

Accepted Manuscript

Ephemeral networks with random availability of links: The case of fast networks

Eleni C. Akrida, Leszek Gaşieniec, George B. Mertzios, Paul G. Spirakis

PII: S0743-7315(15)00187-2

DOI: <http://dx.doi.org/10.1016/j.jpdc.2015.10.002>

Reference: YJPDC 3441

To appear in: *J. Parallel Distrib. Comput.*

Received date: 10 December 2014

Revised date: 25 August 2015

Accepted date: 14 October 2015

Please cite this article as: E.C. Akrida, L. Gaşieniec, G.B. Mertzios, P.G. Spirakis, Ephemeral networks with random availability of links: The case of fast networks, *J. Parallel Distrib. Comput.* (2015), <http://dx.doi.org/10.1016/j.jpdc.2015.10.002>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Highlights

- We define the Temporal Diameter (TD) and the expected TD of temporal networks with random availabilities of links
- We define fast and slow random temporal networks, according to their expected TD
- Random temporal networks on instances of dense (directed) random graphs $G_{n,p}$ are fast
- We introduce the critical availability as a measure of periodic availability of links
- We give a lower and an upper bound on the critical availability

Ephemeral Networks with Random Availability of Links: The case of fast networks[☆]

Eleni C. Akrida^{a,*}, Leszek Gašieniec^a, George B. Mertzios^b, Paul G. Spirakis^{a,c,d}

^a*Department of Computer Science, University of Liverpool
Ashton Building, Ashton Street, Liverpool L69 3BX, UK*

^b*School of Engineering and Computing Sciences, Durham University
Science Site, South Road, Durham DH1 3LE, UK*

^c*Department of Computer Engineering & Informatics, University of Patras
School of Engineering, GR 265 00, Patras, Greece*

^d*Computer Technology Institute
N. Kazantzakis str., University Campus, 265 00 Rio, Patras, Greece*

Abstract

We consider here a model of *temporal* networks, the links of which are available only at certain moments in time, chosen randomly from a subset of the positive integers. We define the notion of the Temporal Diameter of such networks. We also define *fast* and *slow* such temporal networks with respect to the expected value of their temporal diameter. We then provide a partial characterisation of fast random temporal networks. We also define the critical availability as a measure of periodic random availability of the links of a network, required to make the network fast. We finally give a lower bound as well as an upper bound on the (critical) availability.

Keywords: temporal networks, random input, diameter, availability

[☆] A previous version of this work appeared in the 26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '14.

*Please address correspondence to Eleni C. Akrida

Email addresses: Eleni.Akrida@liverpool.ac.uk (Eleni C. Akrida),
L.A.Gasieniec@liverpool.ac.uk (Leszek Gašieniec),
George.Mertzios@durham.ac.uk (George B. Mertzios),
P.Spirakis@liverpool.ac.uk (Paul G. Spirakis)

1. Introduction

1.1. Networks with (periodic) random availabilities

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. [21] and from the sequel [24], 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 (di)graph $G = (V, E)$ (*the links of which can become available*) and an assignment L assigning a set of discrete labels to each edge (resp. arc) of G .

We consider here both the single-label-per-edge model of [21] and the multi-labelled one, which allows links to be available at multiple times (i.e., more than one label per edge).

We first define *temporal networks* by assigning a set $L_e \subseteq \mathbb{N}$ of *time-labels* to every edge e of a (di)graph $G = (V, E)$.

Definition 1 (Temporal Network). *Let $G = (V, E)$ be a (di)graph. A temporal network on G is an ordered triple $N(G) = (V, E, L)$, where $L = \{L_e \subseteq \mathbb{N} : e \in E\}$ is an assignment of labels on the edges of G .*

The values assigned to each edge of the underlying graph, G , are called time labels of the edge and indicate the discrete moments in time 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).

