Sharp Thresholds for Temporal Motifs and Doubling Time in Random Temporal Graphs*

Henry Austin
Department of Computer Science, Durham University, UK

George B. Mertzios
Department of Computer Science, Durham University, UK

Paul G. Spirakis
Department of Computer Science, University of Liverpool, UK

Abstract

In this paper we study two natural models of *random temporal* graphs. In the first, the *continuous* model, each edge e is assigned l_e labels, each drawn uniformly at random from $(0, 1]$, where the numbers l_e are independent random variables following the same discrete probability distribution. In the second, the *discrete* model, the l_e labels of each edge e are chosen uniformly at random from a set $\{1, 2, \ldots, T\}$. In both models we study the existence of *δ-temporal motifs*. Here a δ-temporal motif consists of a pair (H, P), where H is a fixed static graph and P is a partial order over its edges. A temporal graph $\mathcal{G} = (G, \lambda)$ contains (H, P) as a δ-temporal motif if $\mathcal{G}$ has a simple temporal subgraph on the edges of H whose time labels are ordered according to P, and whose life duration is at most δ. We prove *sharp existence thresholds* for all δ-temporal motifs, and we identify a qualitatively different behavior from the analogous static thresholds in Erdős-Rényi random graphs. Applying the same techniques, we then characterize the growth of the largest δ-temporal clique in the continuous variant of our random temporal graphs model. Finally, we consider the *doubling time* of the reachability ball centered on a small set of vertices of the random temporal graph as a natural proxy for temporal expansion. We prove *sharp upper and lower bounds* for the maximum doubling time in the continuous model.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures; Mathematics of computing → Discrete mathematics

Keywords and phrases Random temporal graph, δ-temporal motif, sharp upper and lower bounds, doubling time.

Digital Object Identifier 10.4230/LIPIcs.MFCS.2026.79

Related Version https://arxiv.org/abs/2602.01847

Funding *George B. Mertzios*: Supported by the EPSRC grant EP/P020372/1.
Paul G. Spirakis: Supported by the EPSRC grant EP/P02002X/1.

* The full paper with detailed proofs can be found at [8].

51st International Symposium on Mathematical Foundations of Computer Science (MFCS 2026).
Editors: Michal Koucký and Daniela Petrişan; Article No. 79; pp. 79:1–79:20
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

A temporal (or dynamic) graph is a graph whose topology is subject to discrete changes over time. This paradigm reflects the structure and operation of a vast variety of modern real-life networks; social networks, wired or wireless networks whose links change dynamically, transportation networks, and several physical systems are only a few examples of networks that change over time [28, 34, 46, 49, 53]. Embarking from the fundamental model for temporal graphs, introduced by Kempe et al. [31] with one (integer) time-label per edge, and from its extension which allows multiple (integer) time-edges per edge [40], we consider here a relaxed version of the model where the time-labels are drawn from the positive real numbers in $\mathbb{R}_+$ (see Definition 1). This relaxation of the model allows us to randomly draw each time label uniformly at random from a time interval, e.g. from $(0, 1]$.

▶ **Definition 1** (temporal graph). *A* temporal graph *is a pair $\mathcal{G} = (G, \lambda)$, where $G = (V, E)$ is a footprint (static) graph with n denoting the number of vertices in V, and $\lambda : E \to 2^{\mathbb{R}_+}$ is a* labeling *function that assigns to every edge of G a finite set of time labels, each being a positive real number. $\mathcal{G}$ is a* simple *temporal graph if every edge has exactly one time label, i.e. the labeling function is $\lambda : E \to \mathbb{R}_+$.*

Temporal graphs are traditionally used to model information flow which, due to causality, can only pass along a sequence of edges with increasing time labels. Motivated by this, the majority of work concerning temporal graphs is based on the notion of temporal paths and other "path-related" notions, such as temporal analogues of distance, diameter, reachability, exploration, and centrality [2, 3, 18–21, 32, 33, 40, 45, 48, 52]. Only relatively recently have attempts been made to define "non-path" temporal graph problems, such as temporal analogues of clique [17,27,56,57], vertex cover [4,25], coloring [1,30,39,43], matching [11,16,41], vertex separators [22, 26, 38], and transitivity [42].

Continuing in this line of research on "non-path" temporal problems, we mainly focus in this paper on the existence of a given (arbitrary) *temporal motif* in a temporal graph. A temporal motif consists of a pair (H, P), where H is a fixed (static) graph and P is a partial ordering over the edges of H. A temporal graph $\mathcal{G} = (G, \lambda)$ is said to contain a temporal motif (H, P), if H is contained in the footprint G, and the edges of H have time labels which are ordered in agreement with the partial order P. The existence and frequency of temporal motifs is a question of both theoretical and practical importance. On the one hand, theoretically, temporal motifs stand as an intermediate object containing less temporal information than a temporal subgraph, but more temporal information than a footprint subgraph. In particular, they are well suited for capturing notions of causality and are less sensitive to the precise numerical timings of labels. On the other hand, temporal motifs are widely used in the practical analysis of temporal networks, for both identifying significant interaction patterns and determining the mechanics of network evolution [54]. An especially common technique in this context is comparing the frequency of temporal motifs in a given empirical network with that of some appropriate (and frequently ad hoc) network [24, 35, 54]. A strengthening of this concept, which we study, is the notion of a *δ-temporal motif* (i.e. one where all of its constituent edge labels must lie within some window of length δ). These, in turn, represent a middle ground between a temporal subgraph and a temporal motif. They have seen application both in fraud detection [58] and as a tool for studying different interaction scales in empirical networks [50].

Additionally, the vast majority of theoretical work has focused on deterministic worst-case temporal graphs. While this has led to a relatively good understanding of the worst-case temporal graph properties, the approach provides little information about the vast majority of temporal graphs. This distinction was recently and dramatically demonstrated, by the discovery that almost all simple temporal graphs permit nearly optimal spanners [15], despite the existence of temporal graphs without sub-quadratic spanners [9]. Results such as these raise the general question: to what extent do such unique properties hold for *random temporal graphs* rather than only for carefully constructed ones?

1.1 Our Contribution

In this paper we focus on non-path problems on randomly generated temporal graphs. We expand on earlier work by [15, 44], by exploring the properties of two model variants (one continuous and one discrete model) for random temporal graphs, generalizing the Random Simple Temporal Graph Model (RSTG) [15, 44]. For ease of reading we defer the formal definition of the models (see Definitions 5 and 6). Intuitively, in the continuous model we sample the time labels uniformly at random from $(0, 1]$ for each edge, where the number of labels on each edge is sampled independently for each edge from some discrete distribution. The discrete model is the natural discretization of the continuous one, where the labels of each edge are chosen uniformly at random from a set $\{1, 2, \ldots, T\}$. In our first result, we explore the existence thresholds for all fixed δ-temporal motifs, in terms of the length δ of the permitted time window. Since the temporal graphs we consider here typically have dense foot-print graphs, but disparate time labels, this serves as a useful, if imperfect, analogue for the existence thresholds found in sparse static random graphs. In fact we identify a behavior that is similar, but subtly different from the existence thresholds of fixed static graphs in the Erdős-Rényi random graph model. In particular, the threshold is determined *not* by the reciprocal of the density of the densest subgraph, but instead by the following quantity which we call the *sparsity*.

▶ **Definition 2.** *For a graph $G = (V, E)$, the* sparsity *ρ_G of G is defined as follows,*

$$\rho_G = \min_{H \sqsubseteq G} \frac{|V_H|}{|E_H| - 1},$$

where $H \sqsubseteq G$ denotes that H is a subgraph of G containing at least 2 edges and V_H (resp. E_H) are the set of vertices (resp. edges) comprising H.

Intuitively, this stems from the fact that a random temporal graph has many different intervals which could potentially contain a given δ-temporal motif. Individually, these intervals contain motifs at thresholds roughly in accordance with the existence of their footprint in an Erdős-Rényi random graph and the number of such intervals needed to cover the lifetime of a given temporal graph is inversely proportional to the length of the intervals. However, a direct proof along these lines is made challenging due to the dependency between these intervals, especially when the density of the graph is itself a random variable.

This seemingly minor change, from the reciprocal of density to sparsity, is sufficient to cause qualitatively different threshold behavior from the sparse static case. For example, δ-temporal motifs built from cycles of different sizes have asymptotically distinct existence thresholds in our model, unlike their static counterparts [23] in the Erdős-Rényi random graph. Another consequence of this result is that the asymptotic thresholds are independent of the edge ordering of a given motif. This implies, for example, that time-respecting paths of a given length have the same asymptotic threshold as non-time-respecting paths of the same length. Furthermore, we extend this behavior beyond fixed motifs to consider the size of the largest temporal motif with a clique footprint, and obtain bounds generalizing the results of [44].

