Ephemeral Networks with Random Availability of Links: Diameter and Connectivity*

Eleni C. Akrida
Department of Computer Science
University of Liverpool, UK
eleni.akrida@liverpool.ac.uk

Leszek Gąsieniec
Department of Computer Science
University of Liverpool, UK
l.a.gasieniec@liverpool.ac.uk

George B. Mertzios
School of Engineering and Computing Sciences
Durham University, UK
george.mertzios@durham.ac.uk

Paul G. Spirakis
Department of Computer Science
University of Liverpool, UK
p.spirakis@liverpool.ac.uk

ABSTRACT

In this work we consider temporal networks, the links of which are available only at random times *(randomly available temporal networks)*. Our networks are *ephemeral*: their links appear sporadically, only at certain times, within a given maximum time (*lifetime* of the net). More specifically, our temporal networks notion concerns networks, whose edges (arcs) are assigned one or more random discrete-time labels drawn from a set of natural numbers. The labels of an edge indicate the discrete moments in time at which the edge is available. In such networks, information (e.g., messages) have to follow *temporal paths*, i.e., paths, the edges of which are assigned a strictly increasing sequence of labels. We first examine a very hostile network: a clique, each edge of which is known to be available *only one random time* in the time period $\{1, 2, \ldots, n\}$ (n is the number of vertices). How fast can a vertex send a message to all other vertices in such a network? To answer this, we define the notion of the Temporal Diameter for the random temporal clique and prove that it is $\Theta(\log n)$ with high probability and in expectation. In fact, we show that information dissemination is very fast with high probability even in this hostile network with regard to availability. This result is similar to the results for the random phone-call model. Our model, though, is weaker. Our availability assumptions are different and randomness is provided only by the input. We show here that the temporal diameter of the clique is crucially affected by the clique's lifetime, a, e.g., when a is asymptotically larger than the number of vertices, n, then the temporal diameter must be $\Omega(\frac{a}{n} \log n)$. We, then, consider the least number, r, of *random points in time* at which an edge is available, in order to guarantee at least a temporal path between any pair of vertices of the network (notice that the clique is the only network for which just one instance of availability per edge, even non-random, suffices for this). We show that r is $\Omega(\log n)$ even for some networks of diameter 2. Finally, we compare this cost to an (optimal) deterministic allocation of labels of availability that guarantees a temporal path between any pair of vertices. For this reason, we introduce the notion of the *Price of Randomness* and we show an upper bound for general networks.

*Supported in part by (i) the School of EEE and CS of the Univeristy of Liverpool, (ii) the FET EU IP Project MULTIPLEX under contract No. 317532, and (iii) the EPSRC Grant EP/K022660/1.

SPAA'14, June 23–25, 2014, Prague, Czech Republic.

http://DOI string/url (ACM rightsreview confirmation email).

Categories and Subject Descriptors

G.2.2 [**Discrete Mathematics**]: Graph Theory—*graph labelling, network problems, path and circuit problems*; G.3 [**Probability and Statistics**]: survival analysis

General Terms

Theory

Keywords

Temporal networks, random input, diameter, availability.

1. INTRODUCTION

A temporal network is a network that changes with time. Many networks of today have links that are not always available. In this work, embarking from the foundational work of Kempe et al. [19] and from the sequel [21], we consider time to be discrete, that is, we consider networks in which the links are available *only at certain moments in time*, e.g., days or hours. Such networks can be described via an underlying graph $G = (V, E)$ (*the links of which can become available*) and an assignment L assigning a set of discrete labels to each edge (arc) of G.

We consider here both the single-label-per-edge model of [19] and the multi-labelled one, which allows links to be available at multiple times (i.e., more than one label per edge). These labels are drawn from the natural numbers (in fact from a set of discrete times $S = \{1, 2, \ldots, a\}$) and indicate the discrete moments in time at which the corresponding connection is available. Usually, we take $a = |V(G)|$ (a *normalized* case). Note that such networks are *ephemeral*: No link is available at any time after time (day) a. We call a the *lifetime* of the network.

In many (worst case) situations, availability of links comes at a cost. Available links may correspond, e.g., to connections in physical systems requiring high energy. They may also correspond to very rare moments in time, in which the link of a hostile network is "unguarded" and, thus, one can pass a message at that time without putting the message in danger.

A temporal path (or journey) in such a network is a path on the edges (arcs) of which one can find *strictly increasing* labels. The time label used on the last edge of a temporal path would then indicate the time at which a message would arrive at the last vertex of the path.

Imagine a very hostile *clique-network* G, the links of which are all usually guarded. Whenever a link is guarded it is impossible to pass a message through it. We may pass a message to a neighbour (in G) only when the link to this neighbour is unguarded (i.e., available). Now, let us assume that each link will become available only at a *random time* in S. Let us look at the case where $S = \{1, 2, \ldots, n\}$ (where $n = |V(G)|$). After time n, no link of the clique is ever available! Such a random time indicates a break in the security of the link. How fast can we pass a message (starting from a vertex s) to all the other vertices in the clique? Certainly, one possibility is to wait (for each destination t) for the link (s, t) to become available. But this may mean a passing time equal to $\frac{n}{2}$ in expectation. Can we spread a message faster? In this paper, we show that for the temporal clique with a single random moment of availability per link, one can still pass the message to all vertices in time $\Theta(\log n)$ with high probability. That is, a seemingly very hostile clique (each link of which is unguarded only for one random moment) is, in fact, not so secure with respect to fast dissemination of enemy information.

Note that the clique is the only graph G which achieves *temporal reachability* of all vertices, even when each edge (link) of it is only available at one time, chosen from the set $\{1, 2, \ldots, n\}$.

One now may wish to pay for greater (more) availability of links. In fact, given a graph G, one may wish to find the minimum total number of labels, over all edges, (OPT) that guarantee a temporal path between any $s, t \in V(G)$. But, there are cases in which this OPT quantity is not even approximable in P unless $P = NP$ [21]. Instead, you may bargain locally (per link) to "buy" a *number of random times for which the link is available*. We assume here that we are given a complicated network. Each node has no information about the network topology, but knows the number of vertices, n, and the diameter, d, of the net. We assume that *global coordination over availability of all edges is impossible*. (Otherwise, an assignment of *the same* d consequtive labels per edge would guarantee all-pairs reachability. But this requires global coordination.) However, we allow adjacent vertices to agree on a number, $r(n)$, of *random available times for the edge joining them*. This is a local operation that uses local, random availability to replace the lack of global knowledge and global coordination. What is the least r to guarantee a temporal path between any pair of vertices in G, with high probability? In this paper, we show that r is lower bounded by $\Omega(\log n)$ *even for some graphs of diameter* 2. We then estimate sufficient values of r to guarantee temporal paths between all node-pairs (with high probability) for any graph G. In this paper, we write that an event holds "with high probability" when there exists a constant $c \geq 1$ such that the probability of the event is at least $1 - \frac{1}{n^c}$, where n is the number of vertices.