A continuous-time model has been studied in the context of temporal network flows with time-varying capacities [14, 22], which naturally require such a modeling (think, for example, about pumping crude oil into pipes). There is no unique approach to the continuous-time model of temporal networks. The simplest one is to assume continuous intervals of availability and discretize them; in this case, the techniques and results would be similar to the ones produced in this work. The same effect was observed in [14], where the authors show how a number of results and algorithms for the discrete model of temporal flows can be carried over to the analogous continuous model. Another continuous-time modeling of temporal networks would be to assume intervals of availability on the edges and a transit time, i.e., delay, for each edge; this would dramatically change the results and would potentially require a different approach.

In the context of this paper, we focus on the discrete-time case and we mainly study *random temporal networks*, in which the labels assigned to the edges of the underlying graph are chosen at random from a set of available time labels. Many distributed systems are highly uncoordinated, meaning that the nodes may only have limited knowledge about the network topology; they may, for example, only be aware of the number of vertices and/or the diameter of the network. A way to treat such systems, where global coordination over availability of all edges is unavailable or nonexistent, is to allow adjacent vertices to agree on a number of random times of availability of the edge joining them. This is a local operation that uses local, random availability to replace the lack of global knowledge and global coordination.

In particular, we focus here on temporal networks, the labels of the edges of which are integers randomly chosen by the following procedure:

Label Assignment Procedure(LAP)

Let k, α, ρ be positive integers greater than or equal to 1, with $\rho \leq \alpha$. Let $N_{l_{max}}$ be the set $N_{l_{max}} = \{1, 2, \dots, l_{max}\}$, where $l_{max} = \alpha k$.

- (I) $N_{l_{max}}$ is partitioned into k consecutive *periods* $\Pi_1, \Pi_2, \dots, \Pi_k$, each of α consecutive integers.
- (II) For every edge (arc), e , of $G = (V, E)$ and every given period Π_j , we draw ρ integers independently and uniformly at random from the set $\{(j-1)\alpha + 1, (j-1)\alpha + 2, \dots, j\alpha\}$ and assign the set of those integers (omitting duplicates), denoted by $L(j, e)$, to e . The label set of the edge e , L_e , is the union of $L(j, e)$ over all periods.

We call ρ the *availability* of labels per period, or simply the (period) availability. We call $\frac{\rho}{\alpha}$ the *density* of the periodic random availabilities. We call α the *length of the period*.

Definition 2 (Ephemeral random (temporal) network). Let $G = (V, E)$ (resp. $D = (V, E)$) be a graph (resp. digraph) where V is the set of vertices, $|V| = n$, and E is the set of edges (resp. arcs). Let each edge (resp. arc) be assigned labels (availability instances) by the Label Assignment Procedure, with k periods, Π_i , $i = 1, \dots, k$, length of any period equal to $\alpha \geq 1$ and density $\frac{\rho}{\alpha}$. We call the resulting network an ephemeral random temporal network $N(k, \alpha, \rho)$.

If $\alpha = n$, then the network is called a normalized ephemeral random (temporal) network, $N(k, n, \rho)$.

The number $l_{max} = \alpha k$ is called the (maximum) lifetime of $N(k, \alpha, \rho)$.

Note that such networks are indeed *ephemeral*: no link is available after l_{max} .

The set L_e of labels of any edge (arc) e is the set of exactly all the moments in time at which e is available (for use). One might associate the periods to, say, “months” and the sets $L(j, e)$ (for period Π_j and edge e) as the sets of (randomly chosen) days in which e is available.

In many 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.

1.2. Journeys, the Temporal Diameter and Sparsity

Definition 3 (Time edge). Let $e = \{u, v\}$ (resp. $e = (u, v)$) be an edge (resp. arc) of the underlying (di)graph of a temporal network and consider a label $l \in L_e$. The ordered triplet (u, v, l) is called time edge.

Note that an undirected edge $e = \{u, v\}$ is associated with $2 \cdot |L_e|$ time edges, namely both (u, v, l) and (v, u, l) for every $l \in L_e$.