Our second main result considers a different celebrated property of static random graphs: their higher order connectivity. Since the existence of a good analogue for expansion in temporal graphs remains an open problem, we use the maximum doubling time of the reachability ball centered on small sets of vertices as a natural proxy. In our continuous model, when the number of labels per edge is fixed, we provide sharp upper and lower bounds on the maximum doubling time. As a direct consequence, we immediately obtain a sharp threshold for finite maximum doubling time in the Random Simple Temporal Graph (RSTG) model of [15]. Interestingly, this corresponds precisely to the threshold for the existence of temporal source and sink vertices in that model. In contrast, the doubling time of any specific large set of vertices corresponds to the threshold for both point-to-point reachability and the existence of a giant component in the RSTG.

2 Related Work

Most research concerning temporal motifs has focused on either algorithmic techniques for counting or estimating the number of temporal motifs, or the application of these methodologies for the analysis of real networks [54]. A frequent issue in the latter is the construction of a "null" or "reference" model to compare the relative frequency of motifs against [24, 35, 54]. Despite this there has been very limited research into the distribution of motifs in random temporal graphs. A notable exception to this is [51], which derives algorithms for calculating the expectation and the variance of the number of motifs appearing in a proposed random temporal graph model. While closely related, their work differs from ours in several ways. Firstly, they focus primarily on the construction of a model that can be statistically fit to data, and so use a stochastic block model footprint graph with Poisson distributed inter-edge appearance times. This permits them a greater degree of topological richness than is obtained in our simpler model, however the distribution of the labels on each edge is more restricted. Secondly, they focus primarily on providing an algorithm for the calculation of the first and second moment, rather than their application to find existence thresholds. Thirdly, they work only in the continuous setting and do not consider anything comparable to our discrete model. An interesting open problem is closing this gap, in particular providing equivalent results in a unified model capturing both the topological expressivity of the stochastic block model and temporal expressivity of our label multiplicity distribution based approach. While not presented as a result on temporal motifs, [44] investigates the growth of the largest δ-temporal clique, which implicitly provides thresholds for motifs with a clique footprint and a trivial partial order. In this work, we strictly generalize these results to the more powerful random graph models we consider.

Until recently, as discussed in [15], reachability problems in random temporal graphs have primarily been studied in other fields, such as population protocols, the gossip model or the random edge ordering model. For example, in population protocols the reachability properties of the underlying temporal graph underpin two of the model's most fundamental lower bounds: the time complexity of broadcast [47] and leader election [55]. However, recently, in part due to the seminal work of [15], there has been growing interest in the reachability properties of random temporal graphs from both a structural and algorithmic perspective. While other models have been considered (see for example [7, 10, 13]), there has been a focus on establishing the properties of the so-called "Random Simple Temporal Graph" defined as an Erdős-Rényi random graph with a linear order over its edges. Sharp thresholds have since been established at: $\frac{\log n}{n}$ for point-to-point reachability [15] and the existence of a giant [12] (or even super constant [6]) connected-component; $\frac{2 \log n}{n}$ for the existence of temporal sources and sinks; $\frac{3\log n}{n}$ for full temporal connectivity; and $\frac{4\log n}{n}$ for the existence of pivotal (and non-sharply optimal) spanners [15]. Additionally, there have been several results concerning various characteristics of the length of time-respecting paths in the RSTG and their connections to connectivity [5,14]. Our treatment of doubling time shares the most similarity with [12] as, similarly to there, we require bounds on the growth of reachability forests, however rather than using a foremost forest (as in [12,15]) or the random recursive tree (as in [14]), we instead derive the times directly via a simple coupling with a sum of geometric random variables. An additional distinction is our choice of models which are not necessarily simple and have asymptotically constant density. Despite this, when restricted to the interval $[a, a+p]$ with a single label per edge, sampling from our continuous model is equivalent to sampling from the RSTG with parameter p.

3 Model and Definitions

As a natural generalization of an important tool for the analysis of static graphs, some notion of a "temporal motif" has been studied in a variety of contexts and with a variety of definitions (see [54] for a survey). Particular points of distinction relevant to this work include whether an occurrence of a temporal motif must represent an induced subgraph (see [50] and [51] for a difference in opinion) and whether the associated order must be partial [35], or total [29,37,50,51]. We make use of the following definition for (simple) δ-temporal motifs.

▶ **Definition 3** (δ-temporal motif)**.** *Let $H = (V_H, E_H)$ be a finite static graph, $\delta \in \mathbb{R}$ be a number and $P = (E_H, \prec)$ be a partial order over the edges of H. Then a temporal graph $\mathcal{G} = (V, E, \lambda)$ contains (H, P) as a* δ-temporal motif *if and only if there exists a temporal subgraph $\mathcal{I} = (V', E', \lambda')$ and a mapping $\Phi : V_H \to V'$ such that:*

- ***$\mathcal{I}$ is a simple temporal subgraph of $\mathcal{G}$**: $V' \subseteq V$, $E' \subseteq E[V']$ and $\forall e \in E' : \lambda'(e) \in \lambda(e)$*
- ***Φ is an isomorphism of H and realizes the order over the labels**: Φ is an isomorphism of H, $\forall uv, wx \in E_H : uv \prec_P wx \implies \lambda'(\Phi(u)\Phi(v)) < \lambda'(\Phi(w)\Phi(x))$.* [1]
- ***$\mathcal{I}$ has a life duration of at most δ**: $\max_{e_1,e_2 \in E'} \lambda'(e_1) - \lambda'(e_2) < \delta$.*

To avoid trivialities we restrict ourselves to considering motifs containing at least 2 edges, and describe any graph containing 1 or fewer edges as *trivial* [2]. For two graphs, G and H we use $G \subseteq H$ to indicate that G is a (non-empty) subgraph of H and $G \sqsubseteq H$ to indicate that G is a non-trivial subgraph of H.

In this paper we investigate two models of a random temporal graph. The first is a continuous model, while the second is a natural discretization of it. In both models, there are two forms of randomness: the *number of labels* on each edge (which is sampled identically and independently at random from the "label multiplicity distribution") and the *labels themselves* (which are sampled uniformly at random from some interval depending on the model).

▶ **Definition 4.** *We define the* label multiplicity distribution *ψ to be any distribution over the non-negative integers with a positive and finite second moment.*

---

[1] Definition 3 is given under the assumption that P is a strict partial order. Alternatively a δ-temporal motif could be defined with respect to a *non-strict* partial order P, that is, $uv \prec_P wx \implies \lambda'(\Phi(u)\Phi(v)) \leq \lambda'(\Phi(w)\Phi(x))$. With this alternative definition, we believe, but have not verified, that all results of the paper should remain the same, excluding the case where two labels are required to be equal by the order.

[2] We note that the empty graph is a δ-temporal motif of any sufficiently large graph. Furthermore, all simple motifs containing only a single edge appear as δ-temporal motifs with thresholds 0 and 1 in the continuous and discrete models respectively.

For a random variable $Q \sim \psi$, we denote $r = \mathbb{E}[Q]$ and $r_2 = \mathbb{E}[Q^2]$. Furthermore, we denote $q = \mathbb{P}[Q > 0]$. We can now define both our continuous and discrete models.

▶ **Definition 5** (Doubly Random Temporal Graph; continuous case)**.** *We define the distribution $\Gamma_V(\psi)$ of random temporal graphs on the vertex set V with label multiplicity distribution ψ. A sample $\mathcal{G} = (V, E, \lambda)$ from $\Gamma_V(\psi)$ is obtained as follows: For each edge $e \in \binom{V}{2}$, we sample the set of edge occurrences $\lambda(e)$ by drawing $l_e \sim \psi$ samples uniformly at random from $(0, 1]$.*

▶ **Definition 6** (Doubly Random Temporal Graph; discrete case)**.** *We define the distribution $\Gamma_V(\psi, T)$ of random temporal graphs on the vertex set V with label multiplicity distribution ψ and largest possible time-label T. A sample $\mathcal{G} = (V, E, \lambda)$ from $\Gamma_V(\psi, T)$ is obtained as follows: For each edge $e \in \binom{V}{2}$, we sample the set of edge occurrences $\lambda(e)$ by drawing $l_e \sim \psi$ samples uniformly at random from the set $\{1, ..., T\}$.*

Throughout the paper, in the discrete model, we will only consider β-temporal (H, P) motifs, where β is at least height(P), the length of the longest chain in P. The reason for this is that, if β is smaller than the longest chain in P, then trivially a β-temporal motif (H, P) does not exist. There is a natural relationship between the two models, based on the following notion.

▶ **Definition 7.** *For a continuous temporal graph $\mathcal{G} = (V, E, \lambda)$ with all labels restricted to $(0, 1]$, its T-*discretization *is given by $\mathcal{H} = (V, E, \lambda')$ where $\lambda'(e) = \{\lceil Tl \rceil : l \in \lambda(e)\}$.*

We immediately have the following observation connecting the two models.

▶ **Observation 8.** *For all vertex sets V and label multiplicity distributions ψ the distribution $\Gamma_V(\psi, T)$ is identical to the distribution of T-discretizations of a random sample from $\Gamma_V(\psi)$.*

Furthermore, we consider the higher order connectivity of these temporal graphs. Due to the lack of a good definition for the notion of expansion for temporal graphs, we choose to use the *doubling time* of reachability balls as a proxy for expansion.