In this work we, thus, initiate the study of "random" ephemeral temporal networks. We define notions like the *temporal diameter* (to capture fast dissemination of information) and the *Price of Randomness*, PoR (to capture the cost to pay per link in order to guarantee temporal reachability of all node-pairs by local random available times with high probability). Intuitively, PoR is the ratio $\frac{mr}{OPT}$, where r random times of availability for each and every link are able to guarantee all-pairs temporal reachability with high probability. We believe that our work will motivate further research on both temporal networks and random temporal networks.

1.1 Relation to the Random Phone-Call Model

The first logarithmic time results for probabilistic information dissemination were obtained in the classical Random Phone-Call model defined in [9]. In [9], the authors present a push algorithm that uses $\Theta(\log n)$ time and $\Theta(n \log n)$ message transmissions. For complete graphs of size n, Frieze and Grimmett [15] presented an algorithm that broadcasts in time $\log_2 n + \ln n + o(\log n)$ with a probability of $1 - o(1)$. Later, Pittel [27] showed that (with probability $1 - o(1)$) it is possible to broadcast a message in time $\log_2 n + \ln n + f(n)$, where $f(n)$ can be any slow growing function.

Karp et al. [17] presented a push and pull algorithm which reduces the total number of transmissions to $O(n \log \log n)$, with probability $1 - n^{-1}$, and showed that this result is asymptotically optimal. For sparser graphs it is not possible to stay within $O(n \log \log n)$ message transmissions together with a broadcast time of $O(\log n)$ in this phone-call model, not even for random graphs [11]. However, if each node is allowed to remember a small number of neighbors to which it has communicated in some previous steps, then the number of message transmissions can be reduced to $O(n \log \log n)$, with probability $1 - n^{-1}$ [3, 12].

The network model adopted in this paper resembles the Random Phone-Call model to some extent, however, it is essentially different. The dependence of the temporal diameter (of the hostile clique) on its lifetime, for example, cannot be captured by the random phone-call model. The model described here is, in fact, considerably weaker. In the phone-call model, each node, at each step, can communicate with a random neighbour (in fact, a node may do this at several times). In our model, each link is given a (single) random moment of existence, by the input. A node can send via this link only at that moment. That is, randomness is not a part of our algorithmic techniques and can not be used at arbitrary time steps.

1.2 Other related work

In this section we provide a short survey of papers with studies on networks labelled by time units or segments.

Labelled Graphs. Labelled graphs have been widely used both in Computer Science and in Mathematics, e.g., [25].

Single-labelled and multi-labelled Temporal Networks. The model of temporal networks that we consider in this work is a direct extension of the single-labelled model studied in [19] as well as the multi-labelled model studied in [21]. The prior results of [19,21] do not consider randomness at all, and therefore are different in nature to this work. The initial paper [19] considers the case of one (non-random) label per edge and examines shortest journey algorithms. The second paper [21] extends this (non-random) model to many labels per edge and mainly examines the number of labels needed to guarantee several graph properties with certainty.
Continuous Availabilities (Intervals). Some authors have assumed the availability of an edge for a whole time-interval $[t_1, t_2]$ or multiple such time-intervals and not just for discrete moments as we assume here. Although this is a clearly natural assumption, we design and develop techniques for the discrete case which are quite different from those needed in the continuous case [6,14].
Dynamic Distributed Networks. In recent years, there is a growing interest in distributed computing systems that are inherently dynamic [1,2,4,7,8,10,20,22–24,26,28].
Distance labelling. A distance labelling of a graph G is an assignment of unique labels to vertices of G so that the distance between any two vertices can be inferred from their labels alone [16,18].

2. PRELIMINARIES

In this section we first define *temporal networks* (cf. Definition 1) by assigning a set $L_e \subseteq \mathbb{N}$ of *time-labels* to every edge e of a (di)graph $\widetilde{G} = (V, E)$.

DEFINITION 1. *Let $\widetilde{G} = (V, E)$ be a (di)graph. A temporal network on $\widetilde{G}$ is a triplet $G = (\widetilde{G}, L)$ (also denoted as $G = (V, E, L)$), where $L = \{L_e \subseteq \mathbb{N} : e \in E\}$ is an* assignment *of labels on the edges of $\widetilde{G}$.*

When for every edge e, $L_e \subseteq \{1, 2, \ldots, a\}$, for some $a \in \mathbb{N}$, the network is called *ephemeral* and a is called the *lifetime* of the network.

The values assigned to each edge of the graph are called time labels of the edge and indicate the times at which we can cross it (from one end to the other in arbitrary direction, if the edge is undirected, or from its start to its end, if the edge is directed).

In the context of this paper, we mainly study *random temporal networks*, in which the labels assigned to the edges are chosen at random from a set of available time labels. More specifically, the *model* that we consider is that of random temporal networks, an instance of the labels of each edge of which is given in advance to the algorithm, so that the traveller can see all adjacent edges to all vertices at every moment in time.

2.1 Further Definitions

We can now talk about temporal edges (or time edges) that are considered to be triplets (u, v, l), where u, v are the ends of an edge in the temporal network and $l \in L_{\{u,v\}}$ is a time label of this edge. That is, if an edge $e = \{u, v\}$ has more than one time labels, e.g., has a set of three time labels, $L_e = \{l_1, l_2, l_3\}$, then this edge has three corresponding time edges, (u, v, l_1), (u, v, l_2) and (u, v, l_3).

DEFINITION 2. *A temporal path or journey j from a vertex u to a vertex v ((u, v)-journey) is a sequence of time edges (u, u_1, l_1), $(u_1, u_2, l_2), \ldots,$ (u_{k-1}, v, l_k), such that $l_i < l_{i+1}$, for each $1 \le i \le k - 1$. We call the last time label of journey j, l_k, a*rrival time *of the journey.*

DEFINITION 3. *A (u, v)-journey j in a temporal network is called* f*oremost journey if its arrival time is the minimum arrival time of all (u, v)-journeys' arrival times, under the labels assigned to the graph's edges. We call this arrival time* temporal distance *of target vertex v from source vertex u and we denote $\delta(u, v)$.*

Now, consider any ephemeral temporal network $G = (V, E, L)$. Let every edge receive exactly one time label, chosen randomly, independently from one another from a set $L_0 = \{1, 2, \ldots, a\}$, where $a \in \mathbb{N}$, with the probability of an edge label to be i, $\forall i \in L_0$, equal to $\frac{1}{a}$ (**UNI-CASE**).

DEFINITION 4. *A temporal network that satisfies UNI-CASE is called Uniform Random Temporal Network (U-RTN).*

In the special case, where the largest label, a, that can be assigned to the edges of a graph is equal to the number of its vertices, the formed network is called *Normalized Uniform Random Temporal Network (Normalized U-RTN).*

Note.

There could be prospective study of cases in which each edge of a graph may receive several time labels, selected randomly and independently of one another from the set $L_0 = \{1, 2, \ldots, a\}$, where $a \in \mathbb{N}$, with the selection following a distribution F (**F-CASE**). In such cases, the networks under consideration would be called *F-Random Temporal Networks (F-RTN)* respectively.