Definition 4 (Journey). A temporal path or journey j from a vertex u to a vertex v ((u, v) -journey) is a sequence of time edges $(u, u_1, l_1), (u_1, u_2, l_2), \dots, (u_{k-1}, v, l_k)$, such that $l_i < l_{i+1}$, for each $1 \leq i \leq k - 1$. We call the last time label of journey j , l_k , arrival time of the journey.

Definition 5 (Foremost journey). A (u, v) -journey j in a temporal network is called foremost journey if its arrival time is the minimum arrival time of all (u, v) -journeys’ arrival times, under the labels assigned to the underlying graph’s edges.

Definition 6 (Temporal distance). The temporal distance of a (target) vertex v from a (source) vertex u is the arrival time of the foremost (u, v) -journey and is denoted by $\delta(u, v)$.

Definition 7 (Temporal diameter). Let $N(G)$ be a temporal network. The temporal diameter of $N(G)$, denoted by TD , is the maximum, over all ordered pairs of vertices s, t , $s \neq t$, $\delta(s, t)$.

Note that, for a given labelling L , even for a connected graph G , there may be two vertices s and t so that there is no journey from s to t . In this case, the temporal diameter is $TD(N(G)) = +\infty$

Definition 8 (Expected temporal diameter). Consider instances $N(G)$ of an ephemeral random temporal network $N(k, \alpha, \rho)$. Each instance corresponds to a labelling drawn by the random choices. The expected temporal diameter of N is the expectation $E[TD]$ over all such instances. We denote it by $ETD(N)$.

Note that, for any $N(k, \alpha, \rho)$, over a connected graph G , with k being at least equal to the diameter of G , the $ETD(N)$ is *finite* and it is at most $l_{max} = \alpha k$, since every edge from s to t , $s, t \in V(G)$, will get at least a label of availability in each period. That is, for any connected graph $G = (V, E)$, a number of periods $k \geq \text{diam}(G)$, where $\text{diam}(G)$ is the diameter, suffices to bound the temporal distance between any two vertices $u, v \in V(G)$. This holds even when $\rho = 1$, and even when the single label chosen per period in this case is selected arbitrarily. So, we get the following remark:

Remark 1. Let G be a connected graph (resp. a strongly connected digraph). Let $\text{diam}(G)$ be the diameter of G . Consider the ephemeral random temporal network $N(k, \alpha, \rho)$ with $k \geq \text{diam}(G)$ and $\alpha, \rho \geq 1$. Then, we have that in any instance $I = N(G)$ of it and for every pair of vertices s, t , $\delta_I(s, t) \leq \text{diam}(G)$, where $\delta_I(s, t)$ the temporal distance of t from s in I . Thus, $ETD(N) \leq \alpha \cdot \text{diam}(G)$.

In this work, we consider “sparse” ephemeral random temporal networks (and also the case of a single period) and wish to characterize “fast” random temporal networks.

Definition 9 (Sparse). An ephemeral random temporal network $N(k, \alpha, \rho)$ is called sparse if $\frac{\rho}{\alpha} \in o(1)$.

Definition 10 (Fast). A normalized ephemeral random temporal network $N(k, n, \rho)$ is called fast if $ETD(N) = O(\log n)$.

Clearly, fast networks need just one period for all-to-all journeys to exist with high probability. A major goal of this work is to characterize fast and sparse ephemeral random temporal networks.

Definition 11 (Efficient). An ephemeral (normalized) random temporal network $N(k, \alpha, \rho)$ is called efficient if $ETD(N) = o(\alpha \cdot \text{diam}(G))$, where G is the underlying connected graph (resp. strongly connected digraph).

Definition 12 (Slow). An ephemeral (normalized) random temporal network $N(k, \alpha, \rho)$ is called slow if $ETD(N) = \Theta(\alpha \cdot \text{diam}(G))$, where G is the underlying connected graph (resp. strongly connected digraph).