▶ **Definition 9.** *Let $\mathcal{G} = (V, E, \lambda)$ be a temporal graph and $S \subset V$ be a set of vertices such that $|S| \leq \frac{|V|}{2}$. We define the following notions:*

1. *The t-step reachability ball $\mathcal{B}_{\mathcal{G}}(S, t)$ of S on $\mathcal{G}$ to be all vertices reachable via a time-respecting path initiated from a vertex in S and arriving by time t,*
2. *The doubling time $Double_{\mathcal{G}}(S) = \min\{t \in \mathbb{N}_0 : |\mathcal{B}_{\mathcal{G}}(S, t)| \geq 2|S|\}$ of a set of vertices to be the minimum time required for the reachability ball to double in size. If $\{t \in \mathbb{N}_0 : |\mathcal{B}_{\mathcal{G}}(S, t)| \geq 2|S|\} = \emptyset$, we define $Double_{\mathcal{G}}(S) = \infty$,*
3. *The doubling time $\textbf{Double}(\mathcal{G}) = \max_{S \subset V : |S| \leq \frac{|V|}{2}} Double_{\mathcal{G}}(S)$ of a temporal graph to be the maximum doubling time over all small sets of vertices.*

Finally, we say that an event concerning a temporal graph occurs "with high probability" if the probability of that event approaches 1 as the size of the temporal graph grows.

4 Results

Our first result is on existence thresholds for fixed δ-temporal motifs in the continuous variant of our model.

▶ **Theorem 10.** *Let $H = (V_H, E_H)$ be a non-trivial static graph, $P = (E_H, \prec)$ be any partial order over the edges of H and $\rho_H = \min_{I \sqsubseteq H} \frac{V_I}{E_I - 1}$. A random temporal graph drawn from $\Gamma_{[n]}(\psi)$:*

- *Contains (H, P) as a $\delta(n)$-temporal motif with high probability if $\delta(n) = \omega(n^{-\rho_H})$.*

- *Does not contain* (H,P) *as a* $\delta(n)$*-temporal motif with high probability if* $\delta(n) = o(n^{-\rho_H})$.

For the discrete model, recall that, for any choice of a fixed graph H and a partial order P over its edges, a sample from $\Gamma_{[n]}(\psi, T)$ cannot contain (H,P) as a β-temporal motif, if $\beta < \text{height}(P)$. Therefore we assume below without loss of generality that $\beta \geq \text{height}(P)$. We obtain the following equivalent result for the discrete model.

▶ **Theorem 11.** *Let* $H = (V_H, E_H)$ *be a non-trivial static graph,* $P = (E_H, \prec)$ *be any partial order over the edges of* H *and* $\rho_H = \min_{I \sqsubseteq H} \frac{V_I}{E_I - 1}$. *Suppose that there exists* $n_0 > 0$ *such that* $\min_{n \geq n_0} \beta(n) \geq$ *height*(P). *Then, for* $T(n) \geq \beta(n)$*, a random temporal graph drawn from* $\Gamma_{[n]}(\psi, T(n))$

- *Contains* (H,P) *as a* $\beta(n)$*-temporal motif with high probability if* $\frac{\beta(n)}{T(n)} = \omega(n^{-\rho_H})$.
- *Does not contain* (H,P) *as a* $\beta(n)$*-temporal motif with high probability if* $\frac{\beta(n)}{T(n)} = o(n^{-\rho_H})$.

We also apply our techniques to generalize the previous work of [44] and investigate the growth of the largest δ-temporal clique in the continuous variant of our model. We find that the introduction of more labels produces a subtly different behavior, although still close to the clique number of Erdős-Rényi random graphs.

▶ **Theorem 12.** *Let* K_i *be the clique on* i *vertices and* $P_i^{\emptyset}$ *be the partial order over its edges such that they are all mutually incomparable. Further, let* $0 < \delta < 1$ *be a constant,* $\mathcal{G} = (V, E, \lambda)$ *be a sample from* $\Gamma_{[n]}(\psi)$ *and define* $w_\delta(\mathcal{G})$ *to be the maximum value such that* $\mathcal{G}$ *contains* $(H_{w_\delta(\mathcal{G})}, P^{\emptyset}_{w_\delta(\mathcal{G})})$ *as a* δ*-temporal motif. Then for any* $\epsilon > 0$*, the following holds with high probability,*

$$(1-\epsilon)\frac{2\log n}{\log\left(\frac{\delta r_2 - \delta r + r}{\delta r^2}\right)} \leq w_\delta(\mathcal{G}) \leq (1+\epsilon)\frac{2\log n}{\log\left(\frac{\delta r + (1-\delta)q}{\delta q r}\right)}.$$

Finally, we consider the doubling time of a random sample from $\Gamma_{[n]}(\psi)$ in the special case where ψ is a degenerate distribution, i.e. where every edge has the same number of labels.

▶ **Theorem 13.** *Let* $\mathcal{G} = (V, E, \lambda)$ *be a sample from* $\Gamma_{[n]}(\psi)$*, with* ψ *a degenerate distribution. For any* $\alpha > 0$*,*

$$\mathbb{P}\left[\frac{(2-\alpha)\log n}{rn} \leq \mathbf{Double}(\mathcal{G}) \leq \frac{(2+\alpha)\log n}{rn}\right] = 1 - o(1).$$

We note that when ψ is the degenerate distribution on 1, this statement is equivalent to the following statement about the RSTG model,

▶ **Corollary 14.** *Let* $\mathcal{H} = (V, E, \lambda)$ *be a sample from the RSTG distribution with parameter* $\beta(n)$*. Then the doubling time of* $\mathcal{H}$ *is asymptotically almost surely:*

- *finite, if* $\beta(n) > (2 + \Omega(1))\frac{\log n}{n}$*,*
- *infinite, if* $\beta(n) < (2 - \Omega(1))\frac{\log n}{n}$.

A number of proofs have been omitted from this version due to length constraints. For the detailed proofs of all stated results, please consult the full version [8].

5 Distribution of Temporal Motifs

In this section we outline the proofs of our results (see Theorems 10–12) on the thresholds for the existence of motifs in temporal graphs. The results are obtained via a modification of a first and second moment argument, loosely following the treatment of their static analogues in [23]. Since the proofs of the results concerning both growing motifs and those of constant size share a lot of technical details (simply optimizing for different factors in the associated bounds), we present their shared components together.

5.1 Technical Preliminaries

Before we present our techniques, we will need a few preliminary results and definitions. Our first preliminary is a short definition required for defining the probability of a given collection of labels obeying the partial order. Here it is important to distinguish between the continuous case, where labels are totally ordered almost surely, and the discrete case, where multiple labels may share the same time stamp.

▶ **Definition 15.** *For a poset $P = (A, \prec)$ and $\beta \in [n]$, we take $R(P)$ to be the number of linear extensions of P and $R(P, \beta)$ to be the number of mappings from A to $[\beta]$ such that P is satisfied on $[\beta]$.*

We will also need the following notion of automorphism for temporal motifs. In particular, we require that the mapping over the edges produced by the automorphism of the vertices corresponds to an automorphism of the associated partial order. The number of such automorphisms will play a crucial role in preventing the over-counting of the number of occurrences of temporal motifs when multiple occurrences share the same labels.

▶ **Definition 16.** *For a fixed graph $H = (V, E)$ and a partial order $P = (E, \prec)$, $\phi : V \to V$ is an automorphism of (H, P) if ϕ is an automorphism of H preserving the ordering of P. Explicitly, we require that (i) for all u, v we have $\{u, v\} \in E_H \iff \{\phi(u), \phi(v)\} \in E_H$ and (ii) for all pairs of edges $\{u, v\}, \{w, x\} \in \binom{V}{2}$ we have $\{u, v\} \prec \{w, x\} \iff \{\phi(u), \phi(v)\} \prec \{\phi(w), \phi(x)\}$. We denote the number of such automorphisms by $\#Aut(H, P)$.*

Additionally, we require the following simple domination result, which relates the probability of a temporal motif appearing in two instances of our random graph models based on stochastic domination between their label multiplicity distributions. This will serve as a useful tool which allows us to reduce the problem of obtaining certain bounds for arbitrary label multiplicity distributions to the case of bounded label multiplicity distributions.