In the following section, we focus on the study of uniform random temporal networks, where the underlying graph is complete (clique). Under this scope, we define a statistical property of the uniform random temporal clique, namely its temporal diameter.

DEFINITION 5. *Consider an instance (G, L) of a uniform random temporal clique, G. The Temporal Diameter of G, denoted by TD, is the expected value of the maximum temporal distance over all pairs of vertices in G:*

$$TD(G) = E\big(max_{s,t \in V(G)} \delta(s, t)\big)$$

3. THE TEMPORAL DIAMETER OF THE NORMALIZED UNIFORM RANDOM TEMPORAL CLIQUE

Let $G = K_n$ be a directed clique[1] of n vertices and let us consider its normalized U-version. That is, every edge $e \in E(K_n)$ is given a single availability label, l_e, and those labels are chosen randomly and independently from one another from the set $L_0 = \{1, 2, \ldots, n\}$, with the probability that the label of a particular edge equals i being equal to $\frac{1}{n}$, $\forall i \in L_0$.

We give an algorithmic construction (Algorithm 1) which can, with high probability, find a journey with small expected arrival time from any given source vertex s to any

[1] for every pair of vertices $u, v \in K_n$, there exist both the directed edges $\{u, v\}$ and $\{v, u\}$

given target vertex t in the directed normalized uniform random temporal clique, K_n.

Remark 1. *It is easy to see that the same result holds for the undirected uniform random temporal clique. In this case, an edge e of the clique with a random label l corresponds to two directed edges e' and e'' of the directed clique and the analysis is not significantly affected.*

Algorithm 1 The directed normalized U-RT clique Expansion Process algorithm

Input: An instance of a directed normalized uniform random temporal clique of n vertices, K_n

```
d = Θ(log n); {the exact value of d and of the constants c_1, c_2 below will be determined by the analysis}
Γ_1(s) = {v ∈ V : l_{s,v} ∈ (0, c_1 log n]};
for i = 2, ..., d+1 do
    Γ_i(s) = {v ∈ V : l_{w,v} ∈ (c_1 log n + (i−2)c_2, c_1 log n + (i−1)c_2] for some w ∈ Γ_{i−1}(s)};
Γ'_1(t) = {v ∈ V : l_{v,t} ∈ (2c_1 log n + 2dc_2, 3c_1 log n + 2dc_2]};
for i = 2, ..., d+1 do
    Γ'_i(t) = {v ∈ V : l_{v,w} ∈ (2c_1 log n + (2d − i + 1)c_2, 2c_1 log n + (2d − i + 2)c_2] for some w ∈ Γ'_{i−1}(s)};
if ∃u ∈ Γ_{d+1}(s), v ∈ Γ'_{d+1}(t) such that l_{u,v} ∈ (c_1 log n + dc_2, 2c_1 log n + dc_2] then
    Follow the directed path from s to u, the directed edge (u, v) and the directed path from v to t;
    return success;
else
    return failure;
```

Analysis of the Expansion Process Algorithm.

Next, we analyze the Expansion Process algorithm and we prove that it succeeds with high probability, thus giving a short, $O(\log n)$, journey from s to t.

Note. In the following analysis, we reveal each arc's random label only once, when examined (*delayed revelation of random values*). Thus, we are consistent with the fact that the input is a specific instance (with all random labels drawn).

Denote by p_1 the probability that an outgoing edge of s has a label in the desired interval, i.e., $(0, c_1 \log n]$. p_1 is also the probability that an outgoing edge of some vertex in $\Gamma_{d+1}(s)$ to $\Gamma'_{d+1}(t)$ has label in the desired interval, i.e., $(c_1 \log n + dc_2, 2c_1 \log n + dc_2]$. Finally, p_1 is also the probability that an incoming edge of t has a label in the desired interval, i.e., $(2c_1 \log n + 2dc_2, 3c_1 \log n + 2dc_2]$. It is:

$$p_1 = \frac{c_1 \log n}{n}$$

Denote by p_2 the probability that a vertex $v \in \Gamma_i(s)$ (or a $v \in \Gamma'_i(t)$), $i = 1, 2, \ldots, d$, has an outgoing edge (or incoming edge, respectively) with label that falls in the desired interval, i.e., $(c_1 \log n + (i-1)c_2, c_1 \log n + ic_2]$ (or $(2c_1 \log n + (2d-1)c_2, 2c_1 \log n + (2d-i+1)c_2]$, respectively). It is:

$$p_2 = \frac{c_2}{n}$$

Also, denote by Δ_i, Δ^* and $\Delta'_i, i = 1, 2, \ldots, d+1$ the desired intervals in each case, namely:

$$\begin{aligned}
\Delta_1 &= (0, c_1 \log n] \\
\Delta_2 &= (c_1 \log n, c_1 \log n + c_2] \\
\Delta_3 &= (c_1 \log n + c_2, c_1 \log n + 2c_2] \\
&\ldots \\
\Delta_{d+1} &= (c_1 \log n + (d-1)c_2, c_1 \log n + dc_2] \\
\Delta^* &= (c_1 \log n + dc_2, 2c_1 \log n + dc_2] \\
\Delta'_{d+1} &= (2c_1 \log n + dc_2, 2c_1 \log n + (d+1)c_2] \\
&\ldots \\
\Delta'_2 &= (2c_1 \log n + (2d-1)c_2, 2c_1 \log n + 2dc_2] \\
\Delta'_1 &= (2c_1 \log n + 2dc_2, 3c_1 \log n + 2dc_2]
\end{aligned}$$

Note. If there exists at least one edge with label in the corresponding Δ_i, Δ^* or $\Delta'_i, i = 1, 2, \ldots, d+1$ at each step of the expansion, then the time needed to reach t starting from s is at most $3c_1 \log n + 2dc_2$.

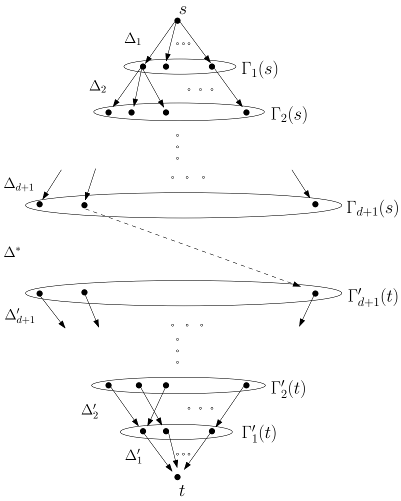

Figure 1: The Expansion Process.

Figure 1 illustrates how the expansion process from s to t works. That is, starting from s, we find the set $\Gamma_1(s)$ of vertices to which there is an edge from s with label within Δ_1, then the set $\Gamma_2(s)$ of vertices to which there is an edge from a vertex in $\Gamma_1(s)$ with label within Δ_2, etc. We show that, with high probability, there is a journey from s to t through vertices in the consecutive Γ_is.

3.1 The first step of the expansion process

The first step of the expansion process aims in establishing with high probability a number of $\Theta(\log n)$ neighbours of s, so that the edge from s to each one of them is in Δ_1. Note that the probability of a label on an edge $(s, u), u \in V$ being in Δ_1 is exactly $p_1 = \frac{c_1 \log n}{n}$, because of the uniform selection of labels.