To further motivate some of the questions raised in this work, consider a very *hostile* clique, G , the edges of which are usually guarded. Whenever an edge is guarded it is impossible to pass a message through. We may pass a message to a neighbour in G only when the link to this neighbour is unguarded (i.e., available). This abstraction of the hostile clique can model various different situations in distributed computing applications. For example in a (very noisy) wireless communication network, whenever the level of noise and interference drops below some level we may consider that each link becomes available only at some random moments in time. Sensor networks with random link failures, where the transmitted data may be corrupted by additive noise, have been studied in [17, 18]. Other applications may include a distributed setting where taxis in a city appear sporadically at random (i.e., unpredictable) points within a period of time, or where different suppliers may put specific commercial items on offer (e.g., air tickets, books, clothes) in unpredictable times, etc.

Now, let us assume that each edge of the (hostile) clique will become available only at a single *random moment* in every period and also $k = 1$, i.e., only one period exists. Let us examine the normalized case, where $\alpha = n$. After time n , no link of the clique is ever available. Such a (random) time of availability 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 and on the average. 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. This means that a sparse normalized random temporal clique is *fast*.

1.3. Our results

In this work, we introduce the model of random ephemeral temporal networks $N(k, \alpha, \rho)$ of k periods of availability of links and ρ moments of

availability per period, chosen uniformly at random from a set of α available labels per period. We define *sparse*, *fast*, *slow* and *efficient* such networks with respect to their expected Temporal Diameter, which is also defined here. We give a partial characterisation of the fast networks of the form $N(k, n, 1)$. Namely, they include the class of networks on instances, $D_{n,p}$, of the directed Erdős-Renyi graphs, with probability of existence of a directed link equal to $p \geq \varepsilon$, for some constant ε independent of the size of the underlying graph. We also give an example of a slow network, namely an ephemeral random temporal network on the line graph. We define the critical availability, ρ^* , of randomly available instances of a link during a period so that the resulting network is fast. We show that ρ^* can be unbounded, using the example of the star graph where we cannot bound the critical availability from above by a constant. Finally, we present a general (non-constant) upper bound on ρ^* . For convenience, we provide a glossary of symbols and notation in Section 6.

1.4. Previous work

A preliminary version of this paper appeared in ACM SPAA 2014[1].

1.4.1. 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 [10]. In [10], 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 [30] 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. [19] 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 [12]. 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 $1 - n^{-1}$ [4, 13].

The network model adopted in this paper resembles the Random Phone-Call model to some extent, however, it is essentially different. The depen-

dence of the temporal diameter on the 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 some (maybe even a single) random moments of existence, by the input. A node can send via this link only at that moment. That is, random availability of links is not a part of our algorithmic techniques and can not be used at arbitrary time steps.

1.4.2. 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., [28].

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 [21] as well as the multi-labelled model studied in [24]. The prior results of [21, 24] do not consider randomness at all, and therefore are different in nature to this work. The initial paper [21] considers the case of one (non-random) label per edge and examines shortest journey algorithms. The second paper [24] 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 [14].

Dynamic Distributed Networks. In recent years, there is a growing interest in distributed computing systems that are inherently dynamic [2, 3, 5, 7, 8, 9, 11, 23, 25, 26, 27, 29, 31].

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, 20].

2. Sparse and fast ephemeral random temporal networks

2.1. The case of an underlying Erdős-Renyi random graph

Definition 13 (Erdős-Renyi graphs). *An instance of $G_{n,p}$ (resp. $D_{n,p}$) is formed when for every pair of vertices (resp. ordered pair of vertices) u, v among a total number of n vertices, the edge $\{u, v\}$ (resp. the arc (u, v)) is chosen to exist with probability p independently of any other edge (resp. arc).*

Let us consider the case where the underlying graph is an instance of the Erdős-Renyi $G_{n,p}$ ($D_{n,p}$ for digraphs) with $p \geq \varepsilon$, where ε is a constant independent of n . We will show:

Theorem 1. *Consider sparse random normalized temporal networks $N(k, n, \rho)$, $1 \leq \rho \leq c$, where c is an integer larger than or equal to 1 and independent of n . Let the underlying graph be any strongly connected instance of the $D_{n,p}$ model, where $p \geq \varepsilon$ ($\varepsilon > 0$ a constant). Then, $ETD(N) = O(\log n)$, i.e., all such networks are fast.*