▶ **Lemma 17.** *Let $I = (V_I, E_I)$ be a fixed graph, $P = (E_I, \prec)$ a partial order over its edges and ψ and χ two valid label multiplicity distributions, such that ψ is stochastically dominated by χ. Further let $\mathcal{G} = (V, E, \lambda)$ be a temporal graph sampled from $\Gamma_V(\psi)$ and $\mathcal{H} = (V, E, \lambda')$ be a temporal graph sampled from $\Gamma_V(\chi)$. Then,*

- *For any $0 < \delta < 1$, the probability (I, P) appears as a δ-temporal motif in $\mathcal{G}$ is at most the probability that (I, P) appears as a δ-temporal motif in $\mathcal{H}$.*
- *For any $0 < \beta < T$, the probability (I, P) appears as a β-temporal motif in the T discretization of $\mathcal{G}$ is at most the probability that (I, P) appears as a β-temporal motif in the T discretization of $\mathcal{H}$.*

We can now construct our random variables of interest. We define $\mathcal{S}$ to be the set of all occurrences of H in the footprint graph $G = (V, E)$ and for any $S \in \mathcal{S}$ we define $\mathcal{T}_S$ to be the set of all mappings from $E_S \to \mathbb{N}_0$. For mathematical convenience we shall treat $\lambda(e)$ as a tuple ordered by the order of sampling rather than a multi-set and we shall assume that for each edge the value of all labels are sampled first, only later followed by how many will actually appear. This allows us to interpret any pair (S, τ) for $S \in \mathcal{S}$ and $\tau \in \mathcal{T}_S$, as a specific candidate for a motif, namely S where each edge $e \in E_S$ has its $\tau(e)$th sampled label. Then, for a random sample $\mathcal{G}$ from $\Gamma_{[n]}(\psi)$, $S \in \mathcal{S}$ and $\tau \in \mathcal{T}_S$, we define $C_{S,\tau}^{(\delta)}$ to be the indicator random variable for the event that (S, τ) describes a δ-temporal (H, P) motif of $\mathcal{G}$. Additionally, $\mathbf{C}_\delta = \sum_{S \in \mathcal{S}, \tau \in \mathcal{T}_S} C_{S,\tau}^{(\delta)}$ corresponds to the total number of appearances of (H, P) as a δ-temporal motif of $\mathcal{G}$. Analogously, for a random sample $\mathcal{H}$ from $\Gamma_{[n]}(\psi, T)$, which we obtain by taking the T-discretization of $\mathcal{G}$, we define $F_{S,\tau}^{(\beta),T}$ to be the indicator for the event that (S,τ) describes a β-temporal (H,P) motif of $\mathcal{H}$. Finally, $\mathbf{F}_{(\beta),T} = \sum_{S\in\mathcal{S},\tau\in\mathcal{T}_S} F^{(\beta),T}_{S,\tau}$ corresponds to the total number of appearances of (H,P) as a β-temporal motif of $\mathcal{H}$.

We now make the following observation relating the two models.

▶ **Observation 18.** *For $\mathcal{G} = (V,E,\lambda)$, $T > \beta \geq 1$, $\mathcal{H}$ the T discretization of $\mathcal{G}$, $H = (V_H, E_H)$ any static graph, $P = (E_H, \prec)$ a partial order over the edges of H, and for every $S \in \mathcal{S}$, $\tau \in \mathcal{T}_S$, we have that $\{F^{(\beta),T}_{S,\tau} = 1\} \subseteq \{C^{(\frac{\beta}{T})}_{S,\tau} = 1\}$. Therefore, with certainty, $\mathbf{F}_{(\beta),T} \leq \mathbf{C}_{\frac{\beta}{T}}$.*

5.1.1 First Moment

Since directly calculating the probability of a temporal motif occurring in a random temporal graph is infeasible, we shall rely on bounds derived from the first moment of the number of such occurrences. In fact, we are able to obtain the following lemma via a relatively standard calculation, exploiting both the linearity of expectation and the independence between the number of labels on each edge and their value. Note that the term concerning the number of realizations of P is absent in the case for $\beta = 1$ as there exists only a single possible realization (all labels in a single snapshot).

▶ **Lemma 19.** *The expectation of $\mathbf{C}_\delta$ is:*

$$\mathbb{E}[\mathbf{C}_\delta] = \delta^{|E_H|} r^{|E_H|} |V_H|! \left(\frac{R(P)}{|E_H|!}\right) \binom{n}{|V_H|} \left(1 + \frac{|E_H|(1-\delta)}{\delta}\right).$$

Similarly, if $\beta = 1$ then $\mathbb{E}[\mathbf{F}_{(1),T}] = r^{|E_H|} T^{1-|E_H|} |V_H|! \binom{n}{|V_H|}$.
On the other hand, if $\beta > 1$ then $\mathbb{E}[\mathbf{F}_{(\beta),T}] \geq \left(\frac{\beta r}{2|E_H|T}\right)^{|E_H|} |V_H|! \left(1 + \frac{|E_H|(1-\frac{\beta}{2T})}{\frac{\beta}{2T}}\right) \binom{n}{|V_H|}$.

With that established, we may turn our attention to the construction of the actual bound. The observant reader may notice that the expectations obtained above will tend to zero whenever $\delta(n) = o(n^{-\rho_H})$ and $|V_H|$ is fixed. In fact, Lemma 19 is sufficient to show the lower bounds of Theorem 10 and Theorem 11 via a straightforward application of the first moment method. Unfortunately, the same cannot be said for Theorem 12. Indeed, even a cursory inspection of the problem will reveal that $C^{(\delta)}_{S_1,\tau_1}$ and $C^{(\delta)}_{S_2,\tau_2}$ have a strong positive association whenever τ_1 and τ_2 share labels or even when S_1 and S_2 share a single edge. Consequently, the first moment produces an extremely weak bound that explodes whenever $r > 1$ and the size of H is unbounded. In order to control these dependencies and limit the contribution of additional occurrences after the first, we make use of a slightly more sophisticated bound based on the ratio of the expectation and the expectation conditioned on being non-zero.

▶ **Lemma 20.** *Let $H = (V_H, E_H)$ be a fixed graph, $P = (E_H, \prec)$, $\mathcal{G} = (V,E,\lambda)$ be a sample from $\Gamma_{[n]}(\psi)$ and $0 < \delta < 1$. Then,*

$$\mathbb{P}[\mathbf{C}_\delta > 0] \leq \binom{n}{|V_H|} \frac{|V_H|!}{\#Aut(H,P)} \left(\frac{\delta r q}{(1-\delta)q + \delta r}\right)^{|E_H|} \left(1 + \frac{|E_H|(1-\delta)}{\delta}\right).$$

5.1.2 Second Moment

Once again we wish to avoid a direct calculation of a lower bound on the probability of a given motif occurring, and so our strategy will be to use the second moment method. Unfortunately, we do not find a unified overarching bound sufficient to obtain our results for both fixed and growing motifs, as positive association rears its head once more (although this time as an obstacle for the fixed case). In particular, when two potential occurrences of a motif share a footprint they remain dependent even if they do not share any labels, due to the inherent randomness of the label multiplicity distribution. This presents an issue whenever ψ is a non-degenerate non-Bernoulli distribution, as we will retain terms in our summations corresponding to such pairs. These terms contribute an additional δ factor to our calculations, which can be absorbed in the case of Theorem 12 where δ is constant but a naive application of the second moment would lead to asymptotically weaker bounds in the case of Theorem 10 and Theorem 11 where δ is decreasing. In order to avoid this pitfall, we shall make use of two bounds, one for the case where ψ is a Bernoulli random variable and one for the more general but still bounded case. We shall warm up with the considerably easier former case.

▶ **Lemma 21.** *Let $\mathcal{G} = (V, E, \lambda)$ be a sample from $\Gamma_{[n]}(\psi)$, for ψ a Bernoulli random variable, then for any $0 < \delta < 1$,*

$$Var[\mathbf{C}_\delta] \leq (\sqrt{2}\delta r)^{2|E_H|}\left(1 + |E_H|\frac{(1-\delta)}{\delta}\right) n^{2|V_H|}(2|V_H|)^{2|V_H|}\sum_{I\subseteq H}(2\delta r)^{-|E_I|}n^{-|V_I|}.$$

Similarly for any $\beta > 0$, $T > \beta$,

$$Var\left[\mathbf{F}_{(\beta),T}\right] \leq (\frac{\sqrt{2}\beta r}{T})^{2|E_H|}\left(1 + |E_H|\frac{(T-\beta)}{\beta}\right) n^{2|V_H|}(2|V_H|)^{2|V_H|}\sum_{I\subseteq H}(\frac{2\beta r}{T})^{-|E_I|}n^{-|V_I|}.$$

With that established, we can present the following more general result, in which we have to be much more careful to control the dependence caused by the label multiplicity distribution. Consequently, the following lemma is restricted to the case where ψ is almost surely bounded.

▶ **Lemma 22.** *Let $\mathcal{G} = (V, E, \lambda)$ be a sample from $\Gamma_{[n]}(\psi)$, where ψ is almost surely bounded, then for any $0 < \delta < 1$,*

$$Var[\mathbf{C}_\delta] \leq \delta^{2|E_H|}r^{2|E_H|}\left(1 + \frac{|E_H|(1-\delta)}{\delta}\right)^2 \sum_{S_1,S_2\in\mathcal{S}:E_{S_1}\cap E_{S_2}\neq\emptyset}\left(\frac{\delta r_2 + r - \delta r}{\delta r^2}\right)^{|E_{S_1}\cap E_{S_2}|}.$$

5.2 Existence Thresholds

We can now turn our attention to the existence thresholds themselves, where we are able to make use of similar approaches for both the fixed and growing motifs. For the continuous model, both the lower bound on the existence threshold for fixed motifs (Theorem 10) and the upper bound on the size of the largest unordered clique motif (Theorem 12) follow from the asymptotics of Lemma 20. Interestingly, this also suffices for the lower bound on the existence threshold for fixed motifs in the discrete case (Theorem 11), as, by Observation 18, the number of occurrences of (H, P) as a $\frac{\beta}{T}$-temporal motif in a sample from $\Gamma_{[n]}(\psi)$ is at least the number of occurrences of (H, P) as a β-temporal motif in the sample's T discretization.