Let $\mathcal{E}_1$ be the event that $\frac{1}{2}E(|\Gamma_1(s)|) \leq |\Gamma_1(s)| \leq \frac{3}{2}E(|\Gamma_1(s)|)$. Note that:

$$E(|\Gamma_1(s)|) = (n-1)p_1 = (n-1)\frac{c_1 \log n}{n}$$

By the Chernoff bound on the Binomial $B(N, p_1)$, where $N = n - 1$, $\forall \beta \in (0, 1)$, it holds:

$$Pr(\#successes \in (1 \pm \beta)Np_1) \geq 1 - e^{-\frac{\beta^2}{2}Np_1}$$

Now, use $\beta = \frac{1}{2}$. We get:

$$\begin{aligned} Pr(\#successes \in (\frac{1}{2}, \frac{3}{2})Np_1) &\geq 1 - e^{-\frac{1}{8}Np_1} \\ &\geq 1 - e^{-\frac{1}{8}(c_1 \log n - \frac{c_1 \log n}{n})} \\ &\geq 1 - e^{-\frac{1}{8}(c_1 - 1)\log n} \\ &\geq 1 - \frac{1}{n^{\frac{c_1 - 1}{8}}} \end{aligned}$$

We now choose $c_1 \geq 33$, and thus $\frac{c_1-1}{8} \geq 4$. So we have established the following:

Lemma 1. *It holds that:*

$$Pr(\mathcal{E}_1) = Pr\left(|\Gamma_1(s)| \in (\frac{1}{2}, \frac{3}{2})(c_1 \log n(1 - \frac{1}{n}))\right) \geq 1 - \frac{1}{n^4}$$

3.2 The expansion process until reaching $\Theta(\sqrt{n})$ vertices

We now show that given:

- $|\Gamma_1(s)| = \Theta(\log n)$, and
- the probability of an edge having a label in a particular interval $\Delta_i, i = 2, \ldots, d+1$ is exactly $p_2 = \frac{|\Delta_i|}{n} = \frac{c_2}{n}$

the vertices reachable from s via temporal paths grow (almost) geometrically.

In particular, let us now condition on the event that $\frac{1}{8}c_1 \log n \leq |\Gamma_i(s)| \leq \lambda\sqrt{n}$, for some fixed $\lambda > 0$. To find the set $\Gamma_{i+1}(s)$, we consider the vertices which are *not* in all the $\Gamma_j(s)$, $j = 1, 2, \ldots, i$ (and the fact that we look for directed edges), i.e.,

$$n_i = n - |\bigcup_{j=1}^{i} \Gamma_j(s)|$$

The probability that a vertex u (out of the n_i vertices) belongs to $\Gamma_{i+1}(s)$ is exactly the probability that the label of some $(v, u) \in E, v \in \Gamma_i(s)$, is in the interval Δ_{i+1}, i.e., equal to:

$$\begin{aligned} q &= 1 - Pr(u \notin \Gamma_{i+1}(s)) \\ &= 1 - (1 - p_2)^{|\Gamma_i(s)|} \\ &= 1 - (1 - \frac{c_2}{n})^{|\Gamma_i(s)|} \end{aligned}$$

We need the following fact:

Fact 1. *It holds that* $(1 - \frac{c_2}{n})^{|\Gamma_i(s)|} \leq 1 - \frac{c_2|\Gamma_i(s)|}{2n}$

Proof. Let $p = \frac{c_2}{n}$ and $k = |\Gamma_i(s)|$. We know that:

$$(1-p)^k \leq 1 - kp + \binom{k}{2}p^2 \tag{1}$$

We will show that:

$$-kp + \binom{k}{2}p^2 \leq -\frac{kp}{2}$$

and, thus, by relation 1 it holds that:

$$(1-p)^k \leq 1 - \frac{kp}{2}$$

Indeed, we have:

$$\begin{aligned} -kp + \binom{k}{2}p^2 &\leq -\frac{kp}{2} \Leftrightarrow \\ \frac{k(k-1)}{2}p &\leq \frac{k}{2} \Leftrightarrow \\ (k-1)p &\leq 1 \Leftrightarrow \\ (|\Gamma_i(s)| - 1)c_2 &\leq n \end{aligned}$$

The latter holds for n sufficiently large. $\square$

Thus, we have:

$$\begin{aligned} q &\geq 1 - (1 - \frac{c_2|\Gamma_i(s)|}{2n}) \\ &= \frac{c_2|\Gamma_i(s)|}{2n} \\ &\geq \frac{c_1 c_2 \log n}{16n} = q' \end{aligned}$$

The random variable $|\Gamma_{i+1}(s)|$ follows the Binomial distribution $B(n_i, q)$ and dominates $B(n_i, q')$. Therefore, by the Chernoff bound (with $\beta = \frac{1}{2}$), we have:

$$Pr\big(|\Gamma_{i+1}(s)| \in (\frac{1}{2}n_i q, \frac{3}{2}n_i q)\big) \geq 1 - e^{-\frac{1}{8}n_i q'} \tag{2}$$

But, $n_i \geq n - (\lambda\sqrt{n})d \geq \frac{n}{2}$. So, relation 2 becomes:

$$\begin{aligned} Pr\big(|\Gamma_{i+1}(s)| \in (\frac{1}{2}n_i q, \frac{3}{2}n_i q)\big) &\geq 1 - e^{-\frac{1}{16}n\frac{c_1 c_2 \log n}{16n}} \\ &\geq 1 - e^{-\frac{1}{256}c_1 c_2 \log n} \\ &\geq 1 - \frac{1}{n^{\frac{c_1 c_2}{256}}} \end{aligned}$$

We will select c_2 so that $\frac{c_1 c_2}{256} \geq 4$. So, with probability at least $1 - \frac{1}{n^4}$, it is:

$$\begin{aligned} \frac{3}{2}n_i q \quad &\geq |\Gamma_{i+1}(s)| \geq \quad \frac{1}{2}n_i q \Rightarrow \\ \frac{3}{2}n_i \frac{c_2|\Gamma_i(s)|}{2n} \quad &\geq |\Gamma_{i+1}(s)| \geq \quad \frac{1}{2}n_i \frac{c_2|\Gamma_i(s)|}{2n} \Rightarrow \\ \frac{3}{4}c_2|\Gamma_i(s)| \quad &\geq |\Gamma_{i+1}(s)| \geq \quad \frac{1}{8}c_2|\Gamma_i(s)| \end{aligned}$$

We have proved that the event

$\mathcal{E}_i$ = "$|\Gamma_{i+1}(s)|$ is at most $\frac{3}{4}c_2|\Gamma_i(s)|$ and at least $\frac{1}{8}c_2|\Gamma_i(s)|$"

holds with probability at least $1 - \frac{1}{n^4}$, provided that $\frac{1}{8}c_1 \log n \leq \Gamma_i(s) \leq \lambda\sqrt{n}$.

