 753–770. DOI: [10.1098/rspa.2009.0456](https://doi.org/10.1098/rspa.2009.0456).
- [37] John Haigh. “Random exchanges of information”. In: *Journal of Applied Probability* 18.3 (1981), pp. 743–746. DOI: [10.2307/3213330](https://doi.org/10.2307/3213330).
- [38] Petter Holme and Jari Saramäki, eds. *Temporal Network Theory*. Springer, 2019. DOI: [10.1007/978-3-030-23495-9](https://doi.org/10.1007/978-3-030-23495-9).
- [39] Svante Janson. “One, Two And Three Times Log N/N For Paths In A Complete Graph With Random Weights”. In: *Combinatorics, Probability and Computing* 8.4 (1999), pp. 347–361. DOI: [10.1017/S0963548399003892](https://doi.org/10.1017/S0963548399003892).
- [40] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. “Randomized rumor spreading”. In: *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)*. 2000, pp. 565–574. DOI: [10.1109/SFCS.2000.892324](https://doi.org/10.1109/SFCS.2000.892324).
- [41] David Kempe, Jon Kleinberg, and Amit Kumar. “Connectivity and Inference Problems for Temporal Networks”. In: *Journal of Computer and System Sciences* 64.4 (2002), pp. 820–842. DOI: [10.1006/jcss.2002.1829](https://doi.org/10.1006/jcss.2002.1829).
- [42] Jon Kleinberg and Éva Tardos. *Algorithm design*. Pearson, 2006.
- [43] Mikhail Lavrov and Po-Shen Loh. “Increasing Hamiltonian paths in random edge orderings”. In: *Random Structures & Algorithms* 48.3 (2016), pp. 588–611. DOI: [10.1002/rsa.20592](https://doi.org/10.1002/rsa.20592).
- [44] Anders Martinsson. “Most edge-orderings of K_n have maximal altitude”. In: *Random Structures & Algorithms* 54.3 (2019), pp. 559–585. DOI: [10.1002/rsa.20803](https://doi.org/10.1002/rsa.20803).
- [45] George B. Mertzios, Hendrik Molter, Rolf Niedermeier, Viktor Zamaraev, and Philipp Zschoche. “Computing Maximum Matchings in Temporal Graphs”. In: *Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science (STACS)*. Vol. 154. 2020, 27:1–27:14. DOI: [10.4230/LIPIcs.STACS.2020.27](https://doi.org/10.4230/LIPIcs.STACS.2020.27).
- [46] George B. Mertzios, Hendrik Molter, and Viktor Zamaraev. “Sliding window temporal graph coloring”. In: *Journal of Computer and System Sciences* 120 (2021), pp. 97–115. DOI: [10.1016/j.jcss.2021.03.005](https://doi.org/10.1016/j.jcss.2021.03.005).
- [47] Yves Mocquard, Bruno Sericola, and Emmanuelle Anceaume. “Probabilistic Analysis of Rumor-Spreading Time”. In: *Inform's Journal on Computing* 32.1 (2020), pp. 172–181. DOI: [10.1287/ijoc.2018.0845](https://doi.org/10.1287/ijoc.2018.0845).

- [48] Yves Mocquard, Bruno Sericola, Samantha Robert, and Emmanuelle Anceaume. “Analysis of the propagation time of a rumour in large-scale distributed systems”. In: *Proceedings of the 15th IEEE International Symposium on Network Computing and Applications*. 2016, pp. 264–271. DOI: [10.1109/NCA.2016.7778629](https://doi.org/10.1109/NCA.2016.7778629).
- [49] John W. Moon. “Random exchanges of information”. In: *Nieuw Archief voor Wiskunde* 20 (1972), pp. 246–249.
- [50] Micha Sharir. “A strong-connectivity algorithm and its applications in data flow analysis”. In: *Computers & Mathematics with Applications* 7.1 (1981), pp. 67–72. DOI: [10.1016/0898-1221\(81\)90008-0](https://doi.org/10.1016/0898-1221(81)90008-0).
- [51] Xiao Zhang, Cristopher Moore, and Mark Newman. “Random graph models for dynamic networks”. In: *The European Physical Journal B* 90.10 (2017), pp. 1–14. DOI: [10.1140/epjb/e2017-80122-8](https://doi.org/10.1140/epjb/e2017-80122-8).