The opposing bounds are somewhat more challenging. In both cases, we are only able to obtain a weak variant of our desired result via a naive application of the second moment method. For fixed motifs, we inherit the restriction of ψ to Bernoulli distributions from Lemma 21, and for the largest unordered clique, we similarly find ourselves restricted to almost surely bounded label multiplicity distributions. However, these restricted versions actually turn out to be sufficient for our purposes. The key idea is to apply Lemma 17 and exploit the relationship between samples from multiple instances of our model with different (and stochastically ordered) label multiplicity distributions. On the one hand, in the case of fixed motifs, observe that any non-negative non-degenerate discrete distribution ψ over $\mathbb{N}$ stochastically dominates some Bernoulli distribution σ. Since we know from Lemma 17 that the number of instances of a temporal motif in two coupled samples from the continuous (resp. discrete) model with different label multiplicity distributions inherit the stochastic order of their corresponding label multiplicity distributions, it follows that the probability of a given temporal motif existing in the sample from $\Gamma_{[n]}(\psi)$ (resp. $\Gamma_{[n]}(\psi, T)$) is lower bounded by the probability of one existing in the sample from $\Gamma_{[n]}(\sigma)$ (resp. $\Gamma_{[n]}(\sigma, T)$). Combining this with the bound from the second moment method, even when restricted only to Bernoulli label multiplicity distributions, is then sufficient to obtain our threshold. On the other hand, we have to be a bit more careful in the case of the largest unordered clique, as the error introduced from the discrepancy between the actual label multiplicity distribution and a Bernoulli lower bound is too high. Fortunately, for any valid choice ψ of label multiplicity distribution, we can find a second almost surely bounded distribution χ such that ψ stochastically dominates χ and χ has first and second moments arbitrarily close to those of ψ. In particular, for a suitable choice of χ, we are able to achieve the same bound in Theorem 12 for both samples from $\Gamma_n(\psi)$ and $\Gamma_{[n]}(\chi)$. The result then follows from the same stochastic order argument as for fixed motifs.

6 Doubling Times

In this section, we outline the ideas behind Theorem 13 characterizing the doubling time of $\Gamma_{[n]}(\psi)$ whenever ψ is a degenerate random variable. The proof largely consists of obtaining probabilistic bounds on the doubling time of small individual sets followed by a more careful analysis of the particularly early and late participants. This approach is inspired by [15], although as we deal with relatively large sets of vertices rather than individual vertices, we have to put in a lot of extra-leg work in the process. Our main strategy for obtaining bounds on the doubling time of any specific set is to bound the time required for reachability balls to cross a sequence of cuts in the underlying graph. This introduces a great deal of dependency, which we are fortunately able to control (see the full version for more details).

For the sake of mathematical convenience, we first establish the results for a slightly different random temporal graph model and then show how they can be extended to our original setting. The model in question is a different discretization to the one we have considered thus far: the so-called "order discretization". In constructing the order discretization of $\mathcal{G}$, we strip away all information about the precise timing of edge labels and instead only retain the relative order between them. The order discretization of $\mathcal{G}$ contains only a single edge at any given time and the edge at time k corresponds to the kth edge to appear in $\mathcal{G}$. Formally, it is defined as follows.

▶ **Definition 23.** *For a temporal graph $\mathcal{G} = (V, E, \lambda)$ where no two edges share the same label, we define $\mathcal{G}^O = (V, E, \lambda^O)$ the* order discretization *of $\mathcal{G}$ such that $\lambda^O(e)$ corresponds to the set of relative positions of each label in $\lambda(e)$ in the overall set of labels. Furthermore, for $0 \leq a < b$ we define $\mathcal{G}^O_{[a,b]} = (V, E_{[a,b]}, \lambda^O_{[a,b]})$ to be $\mathcal{G}^O$ with all labels outside of $[a, b]$ removed.*

Since the number of labels falling within a given interval of a sample from $\Gamma_{[n]}(\psi)$ is highly concentrated, we are easily able to convert between times in a sample from $\Gamma_{[n]}(\psi)$ and its order discretization.

6.1 Upper bounds

We will begin with our argument for obtaining an upper bound on the doubling time. However, first, we will need a few preliminary ideas. A key property of our model is that it is in some sense "time reversible" as the probability of any set of edge labels appearing in a given order is the same as that of them falling in the reverse (or in fact any) order. This provides us with a very useful tool as we are able to consider our normal temporal graph and variants of it where we have partially reversed its labeling interchangeably. The next definition and lemma (see also Definition 9) exploits this property to reduce bounding the doubling time of small sets of vertices to bounding the doubling time of large sets of vertices in the partially reversed graph.

▶ **Definition 24.** *For any temporal graph $\mathcal{G} = (V, E, \lambda)$, we define:*

- $\textbf{\textit{Double}}^-(\mathcal{G}) = \max_{S \subset V : 2 \leq |S| \leq \frac{n}{4}} (\textit{Double}_{\mathcal{G}}(S))$ *to be the doubling time over small sets.*
- $\textbf{\textit{Double}}^+(\mathcal{G}) = \max_{S \subset V : \frac{n}{4} < |S| \leq \frac{n}{2}} (\textit{Double}_{\mathcal{G}}(S))$ *to be the doubling time over large sets.*
- $\mathcal{R}(\mathcal{G}, t) = (V, E, \lambda_{\mathcal{R}(\mathcal{G},t)})$ *to be the partially reversed temporal graph obtained by reversing the occurrence times of all labels before t, i.e. taking $\lambda_{\mathcal{R}(\mathcal{G},t)}(e) = \{R_t(l) : l \in \lambda(e)\}$, where* $R_t(l) = \begin{cases} t - l & \text{if } t > l \\ l & \text{otherwise} \end{cases}$.

▶ **Lemma 25.** *For any temporal graph $\mathcal{G} = (V, E, \lambda)$, $\textbf{\textit{Double}}^-(\mathcal{R}(\mathcal{G}, \textbf{\textit{Double}}^+(\mathcal{G})) \leq \textbf{\textit{Double}}^+(\mathcal{G})$.*

Proof. For ease of notation we will denote $D = \mathbf{Double}^+(\mathcal{G})$ and $\mathcal{H} = \mathcal{R}(\mathcal{G}, \mathbf{Double}^+(\mathcal{G}))$. Now for the sake of contradiction, assume that there exists a set S with $2 \leq |S| \leq \frac{n}{4}$ such that $\text{Double}_{\mathcal{H}}(S) \geq D$. Consider that for any two vertices u and v if there exists a path from u to v in $\mathcal{G}$ arriving before time D there must exist a path from v to u in $\mathcal{H}$ arriving by time D. Thus, we must have that T the set of vertices outside of S that can reach S in $\mathcal{G}$ before time D has size at most $|S| - 1$. Furthermore, since $|S| \leq \frac{n}{4}$, we must have that $|S \cup T| \leq \frac{n}{2}$ and so we can take $U \subset V$ such that $|U| = \lceil \frac{n}{2} \rceil$ and $S \cup T \subseteq U$. Now by definition $V \setminus U$ has doubling time at most D in $\mathcal{G}$. Since $V \setminus U$ has size $\lfloor \frac{n}{2} \rfloor$, this implies that at most one vertex cannot be reached from $V \setminus U$ via a time-respecting path before time D in $\mathcal{G}$. So, in particular, there must be at least one vertex from $V \setminus U$ with a path to some member of S in $\mathcal{G}$ arriving by time D. This contradicts the definition of T and so no such set S can exist. Therefore every set with size at least 2 and at most $\frac{n}{4}$ must have a doubling time in $\mathcal{H}$ of at most D. ◀

While stated for arbitrary temporal graphs, in our setting, given the equal probability of sampling both $\mathcal{G}$ and $\mathcal{R}(\mathcal{G}, \mathbf{Double}^+(\mathcal{G}))$, this turns any upper bound on the doubling time of large sets into an upper bound on the doubling time of small sets. Our second key preliminary is an upper bound on the time taken for the reachability ball centered on any specific set of vertices to grow to a certain size, obtained from a carefully constructed coupling between the total time between expansions of the reachability ball and a sum of geometric random variables. We defer its precise (and somewhat involved) statement to the full version [8] and will refer to it here as our *single-set growth upper bound*.

With these tools established, we may give our upper bound on the doubling time in $\mathcal{G}^O$.