Thus, by conditioning on the event $\mathcal{E} = \bigcap_{i=1}^{d} \mathcal{E}_i$, we have that:

$$|\Gamma_{d+1}(s)| \geq \frac{1}{8}(\frac{c_2}{8})^d c_1 \log n$$

and also

$$|\Gamma_{d+1}(s)| \leq \frac{1}{8}(\frac{3c_2}{4})^d c_1 \log n$$

Choose d so that:

$$\frac{1}{8}(\frac{3c_2}{4})^d c_1 \log n \le \lambda' \sqrt{n}, \text{ for some constant } \lambda' > 0$$

$$\Rightarrow d \le \frac{\log \frac{8\lambda'\sqrt{n}}{c_1 \log n}}{\log \frac{3c_2}{4}}$$

and also:

$$\frac{1}{8}(\frac{c_2}{8})^d c_1 \log n > \sqrt{n}$$

$$\Rightarrow d > \frac{\log \frac{8\sqrt{n}}{c_1 \log n}}{\log \frac{c_2}{8}}$$

The probability that one or more of the events $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_d$ fail is (by the union bound) at most:

$$d\frac{1}{n^4} \le c' \log n \frac{1}{n^4} \le \frac{1}{n^3}, \text{ for some } c' > 0$$

Thus, we have shown the following:

Theorem 1. *With probability at least* $1 - \frac{1}{n^3}$*, the expansion process out of* s *arrives at* $\Theta(\sqrt{n})$ *vertices with temporal paths of length* $d + 1 = \Theta(\log n)$*, consistently labelled in the intervals* Δ_i*,* $i = 1, 2, \ldots, d + 1$*, in time at most* $c_1 \log n + dc_2 = \Theta(\log n)$*.*

3.3 The reverse expansion process (out of t)

Consider the edges reaching t *reversed* and consider the process that labels them in Δ'_1. Let $\Gamma'_1(t)$ be the vertices derived in this way, i.e., reaching t with an edge labelled in Δ'_1. Continue the reverse expansion process until we reach $\Theta(\sqrt{n})$ vertices. By symmetry and independence, we get exactly the same result as in Theorem 1:

Theorem 2. *The expansion process out of* t *arrives at* $\Theta(\sqrt{n})$ *vertices with temporal paths (reverse direction) of length* $d+1 = \Theta(\log n)$*, consistently labelled in the intervals* Δ'_i*,* $i = 1, 2, \ldots, d + 1$*. Thus, it arrives to each of these vertices in time at most* $c_1 \log n + dc_2 = \Theta(\log n)$ *with probability at least* $1 - \frac{1}{n^3}$*.*

3.4 The matching argument

The probability that both $|\Gamma_{d+1}(s)|$ and $|\Gamma'_{d+1}(t)|$ are of size at least $\lambda'\sqrt{n}$, $\lambda' > 0$ is at least $1 - 2\frac{1}{n^3}$. Note that we just need one edge $(v_1, v_2), v_1 \in \Gamma_{d+1}(s), v_2 \in \Gamma'_{d+1}(t)$ with label in the interval Δ^* in order to demonstrate the existence of a temporal path of largest label at most $\Theta(\log n)$ from s to t. Note also that for a *given edge* $(v_1, v_2), v_1 \in \Gamma_{d+1}(s), v_2 \in \Gamma'_{d+1}(t)$, its label is in Δ^* with probability exactly:

$$p_1 = \frac{|\Delta^*|}{n} = \frac{c_1 \log n}{n}$$

Thus, the probability of the event A = "existence of such an edge" is:

$$p = 1 - \left(1 - \frac{c_1 \log n}{n}\right)^{|\Gamma_{d+1}(s)|\cdot|\Gamma'_{d+1}(t)|}$$

and due to Theorems 1 and 2, it is:

$$\begin{aligned} p &\ge 1 - \left(1 - \frac{c_1 \log n}{n}\right)^{(\lambda')^2 n} \\ &\ge 1 - e^{-(\lambda')^2 c_1 \log n} \\ &= 1 - \frac{1}{n^{(\lambda')^2 c_1}} \end{aligned}$$

We can choose c_1 through the analysis so that we have:

$$p \ge 1 - \frac{1}{n^3}$$

The probability of any of the events of Theorems 1 and 2 or event A failing is at most $3\frac{1}{n^3}$. Thus,

Theorem 3. *In the directed normalized uniform random temporal clique, given any vertices* s, t*, we can go from* s *to* t *via a temporal path of length at most* $\gamma \log n$*, for some constant* $\gamma > 1$*, with probability at least* $1 - \frac{3}{n^3}$*.*

But, then we get our main theorem, as follows:

Theorem 4. *The Temporal Diameter of the directed normalized uniform random temporal clique is (with high probability) at most* $\gamma \log n$*, for some constant* $\gamma > 1$*.*

Proof. The probability that there exists a pair of vertices $s, t \in V$ so that Theorem 3 fails is less than $n^2 \frac{3}{n^3} = \frac{3}{n}$ (by the union bound). So, with probability at least $1 - \frac{3}{n}$, it holds that:

$$max_{s,t\in V}\{\text{temporal distance of } t \text{ from } s\} \le \gamma \log n$$

Thus,

$$TD \le \gamma \log n, \text{ with probability at least } 1 - \frac{3}{n}$$

and

$$TD > \gamma \log n \text{ (but still } \le n\text{), with probability at most } \frac{3}{n}$$

Therefore,

$$\begin{aligned} TD &\le (1 - \frac{3}{n})\gamma \log n + \frac{3}{n} n \\ &\le \gamma \log n - \frac{3 \log n}{n} + 3 \end{aligned}$$

□

Remark.

One can easily see that the latter is a *threshold* and that the Temporal Diameter of the directed normalized uniform random temporal clique cannot be any less than $\Omega(\log n)$. Assume the event E_1, where the temporal diameter of the directed normalized uniform random temporal clique G is $TD(G) = o(\log n)$, i.e.,

$$\exists \alpha(n) \xrightarrow[n\to+\infty]{} +\infty : TD(G) \le \frac{\log n}{\alpha(n)}$$

Conditional on E_1, the label in every edge of G realizing the diameter is within the interval $(0, \frac{\log n}{\alpha(n)})$. Then, the probability of an edge e "existing" at some moment within the interval $(0, n)$ is:

$$\frac{l_e}{n} \le \frac{\log n}{n\alpha(n)} = p$$

The temporal connectivity of G is dominated by the probability that $G_{n,p}$ is connected. But when $p = o(\frac{\log n}{n})$, then $G_{n,p}$ will almost surely be disconnected [5].

3.5 Spreading a message in the directed uniform random temporal clique

Let us consider again the very hostile clique network G of n vertices in which each edge is available only one random time in the time period $\{1, 2, \ldots, n\}$, and let us consider the case where a vertex s wishes to propagate a message to all other vertices. How fast can this message from s be disseminated to the whole network? Consider the following protocol:

$\forall u \in V(G)$, if u has the message from s, then:
when an arc out of u becomes available, send the message through that arc;

The expansion process described in Algorithm 1 is a construction that demonstrates a temporal path with $O(\log n)$ arrival time from any vertex of the directed uniform random temporal clique to any other vertex with high probability. The above protocol merely exploits Algorithm 1. Thus, it will achieve the dissemination of the message from a specific vertex s to all other vertices of the directed uniform random temporal clique network in *logarithmic time*, $O(\log n)$.