▶ **Lemma 26.** *Let $\mathcal{G} = (V, E, \lambda)$ be a sample from $\Gamma_{[n]}(\psi)$, for ψ a degenerate distribution. Then for any $\beta > 0$, we have that $\textbf{\textit{Double}}(\mathcal{G}^O) \leq (1 + \beta) n \log n$.*

The proof of this result consists largely of applying the single-set growth upper bound, along with a concentration inequality and Lemma 25 in order to deal with the case of small sets. More specifically, observe that once all vertices can be reached by at least $\frac{3n}{4}$ other vertices, every set of size at least $\frac{n}{4}$ can reach all vertices. Thus, once this has occurred the reachability ball of all large sets of vertices must have doubled in size. In order to bound this time, we first show via an application of the single-set growth upper bound and Markov's inequality that almost all vertices can be reached by $\frac{3n}{4}$ other vertices by roughly time $n \log n$. From a second application of the single-set growth upper bound, we find that the few remaining vertices are reached from the already well connected vertices within only approximately a further $n \log n$ rounds. The more general result then follows from using Lemma 25 to extend the bound to cover all smaller sets of vertices.

Proof. We shall make use of the following notation. Denote by $\mathbf{H}(a,b) = H_{n-a} + H_{b-1} - H_{a-1} - H_{n-b}$, where H_a is the ath harmonic number with the convention that $H_0 = 0$. Further, denote by T the first time step such that all vertices can be reached by at least $\frac{3n}{4}$ other vertices. For a fixed $\gamma, \epsilon > 0$ small, denote by X the number of vertices that are not reachable by $\frac{3n}{4}$ other vertices by time $T_1 = \frac{(1+\epsilon)n\mathbf{H}(1,\frac{3n}{4}+1)}{2(1-\gamma)}$. It follows from our single-set growth upper bound and the linearity of expectation that,

$$\mathbb{E}[X] \leq n \cdot \left(\exp\left[-\frac{n-1}{n}\mathbf{H}(1, \frac{3n+4}{4})(\epsilon - \log(1+\epsilon))\right] + o(n^{-1})\right).$$

For n sufficiently large and ϵ sufficiently small,

$$\leq n^{1-\epsilon+\log(1+\epsilon)}.$$

Now applying Markov's inequality, we have that,

$$\mathbb{P}[X > n^{1-\epsilon+\log 1+\epsilon} \log n] \leq \frac{1}{\log n}.$$

Thus with probability $1 - \frac{1}{\log n}$ we have a set of at least $n - n^{1-\epsilon+\log 1+\epsilon} \log n$ vertices that can be reached by at least $\frac{3n}{4}$ other vertices which we shall denote by A. Trivially, T must occur at the very latest when every member of $V \setminus A$ is reachable by a member of A via a path starting after T_1. Since $T_1 = O(n \log n)$, we know from a second application of our single-set growth upper bound that this occurs at the very latest by $T_1 + \frac{(1+\epsilon)n\mathbf{H}(n-n^{1-\epsilon+\log 1+\epsilon}\log n, n)}{2(1-\gamma)}$ with probability $1 - o(1)$. Taking a union bound over the probability of failures we get that with probability $1 - o(1)$,

$$T < \frac{(1+\epsilon)n}{2(1-\gamma)}\left(\mathbf{H}(1, \frac{3n}{4}+1) + \mathbf{H}(n - n^{1-\epsilon+\log 1+\epsilon}\log n, n)\right)$$

$$= \frac{(1+\epsilon)n}{2(1-\gamma)}\left(H_{n-1} + H_{\frac{3n}{4}} - H_{\frac{n}{4}-1} + H_{n^{1-\epsilon+\log 1+\epsilon}\log n} + H_{n-1} - H_{n-n^{1-\epsilon+\log 1+\epsilon}\log n}\right).$$

For any $\alpha > 0$ small, and all n large enough, we can apply a folklore approximation of the harmonic series (see the full version [8] for the precise statement),

$$\leq \frac{(1+\epsilon)n\log n}{2(1-\gamma)}\left(\left(4 - \epsilon + \log(1+\epsilon) + \frac{\log\log n}{\log n}\right)(1+\alpha) - 2(1-\alpha)\right)$$

$$= \frac{(1+\epsilon)n\log n}{2(1-\gamma)}\left(2 + 2\alpha + (1+\alpha)\left(-\epsilon + \log(1+\epsilon) + \frac{\log\log n}{\log n}\right)\right).$$

For any $\beta > 0$, taking $\alpha = \gamma = \epsilon > 0$ sufficiently small, we find,

$$\leq (1+\beta)n \log n.$$

Thus, we find that, with probability $1 - o(1)$, all vertices can be reached by at least $\frac{3n}{4}$ other vertices by time $(1+\beta)n\log n$. Therefore, we must have that by time $(1+\beta)n\log n$ all sets of size at least $\frac{n}{4}$ can reach all vertices. Thus we have that with probability $1-o(1)$, $\mathbf{Double}^+(\mathcal{G}) \leq (1+\beta)n\log n$ and by the time reversible nature of the distribution that $\mathbf{Double}^+(\mathcal{R}(\mathcal{G}, (1+\beta)n\log n)) \leq (1+\beta)n\log n$ with the same probability. In which case it follows from Lemma 25 that $\mathbf{Double}^-(\mathcal{G}) \leq (1+\beta)n\log n$. Thus taking a union bound over the failure probabilities, we find that $\mathbf{Double}(\mathcal{G}) \leq (1+\beta)n\log n$ ◀

6.2 Lower Bounds

In this section we present a lower bound to match the upper bound of Lemma 26. However, this time, our lives will be significantly easier, as it suffices to consider only the doubling time of sets containing precisely $\frac{n}{2}$ vertices. As before, we need a bound on the doubling time of any specific set, and so make use of the following result analogous to our single-set growth upper bound, the proof of which we also defer.

▶ **Lemma 27.** *Let $\mathcal{G} = (V, E, \lambda)$ be a sample from $\Gamma_{[n]}(\psi)$, for ψ a degenerate distribution, and let $S \subseteq V$ be a set of vertices such that $|S| = \frac{n}{2}$. Then for any $\alpha, \beta > 0$ sufficiently small we have that,*

$$\mathbb{P}\left[Double_{\mathcal{G}^O}(S) < \frac{(1-\alpha)n\log n}{2(1+\beta)}\right] = O(n^{\frac{\log 2 - 1}{5}}) + O(n^{2(1-\beta)(\alpha + \log(1-\alpha))}).$$

With that established we may now give the main result of the subsection, a lower bound on the doubling time of $\mathcal{G}^O$.

▶ **Lemma 28.** *Let $\mathcal{G} = (V, E, \lambda)$ be a sample from $\Gamma_{[n]}(\psi)$, with ψ a degenerate distribution, then $\boldsymbol{Double}(\mathcal{G}^O) \geq (1-\epsilon)n\log n$.*

The key notion for the proof of this result is the idea of a "blocker vertex", i.e. a vertex who is not yet reachable by any member of some large set of vertices. By applying Lemma 27 and Markov's inequality, followed by a simple counting argument, we are able to show that at least a constant fraction of vertices must be blockers around time $\frac{n\log n}{2}$. Then, via drift analysis, we are able to show that at least one blocker vertex will not be included in an edge before time $(1-\epsilon)n\log n$. Given the fact that the doubling time is lower bounded by the time when the last vertex stops being a blocker, and that a vertex must be included in an edge in order to stop being a blocker, we immediately obtain the bound.

Proof. Let X be the number of sets of vertices of size $\frac{n}{2}$ which have a doubling time of at most $\frac{(1-\alpha)n\log n}{2(1+\alpha)}$ for some $\alpha > 0$ small. We know from Lemma 27 (where the $O(n^{\frac{\log 2-1}{5}})$ vanishes into the term depending on α) and the linearity of expectation that for some constant $c > 0$,

$$\mathbb{E}[X] \leq cn^{2(1-\alpha)(\alpha+\log(1-\alpha))} \cdot \binom{n}{\frac{n}{2}}.$$

Therefore by Markov's inequality it holds that,

$$\mathbb{P}\left[X \geq cn^{2(1-\alpha)(\alpha+\log(1-\alpha))}\log n \cdot \binom{n}{\frac{n}{2}}\right] \leq \frac{1}{\log n}.$$

Therefore with probability $1 - \frac{1}{\log n}$ there are at least $(1 - cn^{(1-\alpha)(\alpha+\log(1-\alpha))})\binom{n}{\frac{n}{2}}$ sets with a doubling time of at least $\frac{(1-\alpha)n\log n}{2(1+\alpha)}$.

Now, a set of size $\frac{n}{2}$ has not yet doubled its reachability ball, if and only if, there is some vertex outside of the set that it cannot yet reach. We call such an unreachable vertex a "blocker" vertex for that set. Furthermore, any individual blocker vertex can block at most $\binom{n-1}{\frac{n}{2}} = \frac{1}{n}\binom{n}{\frac{n}{2}}$ sets (precisely those that do not contain it). Since we know that with high probability there are $(1 - n^{-\beta})\binom{n}{\frac{n}{2}}$ sets which are blocked for some $\beta > 0$, for any constant $0 < d < 1$ it holds that there must be at least dn blocker vertices.

In order for a set to have doubled its reachability ball, it must have no more blocker vertices and a blocker vertex can only stop being a blocker when it is included in an edge. What we will now show is that, with high probability, at least one blocker vertex is not included in an edge between $\frac{(1-\alpha)}{2(1+\alpha)}n\log n$ and $(1-\epsilon)n\log n$.

▷ Claim 29. There exists at least one blocker vertex that does not receive an edge between time $\frac{(1-\alpha)}{2(1+\alpha)}n\log n$ and $(1-\epsilon)n\log n$ with probability $1 - o(1)$.

Proof. Let B be a set of dn vertices chosen arbitrarily at time $\frac{(1-\alpha)}{2(1+\alpha)}n\log n$ to either be a subset or super-set of the set of the currently blocking vertices. Let Y_t be the set of vertices from B that have not yet received an edge between time $\frac{(1-\alpha)}{2(1+\alpha)}n\log n$ and $\frac{(1-\alpha)}{2(1+\alpha)}n\log n + t$, and $\mathcal{F}_t$ be its corresponding filtration. We have the following two bounds, for $t < 2n\log n$ and any fixed $\gamma > 0$ immediately,

$$\mathbb{P}[Y_t - Y_{t+1} = 1 | \mathcal{F}_t, Y_t = y \geq 1] \leq \frac{2y(1+\gamma)(n-y)}{n(n-1)},$$

$$\mathbb{P}[Y_t - Y_{t+1} = 2 | \mathcal{F}_t, Y_t = y \geq 2] \leq \frac{y(1+\gamma)(y-1)}{n(n-1)}.$$

We define $T_{\text{connect}} = \inf\{t \geq 0 : Y_t \leq 2\}$, take $\zeta = 1$ and $g(y) = \log y + 1$, then

$$\mathbb{E}\left[e^{\zeta(g(Y_t)-g(Y_{t+1}))} | \mathcal{F}_t, Y_t = y \geq 3\right] = \mathbb{P}[Y_t - Y_{t+1} = 0 | \mathcal{F}_t, Y_t = y \geq 3] \quad (1)$$

$$+ \mathbb{P}[Y_t - Y_{t+1} = 1 | \mathcal{F}_t, Y_t = y \geq 3]\frac{y+1}{y} \quad (2)$$

$$+ \mathbb{P}[Y_t - Y_{t+1} = 2 | \mathcal{F}_t, Y_t = y \geq 3]\frac{y+1}{y-1}. \quad (3)$$

Since the probabilities must sum to one, and the coefficient of the term corresponding to the event with no change is minimal, we obtain for all $t < n\log n$,

$$\leq 1 - \frac{2y(1+\gamma)(n-y)}{n(n-1)} - \frac{y(1+\gamma)(y-1)}{n(n-1)} + \frac{2(y+1)(1+\gamma)(n-y)}{n(n-1)} + \frac{(1+\gamma)(y+1)y}{n(n-1)}$$

$$\leq 1 + \frac{2(1+2\gamma)}{n}.$$

We now apply Theorem 2 of [36], to obtain that, for any $\eta > 0$ small,

$$\mathbb{P}\left[T_{\text{connect}} < \frac{(1-\eta)n\log n}{2}\right] \leq \left(1 + \frac{2(1+2\gamma)}{n}\right)^{\frac{(1-\eta)n\log n}{2}} \cdot e^{\log 3 - \log dn + 1}$$

$$\leq \frac{3}{d}n^{(1-\eta)(1+2\gamma)-1}$$

$$\leq \frac{3}{d}n^{2\gamma-\eta-2\eta\gamma}.$$

Taking $\gamma < \frac{\eta}{2}$ gives us that with probability $1 - o(1)$ there is at least one member of B, that has not been included in an edge after $\frac{(1-\alpha)}{2(1+\alpha)} n \log n$ and before $\frac{(1-\alpha)}{2(1+\alpha)} n \log n + \frac{1-\eta}{2} n \log n$. However, we also know that with probability $1 - o(1)$ there are at least dn blocker vertices. Therefore, by construction, with probability $1 - o(1)$, B only contains blocker vertices. A union bound over the probability of failure gives the claim. ◀

The main claim follows immediately from this, the consolidation of constant terms and the observation that if there is at least one blocker vertex that has not received an edge, there must be at least one set that has not yet doubled its size. ◀

6.3 Tidying Up

In this section, we prove our main theorem on the doubling time in the continuous model. In order to achieve this we will convert our upper and lower bounds on the order discretization of a sample from the continuous model back to the undiscretized version. This is made quite easy by the strong concentration of the number of labels occurring within a given interval. In fact, since ψ is degenerate we obtain the ever pleasant binomial distribution. We are able, therefore, to obtain the following simple precursor, showing that with high probability the number of labels appearing in the initial interval is highly concentrated.

▶ **Lemma 30.** *For $\alpha, \beta > 0$ with β small, the following holds, for any $\mathcal{G} = (V, E, \lambda)$,*

$$\mathbb{P}\left[\sum_{e \in E} |\lambda^O_{[0, \frac{\alpha \log n}{r(n-1)}]}(e)| \leq \frac{\alpha(1-\beta) n \log n}{2}\right] = o(n^{-1}).$$

Similarly,

$$\mathbb{P}\left[\sum_{e \in E} |\lambda^O_{[0, \frac{\alpha \log n}{r(n-1)}]}(e)| \geq \frac{\alpha(1+\beta) n \log n}{2}\right] = o(n^{-1}).$$

The proof of the result itself is then simply an application of Lemma 30 to the bounds implied by Lemma 28 and Lemma 26.

Proof of Theorem 13. It follows from Lemma 26, Lemma 28 and a union bound that for any $\epsilon > 0$,

$$\mathbb{P}[(1-\epsilon) n \log n < \mathbf{Double}(\mathcal{G}^O) < (1+\epsilon) n \log n] = 1 - o(1).$$

Furthermore, we know from Lemma 30 that for any $\delta > 0$ the $(1-\epsilon) n \log n$th label falls after $\frac{2(1-\delta)(1-\epsilon)\log n}{r(n-1)}$ and that the $(1+\epsilon) n \log n$th labels falls before $\frac{2(1+\delta)(1+\epsilon)\log n}{r(n-1)}$ with probability $1 - o(1)$. Thus taking a union bound over the failure probabilities, and $\alpha > \delta + \epsilon + \delta\epsilon$, we obtain the claim. ◀

References

1. Duncan Adamson, George B. Mertzios, and Paul G. Spirakis. Maintaining bipartite colourings on temporal graphs on a budget. In *Proeedings of the 33rd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 2026. To appear.
2. Eleni C. Akrida, Leszek Gasieniec, George B. Mertzios, and Paul G. Spirakis. Ephemeral networks with random availability of links: The case of fast networks. *Journal of Parallel and Distributed Computing*, 87:109–120, 2016.

3 Eleni C. Akrida, Leszek Gasieniec, George B. Mertzios, and Paul G. Spirakis. The complexity of optimal design of temporally connected graphs. *Theory of Computing Systems*, 61(3):907–944, 2017.

4 Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis, and Viktor Zamaraev. Temporal vertex cover with a sliding time window. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107, pages 148:1–148:14, 2018. Journal version appeared at JCSS in 2020.

5 Omer Angel, Asaf Ferber, Benny Sudakov, and Vincent Tassion. Long monotone trails in random edge-labellings of random graphs. *Combinatorics, Probability and Computing*, 29(1):22–30, 2020.

6 Caelan Atamanchuk, Luc Devroye, and Gábor Lugosi. On the size of temporal cliques in subcritical random temporal graphs. *Combinatorics, Probability and Computing*, pages 1–9, 2024.

7 Caelan Atamanchuk, Luc Devroye, and Gábor Lugosi. Uniform temporal trees. *Random Structures & Algorithms*, 67(4):e70040, 2025.

8 Henry Austin, George B. Mertzios, and Paul G. Spirakis. Sharp thresholds for temporal motifs and doubling time in random temporal graphs. *CoRR*, 2026. URL: `http://arxiv.org/abs/2602.01847`.

9 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)*, pages 149:1–149:14, 2016.

10 Samuel Baguley, Andreas Göbel, Nicolas Klodt, George Skretas, John Sylvester, and Viktor Zamaraev. Temporal exploration of random spanning tree models. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2876–2887, 2026.

11 Julien Baste, Binh-Minh Bui-Xuan, and Antoine Roux. Temporal matching. *Theoretical Computer Science*, 806:184–196, 2020.

12 Ruben Becker, Arnaud Casteigts, Pierluigi Crescenzi, Bojana Kodric, Michael Raskin, Malte Renken, and Viktor Zamaraev. Giant components in random temporal graphs. *SIAM Journal on Discrete Mathematics*, 40(2):449–485, 2026. `arXiv:https://doi.org/10.1137/24M1673206`, `doi:10.1137/24M1673206`.

13 Anna Brandenberger, Serte Donderwinkel, Céline Kerriou, Gábor Lugosi, and Rivka Mitchell. Temporal connectivity of random geometric graphs. *arXiv preprint arXiv:2502.15274*, 2025.