3.6 Temporal Diameter and lifetime - A lower bound

It is easy to see that the following theorem holds:

THEOREM 5. *Let G be the uniform random temporal clique network of n vertices and of lifetime a, i.e., each edge is available exactly one random time within the time period $\{1, 2, \ldots, a\}$, for some $a \in \mathbb{N}$. If a is asymptotically larger than n, then the temporal diameter must be $\Omega(\frac{a}{n} \log n)$.*

PROOF. Assume that the temporal diameter was $k < \frac{a}{n} \log n$. Now, consider (only) the arcs with labels up to k. Since the probability distribution of the labels on the edges of G is uniform, this edge-induced subgraph is the Erdös-Rényi random graph $G_{n,p}$ [5,13], where $p = \frac{k}{a} < \frac{\log n}{n}$. However, it is well known that for such p, $G_{n,p}$ is disconnected with high probability. Therefore, with high probability, the maximum label in a temporal path between at least one pair of vertices is at least $k+1$, i.e. $\Omega(\frac{a}{n} \log n)$. $\square$

The dependence of the Temporal Diameter on the lifetime is a phenomenon that is not captured by static models (such as the random phone-call model).

4. GUARANTEEING TEMPORAL REACHABILITY WITH HIGH PROBABILITY: GRAPHS OF SMALL DIAMETER

4.1 Definitions

Note that the clique is the only graph for which temporal reachability is guaranteed even with 1 random label per edge (drawn from any distribution). This is the case, because one can always follow the edge (s, t) from any s to any t at the time given by the label. For other networks, one may hope that temporal reachability can be guaranteed (whp) with a number of *random* labels per edge.

In the following, we consider selection of labels from the set $\{1, 2, \ldots, n\}$ for a graph $G = (V, E)$ with $|V| = n$ (*normalized* case). We focus on *independent* and *uniformly random* selection of labels (available times) for each edge.

Let $G = (V, E)$ be a (di)graph and L be an assignment of time labels on the edges of G. Consider the property T_{reach} = "$\forall u, v \in V$, $\exists(u, v)$-path in $G \Leftrightarrow \exists(u, v)$-journey in (G, L)".

DEFINITION 6. *An assignment L of temporal labels to the edges of a graph G preserves the* reachability *of G if (G, L) has the property T_{reach}.*

DEFINITION 7. *Let $G = (V, E)$ be a connected (di)graph with $|V| = n$. A random experiment E which assigns $r(n)$ independent random labels to every edge of G strongly guarantees temporal reachability with high probability, if the probability of the property T_{reach} (in the experiment E) is at least $1 - \frac{1}{n^a}$, for some $a \geq 1$.*

DEFINITION 8. *Let $G = (V, E)$ be a connected (di)graph with $|V| = n$ and $E = m$. Let $r(n)$ be the smallest number of random labels per edge which, when assigned to the edges of G, strongly guarantees temporal reachability with high probability. Let, also, $OPT = \sum_{e \in E} |L_e|$ be the total number of labels assigned to the edges of G in the optimal*[2] *(deterministic) assignment which preserves the reachability of G. The Price of Randomness for G is:*

$$PoR(G) = m \frac{r(n)}{OPT}$$

Note that, for some cases, it has been shown that OPT is hard to approximate (there exists no $PTAS$) unless $P = NP$ [21].

4.2 The Price of Randomness can be high

We show here that $PoR(G)$ is *not bounded by any constant* even for graphs G of *diameter* 2.

THEOREM 6. *There is a graph $G = (V, E)$ of n vertices and diameter 2 for which:*

$$PoR(G) = \Theta(\log n)$$

PROOF. Let us consider the *star* graph G_n of n vertices, that is the complete bipartite graph $K_{1,n-1}$: a tree with one internal node and $n - 1$ leaves. Note that $OPT = 2m$, since there exists an assignment of 2 labels per edge (e.g., labels 1, 2 for every edge) which preserves the reachability of G_n, and obviously any assignment of 1 label per edge does not. We will show that $PoR(G_n) = \Theta(\log n)$.

(a) First, we establish that $r(n) = \Theta(\log n)$ random labels per edge are enough to strongly guarantee temporal reachability whp. Let us use $r(n) = \rho \log n$ ($\rho > 8$) random labels per edge. Denote by c the center vertex of G_n. Now consider two fixed leafs, u_1, u_2, of G_n.

Each of the edges $e_1 = \{u_1, c\}$ and $e_2 = \{c, u_2\}$ is assigned $r(n)$ random labels. Let us denote by s_1, s_2 the sets of labels assigned to e_1 and e_2 respectively. We call *2-split* (u_1, u_2)-journey any (u_1, u_2)-journey, where the first temporal edge has a label within the interval $(0, \frac{n}{2})$ and the second temporal edge has a label within the interval $(\frac{n}{2}, n)$(see Figure 2).

[2]By optimal assignment, we mean the assignment with the least total number of labels.

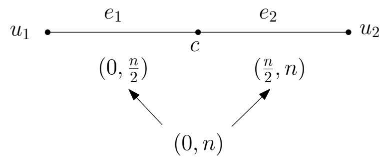

Figure 2: 2-split journey in a star graph.

The probability that an element of s_1 falls within the interval $(0, \frac{n}{2})$ is $\frac{1}{2}$. So, the probability that no element of s_1 falls within this interval is:

$$\begin{aligned} Pr(\text{all labels of } e_1 \geq \frac{n}{2}) &= (1 - \frac{1}{2})^{\rho \log n} \\ &\leq e^{-\frac{\rho \log n}{2}} \\ &= \frac{1}{n^{\frac{\rho}{2}}} \end{aligned}$$

Similarly, the probability that an element of s_2 falls within the interval $(\frac{n}{2}, n)$ is $\frac{1}{2}$. So, the probability that no element of s_2 falls within this interval is:

$$\begin{aligned} Pr(\text{all labels of } e_2 \leq \frac{n}{2}) &= (1 - \frac{1}{2})^{\rho \log n} \\ &\leq \frac{1}{n^{\frac{\rho}{2}}} \end{aligned}$$

Hence, the probability that we can find a label $l_1 \in s_1$ and a label $l_2 \in s_2$ such that $l_1 \in (0, \frac{n}{2})$ and $l_2 \in (\frac{n}{2}, n)$, i.e., the probability that there exists a *2-split* (u_1, u_2)-journey, is:

$$\begin{aligned} & Pr(\exists \textit{2-split}(u_1, u_2) - \textit{journey}) \\ = \; & Pr(\exists l_1 \in s_1, l_2 \in s_2 : l_1 \in (0, \frac{n}{2}) \;\&\; l_2 \in (\frac{n}{2}, n)) \\ = \; & Pr(\exists l_1 \in s_1 : l_1 \in (0, \frac{n}{2})) \\ & \cdot Pr(\exists l_2 \in s_2 : l_2 \in (\frac{n}{2}, n)) \\ = \; & (1 - Pr(\text{all labels of } e_1 \geq \frac{n}{2})) \\ & \cdot (1 - Pr(\text{all labels of } e_2 \leq \frac{n}{2})) \\ \geq \; & (1 - \frac{1}{n^{\frac{\rho}{2}}})^2 \geq 1 - \frac{2}{n^{\frac{\rho}{2}}} \end{aligned}$$