14 Nicolas Broutin, Nina Kamčev, and Gábor Lugosi. Increasing paths in random temporal graphs. *The Annals of Applied Probability*, 34(6):5498–5521, 2024.

15 Arnaud Casteigts, Michael Raskin, Malte Renken, and Viktor Zamaraev. Sharp thresholds in random simple temporal graphs. *SIAM Journal on Computing*, 53(2):346–388, 2024.

16 Lapo Cioni, Riccardo Dondi, Andrea Marino, Jason Schoeters, and Ana Silva. Matching and edge cover in temporal graphs. In *Proceedings of the 4th Symposium on Algorithmic Foundations of Dynamic Networks (SAND)*, volume 330, pages 8:1–8:16, 2025.

17 Tom Davot, Jessica A. Enright, Jayakrishnan Madathil, and Kitty Meeks. Temporal triadic closure: Finding dense substructures in social networks that evolve over time. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, pages 26939–26946, 2025.

18 Konstantinos Dogeas, Thomas Erlebach, Frank Kammer, Johannes Meintrup, and William K. Moses Jr. Exploiting automorphisms of temporal graphs for fast exploration and rendezvous. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 297, pages 55:1–55:18, 2024.

19 Jessica A. Enright, Kitty Meeks, George B. Mertzios, and Viktor Zamaraev. Deleting edges to restrict the size of an epidemic in temporal networks. *Journal of Computer and System Sciences*, 119:60–77, 2021.

20 Jessica A. Enright, Kitty Meeks, and Hendrik Molter. Counting temporal paths. *Algorithmica*, 87(5):736–782, 2025.

21 Thomas Erlebach, Michael Hoffmann, and Frank Kammer. On temporal graph exploration. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 444–455, 2015. Journal version appeared at JCSS in 2021.

22 Till Fluschnik, Hendrik Molter, Rolf Niedermeier, Malte Renken, and Philipp Zschoche. Temporal graph classes: A view through temporal separators. *Theoretical Computer Science*, 806:197–218, 2020.

23 Alan Frieze and Michał Karoński. *Introduction to random graphs*. Cambridge University Press, 2015.

24 Laetitia Gauvin, Mathieu Génois, Márton Karsai, Mikko Kivelä, Taro Takaguchi, Eugenio Valdano, and Christian L. Vestergaard. Randomized reference models for temporal networks. *SIAM Review*, 64(4):763–830, 2022. arXiv:https://doi.org/10.1137/19M1242252, doi:10.1137/19M1242252.

25 Thekla Hamm, Nina Klobas, George B. Mertzios, and Paul G. Spirakis. The complexity of temporal vertex cover in small-degree graphs. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 10193–10201, 2022.

26 Hovhannes A. Harutyunyan, Kamran Koupayi, and Denis Pankratov. Temporal separators with deadlines. In Satoru Iwata and Naonori Kakimura, editors, *Proceedings of the 34th International Symposium on Algorithms and Computation (ISAAC)*, volume 283, pages 38:1–38:19, 2023.

27 Anne-Sophie Himmel, Hendrik Molter, Rolf Niedermeier, and Manuel Sorge. Adapting the bron-kerbosch algorithm for enumerating maximal cliques in temporal graphs. *Social Network Analysis and Mining*, 7(1):35:1–35:16, 2017.

28 Petter Holme and Jari Saramäki. *Temporal network theory*, volume 2. Springer, 2019.

29 Y. Hulovatyy, H. Chen, and T. Milenković. Exploring the structure and function of temporal networks with dynamic graphlets. *Bioinformatics*, 31(12):i171–i180, 06 2015.

30 Allen Ibiapina, Minh-Hang Nguyen, Mikaël Rabie, and Cléophée Robin. How to color temporal graphs to ensure proper transitions. *CoRR*, abs/2505.10207, 2025. URL: https://doi.org/10.48550/arXiv.2505.10207.

31 D. Kempe, J. M. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of the 32nd Annual ACM symposium on Theory of computing (STOC)*, pages 504–513, 2000.

32 Nina Klobas, George B. Mertzios, Hendrik Molter, Rolf Niedermeier, and Philipp Zschoche. Interference-free walks in time: temporally disjoint paths. *Autonomous Agents and Multi-Agent Systems*, 37(1):1, 2023.

33 Nina Klobas, George B. Mertzios, Hendrik Molter, and Paul G. Spirakis. Temporal graph realization from fastest paths. *Theoretical Computer Science*, 1056:115508, 2025.

34 Nina Klobas and George B. Mertzios Paul G. Spirakis. Sliding into the future: Investigating sliding windows in temporal graphs. In *Proceedings of the 48th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 5:1–5:12, 2023.

35 Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(11):P11005, 2011.

36 Per Kristian Lehre and Carsten Witt. General drift analysis with tail bounds. *CoRR*, abs/1307.2559, 2013.

37 Paul Liu, Austin R. Benson, and Moses Charikar. Sampling methods for counting temporal motifs. In J. Shane Culpepper, Alistair Moffat, Paul N. Bennett, and Kristina Lerman, editors, *Proceedings of the 12th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 294–302, 2019.

38 Nicolas Maack, Hendrik Molter, Rolf Niedermeier, and Malte Renken. On finding separators in temporal split and permutation graphs. *Journal of Computer and System Sciences*, 135:1–14, 2023.

39 Andrea Marino and Ana Silva. Coloring temporal graphs. *Journal of Computer and System Sciences*, 123:171–185, 2022.

40 George B. Mertzios, Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Temporal network optimization subject to connectivity constraints. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 657–668, 2013. Journal version appeared at Algorithmica in 2019.

41 George B. Mertzios, Hendrik Molter, Rolf Niedermeier, Viktor Zamaraev, and Philipp Zschoche. Computing maximum matchings in temporal graphs. In *Proeedings of the 37th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 154, pages 27:1–27:14, 2020. Journal version appeared at JCSS in 2023.

42 George B. Mertzios, Hendrik Molter, Malte Renken, Paul G. Spirakis, and Philipp Zschoche. The complexity of transitively orienting temporal graphs. *Journal of Computer and System Sciences*, 150:103630, 2025.

43 George B. Mertzios, Hendrik Molter, and Viktor Zamaraev. Sliding window temporal graph coloring. *Journal of Computer and System Sciences*, 120:97–115, 2021.

44 George B. Mertzios, Sotiris E. Nikoletseas, Christoforos L. Raptopoulos, and Paul G. Spirakis. The threshold of existence of δ-temporal cliques in random simple temporal graphs. In *Proceedings of the 20th International Symposium on Algorithmics of Wireless Networks (ALGOWIN)*, pages 131–143, 2024.

45 Othon Michail and Paul G. Spirakis. Traveling salesman problems in temporal graphs. *Theoretical Computer Science*, 634:1–23, 2016.

46 Othon Michail and Paul G. Spirakis. Elements of the theory of dynamic networks. *Communications of the ACM*, 61(2):72–72, January 2018.

47 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 (NCA)*, pages 264–271, 2016.

48 Hendrik Molter, Malte Renken, and Philipp Zschoche. Temporal reachability minimization: Delaying vs. deleting. *Journal of Computer and System Sciences*, 144:103549, 2024.

49 Vincenzo Nicosia, John Tang, Cecilia Mascolo, Mirco Musolesi, Giovanni Russo, and Vito Latora. Graph metrics for temporal networks. In *Temporal Networks*. Springer, 2013.

50 Ashwin Paranjape, Austin R. Benson, and Jure Leskovec. Motifs in temporal networks. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 601–610, 2017.

51 Alexandra M. Porter, Baharan Mirzasoleiman, and Jure Leskovec. Analytical models for motifs in temporal networks. In *Proceedings of the Companion of The Web Conference 2022*, pages 903–909. ACM, 2022.

52 Maciej Rymar, Hendrik Molter, André Nichterlein, and Rolf Niedermeier. Towards classifying the polynomial-time solvability of temporal betweenness centrality. *Journal of Graph Algorithms and Applications*, 27(3):173–194, 2023.

53 N. Santoro. Computing in time-varying networks. In *Proceedings of the 13th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, page 4, 2011.

54 Ahmet Erdem Sarıyüce. A powerful lens for temporal network analysis: temporal motifs. *Discover Data*, 3(1):14, 2025.

55 Yuichi Sudo and Toshimitsu Masuzawa. Leader election requires logarithmic time in population protocols. *CoRR*, abs/1906.11121, 2019. URL: `http://arxiv.org/abs/1906.11121`, `arXiv:1906.11121`.

56 Jordan Viard, Matthieu Latapy, and Clémence Magnien. Revealing contact patterns among high-school students using maximal cliques in link streams. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1517–1522, 2015.

57 Tiphaine Viard, Matthieu Latapy, and Clémence Magnien. Computing maximal cliques in link streams. *Theoretical Computer Science*, 609:245–252, 2016.

58 Jiajing Wu, Jieli Liu, Weili Chen, Huawei Huang, Zibin Zheng, and Yan Zhang. Detecting mixing services via mining bitcoin transaction network with hybrid motifs. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 52(4):2237–2249, 2022.