Therefore, it is almost sure that we can find a *2-split* (u_1, u_2)-journey in (G, L). Now, the probability that there exists a pair of vertices $s, t \in V(G)$ such that there is no *2-split* (s, t)-journey in (G, L) is:

$$\begin{aligned} & Pr(\exists s, t \in V(G) : \nexists \textit{2-split}(s, t) - \textit{journey}) \\ \leq \; & \sum_{s,t \in V(G)} Pr(\nexists \textit{2-split}(s, t) - \textit{journey}) \\ \leq \; & n(n-1) \frac{2}{n^{\frac{\rho}{2}}} \\ < \; & \frac{2}{n^2}, \text{ for } \rho > 8 \end{aligned}$$

We conclude that almost surely $r(n) = \rho \log n, \; \rho > 8$, random labels per edge[3] suffice for an assignment to strongly guarantee temporal reachability with high probability in the star graph G_n.

(b) Surprisingly, we now show that, for the assignment to strongly guarantee temporal reachability whp, $r(n)$ has to be $\Omega(\log n)$. Suppose that, through an assignment L, each edge of G_n now receives $k = \frac{\log n}{\beta(n)}$ random labels (from the set $\{1, 2, \ldots, n\}$), where $\beta(n) \to +\infty$ as $n \to +\infty$. Consider two fixed leafs $u_1, u_2 \in V(G)$ and let $e_1 = \{u_1, c\}, e_2 = \{c, u_2\}$ and E_{u_1,u_2} be the following event:

$$\begin{aligned} & \textit{There exists no } (u_1, u_2)\textit{-journey in } (G_n, L) \\ \equiv \; & \exists a \in \{2, 3, \ldots, n-2\} : \textit{ all of } e_1\textit{'s labels fall} \\ & \textit{within } (a, n) \textit{ and all of } e_2\textit{'s labels within} (0, a) \end{aligned}$$

Given a specific $a \in \{2, 3, \ldots, n-2\}$, the probability that all of e_1's labels fall within (a, n) *and* all of e_2's labels fall within $(0, a)$ is:

$$\begin{aligned} & Pr(\text{all of } e_1\text{'s labels fall within } (a, n) \\ & \quad \text{and all of } e_2\text{'s labels fall within } (0, a)) \\ & = (1 - \frac{a}{n})^k (\frac{a}{n})^k \end{aligned}$$

Now, the probability that event E_{u_1,u_2} occurs is at least as large as the probability that all of e_1's labels fall within (a, n) *and* all of e_2's labels fall within $(0, a)$, for a specific $a \in \{2, 3, \ldots, n-2\}$, e.g., for $a = \frac{n}{2}$. So:

$$\begin{aligned} Pr(E_{u_1,u_2}) &\geq Pr(e_1\text{'s labels fall within } (\frac{n}{2}, n) \text{ and} \\ & \qquad e_2\text{'s labels fall within } (0, \frac{n}{2})) \\ &= (\frac{1}{2})^k (\frac{1}{2})^k = (\frac{1}{2})^{2k} = \frac{1}{2^{2k}} \end{aligned}$$

The probability that no a, such that all of e_1's labels fall within (a, n) *and* all of e_2's labels fall within $(0, a)$, exists is:

$$Pr(\neg E_{u_1,u_2}) = 1 - Pr(E_{u_1,u_2}) \leq 1 - \frac{1}{2^{2k}}$$

Note that also $Pr(E_{u_2,u_1}) \geq \frac{1}{2^{2k}}$ (by symmetry).

In the star graph G_n, we can group the leafs in $\lfloor \frac{n-1}{2} \rfloor = n'$ disjoint pairs $\{u_1, u_2\}, \{u_3, u_4\}, \ldots, \{u_{n'-1}, u_{n'}\}$ defining the paths $(start, center, end)$ $P_1 = (u_1, c, u_2), P_2 = (u_3, c, u_4), \ldots, P_{n'} = (u_{n'-1}, c, u_{n'})$. These paths receive *independent labels since no edges of P_i overlap with any edge of P_j, $i, j = 1, 2, \ldots, n'$, $i \neq j$*. So:

$$\begin{aligned} Pr(\neg E \text{ holds for all these pairs}) &\leq (1 - \frac{1}{2^{2k}})^{n'} \\ &\leq e^{-\frac{n'}{2^{2k}}} \end{aligned}$$

i.e.,

$$\begin{aligned} & Pr(\text{there are temporal paths between each } (u_i, u_j), \\ & \quad i = 1, 3, \ldots, j = 2, 4, \ldots) \leq e^{-\frac{n'}{2^{2k}}} \end{aligned}$$

[3]the labels are selected uniformly at random from the set $L_0 = \{1, 2, \ldots, n\}$ and the edges receive their labels independently of one another

Since $k = \frac{\log n}{\beta(n)}$, we get:

$$\frac{n'}{2^{2k}} = \frac{\lfloor \frac{n-1}{2} \rfloor}{2^{\frac{2\log n}{\beta n}}} = \lfloor \frac{n-1}{2} \rfloor \left(4^{-\log n}\right)^{\frac{1}{\beta(n)}}$$
$$= \lfloor \frac{n-1}{2} \rfloor \left(\frac{1}{n^2}\right)^{\frac{1}{\beta(n)}}$$

So:

$$\frac{n'}{2^{2k}} \geq \frac{n}{3}\left(\frac{1}{n^2}\right)^{\frac{1}{\beta(n)}} > \log n \qquad (3)$$

Relation (3) holds, since:

$$\frac{n}{3}\left(\frac{1}{n^2}\right)^{\frac{1}{\beta(n)}} > \log n \Leftrightarrow \left(\frac{3\log n}{n}\right)^{\beta(n)} < \frac{1}{n^2}$$

But:

$$\left(\frac{3\log n}{n}\right)^{\beta(n)} < \left(\frac{1}{\sqrt{n}}\right)^{\beta(n)} = \left(\frac{1}{n}\right)^{\frac{\beta(n)}{2}} < \frac{1}{n^2},$$

because $\frac{\beta(n)}{2} > 2$. So, by relation (3), we have:

$$-\frac{n'}{2^{2k}} < -\log n \Rightarrow e^{-\frac{n'}{2^{2k}}} < e^{-\log n} = \frac{1}{n}$$
$$\Rightarrow \quad Pr(\exists \text{temporal paths } P_i, \forall i = 1, 2, \ldots, n') \leq \frac{1}{n}$$

Thus, it must be:

$$r(n) > \frac{\log n}{\beta(n)}, \text{ for every such } \beta(n) \to +\infty$$

i.e.,

$$PoR(G) > \frac{\log n}{2\beta(n)}, \text{ for all } \beta(n) \to +\infty$$
$$\Rightarrow \quad PoR(G) \geq c\log n - o(n), \text{ for some } c > 0$$

By parts (a) and (b) we obtain that $PoR(G) = \Theta(\log n)$, for the star graph G_n. $\square$

5. THE PRICE OF RANDOMNESS IN GENERAL GRAPHS

Let $G = (V, E)$ be an arbitrary connected graph of $|V| = n$ vertices. Let the set of available times (labels) be $S = \{1, 2, \ldots, q\}$. Let $d(G) = diam(G)$ be the diameter of G. Clearly, for an assignment L to preserve the reachability of G, there must be a label in (G, L) at least equal to G's diameter, since otherwise a path in G realizing the diameter cannot be made a journey. So, $q \geq d(G)$.

For each edge e of G, consider a structure $s(e)$ being a sequence of boxes $B_1(e), B_2(e), \ldots, B_{d(G)}(e)$ (see Figure 3).

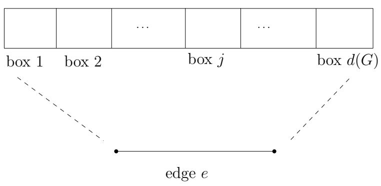

Figure 3: Structure $s(e)$.

Let each Box_i of e be assigned to a corresponding range (sequence) $L_i(e)$ of labels, each of size (#*labels*) equal to $\lambda = \frac{q}{d(G)}$, so that:

$$\forall i = 1, 2, \ldots, d(G),$$
$$Box_i \text{ corresponds to } L_i(e) = \{(i-1)\lambda + 1, \ldots, i\lambda\}$$

CLAIM 1. *If $\forall e \in E(G), \forall Box_i(e)$ we put in $Box_i(e)$ one of the labels of $L_i(e)$, then temporal reachability is guaranteed in G.*

PROOF. For any s, t, any shortest path p *from* s *to* t will be of length $|p| \leq d(G)$. Any edge e may be at any "place" in p (first, second, ..., even last) or not belong to p at all. The journey from s to t is the path $p = \Big(e_{p_1} = \{s, u_1\}, e_{p_2} = \{u_1, u_2\}, \ldots, e_{p_{last}} = \{u_{|p|-1}, t\}\Big)$ with the label per edge e_{p_i} in the $(Box_{p_i}(e_{p_i}))$. $\square$

Note now that when we assign a random label in edge e (drawn uniformly, independently from L), the probability that this label falls in $Box_i(e)$ is exactly $\frac{\lambda}{q}$. For r random labels assigned to e, the probability that none of them falls is $Box_i(e)$ is $\left(1 - \frac{\lambda}{q}\right)^r$. Thus, the probability of the event:

$$A(e) = \text{"there exists a box of } e \text{ without a label"}$$

is at most $d(G)\left(1 - \frac{\lambda}{q}\right)^r$.

Clearly,

$$\left(1 - \frac{\lambda}{q}\right)^r \leq e^{-\frac{\lambda r}{q}}$$

and since $d(G) \leq n$, it is enough to have $\frac{\lambda r}{q} > 2\log n$ to get $d(G)\left(1 - \frac{\lambda}{q}\right)^r < n\frac{1}{n^2} = \frac{1}{n}$. But,

$$\frac{\lambda r}{q} > 2\log n \Leftrightarrow r > 2d(G)\log n$$

So, we have shown the following.

THEOREM 7. *If we assign any $r > 2d(G)\log n$ random labels at each edge of G, then temporal reachability is guaranteed with high probability. So,*

$$r(n) \leq 2d(G)\log n + \varepsilon, \textit{ for some } \varepsilon > 0.$$

Since $OPT \geq n-1$ (at least $n-1$ edges must be labelled in order to have a labelled spanning tree), we get the following.

THEOREM 8. *For every connected graph G, it holds that:*

$$PoR(G) \leq (2d(G)\log n + \varepsilon)\frac{m}{n-1}, \textit{ for some } \varepsilon > 0.$$

Note. The upper bound on the PoR for general graphs can be improved slightly by the *Coupon Collector theorem.*

6. CONCLUSIONS AND FURTHER RESEARCH

In this work, we initiated research on networks with *sparse random* availability of links. The subject of designing the availability of a net (by combining random availabilities and optimal local availabilities) is a subject of our current research.

7. REFERENCES

[1] D. Angluin, J. Aspnes, Z. Diamadi, M. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, pages 235–253, 2006.

[2] C. Avin, M. Koucký, and Z. 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)*, pages 121–132, 2008.

[3] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 155–164, 2008.

[4] B. Bollobás. *Modern Graph Theory*. Springer, 1998.

[5] B. Bollobás. *Random Graphs, 2nd Edition*. Cambridge University Press, 2001.

[6] B.-M. Bui-Xuan, A. Ferreira, and A. Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *International Journal of Foundations of Computer Science*, 14(2):267–285, 2003.

[7] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems (IJPEDS)*, 27(5):387–408, 2012.

[8] A. E. F. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time of edge-markovian evolving graphs. *SIAM Journal on Discrete Mathematics (SIDMA)*, 24(4):1694–1712, 2010.

[9] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the 6th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 1–12, 1987.

[10] C. Dutta, G. Pandurangan, R. Rajaraman, Z. Sun, and E. Viola. On the complexity of information spreading in dynamic networks. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 717–736, 2013.

[11] R. Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proceedings of the 18th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 148–157, 2006.

[12] R. Elsässer and T. Sauerwald. The power of memory in randomized broadcasting. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 218–227, 2008.

[13] P. Erdös and A. Rényi. On random graphs, i. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[14] L. Fleischer and É. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3-5):71–80, 1998.

[15] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985.

[16] C. Gavoille, D. Peleg, S. Perennes, and R. Raz. Distance labeling in graphs. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 210–219, 2001.

[17] R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 565–574, 2000.

[18] M. Katz, N. A. Katz, A. Korman, and D. Peleg. Labeling schemes for flow and connectivity. *SIAM Journal on Computing*, 34(1):23–40, 2004.

[19] 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.

[20] F. Kuhn, N. A. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2010.

[21] G. B. Mertzios, O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Temporal network optimization subject to connectivity constraints. In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP), Part II*, pages 657–668, 2013.

[22] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Mediated population protocols. *Theoretical Computer Science*, 412(22):2434–2450, 2011.

[23] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.

[24] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Causality, influence, and computation in possibly disconnected synchronous dynamic networks. In *OPODIS*, pages 269–283, 2012.

[25] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer, 2002.

[26] R. O'Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proceedings of the 2005 joint workshop on Foundations of Mobile Computing (DIALM-POMC)*, pages 104–110, 2005.

[27] B. Pittel. Linear probing: The probable largest search time grows logarithmically with the number of records. *J. Algorithms*, 8(2):236–249, 1987.

[28] C. Scheideler. Models and techniques for communication in dynamic networks. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 2002.