Proof. Clearly, if N_1 is such a network and N_2 is defined on the same $D_{n,p}$ but with $\rho = c = 1$, then $ETD(N_1) \leq ETD(N_2)$, since the increased availability of each edge per period may only introduce better journeys, i.e., journeys with smaller arrival time, to the network. So, we fix $\rho = c = 1$. We will first consider the case, where the total number of periods is $k = 1$, and show that, for any two particular vertices s and t , it holds that:

$$Pr[\delta(s, t) = \Theta(\log n)] \geq 1 - \frac{1}{n^4}$$

The result for the expectation of the maximum value of $\delta(s, t)$ will follow, by noticing that in this case, $\delta(s, t) \leq n \cdot \text{diam}(D)$, where D is any strongly connected instance of $D_{n,p}$, and thus $\delta(s, t) \leq n^2$. The probability space of the examined case is the joint space S obtained by two independent experiments:

- (a) Draw an instance D of $D_{n,p}$.
- (b) Assign a single label chosen uniformly at random from $\{1, 2, \dots, n\}$ to each arc of D independently of the assignments to the other arcs.

We first show the following lemma:

Lemma 1. For any two particular vertices s, t , $s \neq t$, and for the probability space S , it holds:

$$Pr[\delta(s, t) \leq \gamma \log n] \geq 1 - \frac{3}{n^5}, \text{ for some constant } \gamma > 1$$

Proof. We will use the method of *deferred decisions* to prove our result. This means that when we first examine a possible arc (u, v) , the probability that the arc exists is p and the probability that it is assigned a particular label l , given its existence, is $Pr[(u, v) \text{ is assigned label } l / (u, v) \text{ exists}] = \frac{1}{n}$. So, if Δ is any sub-set $\{i, i+1, \dots, j\}$, $j > i$ of $\{1, 2, \dots, n\}$ and \mathcal{E} is the event that a message can be sent via (u, v) in a specific random instance in Δ , then we have:

$$\begin{aligned} Pr[\mathcal{E}] &= Pr[(u, v) \text{ exists}] \cdot Pr[\text{the label of } (u, v) \text{ is in } \Delta / (u, v) \text{ exists}] \\ &= \frac{\Delta}{n} p \end{aligned} \quad (1)$$

We will use relation (1) repeatedly below. We now fix p to be $p = \varepsilon$, for some $0 < \varepsilon < 1$ constant. If p is larger, then the possibility of existence of every arc increases, thus the probability $Pr[\delta(s, t) \leq a]$ increases, for all a and all ordered pairs of vertices s, t which are connected via a (directed) path in the underlying digraph.

We will provide an algorithmic construction (Algorithm 1) which constructs a journey with logarithmic arrival time from any given vertex s to any given vertex t , with high probability. Let $d \in \Theta(\log n)$ and γ_i , $i = 1, 2$, suitable constants. We will only consider labels which belong in the following sequence of consecutive intervals (all of which are in the first period):

$$\begin{aligned} \Delta_1 &= (0, \gamma_1 \log n] \\ \Delta_2 &= (\gamma_1 \log n, \gamma_1 \log n + \gamma_2] \\ \Delta_3 &= (\gamma_1 \log n + \gamma_2, \gamma_1 \log n + 2\gamma_2] \\ &\dots \\ \Delta_{d+1} &= (\gamma_1 \log n + (d-1)\gamma_2, \gamma_1 \log n + d\gamma_2] \\ \Delta^* &= (\gamma_1 \log n + d\gamma_2, 2\gamma_1 \log n + d\gamma_2] \\ \Delta'_{d+1} &= (2\gamma_1 \log n + d\gamma_2, 2\gamma_1 \log n + (d+1)\gamma_2] \\ &\dots \end{aligned}$$

$$\begin{aligned}\Delta'_2 &= (2\gamma_1 \log n + (2d - 1)\gamma_2, 2\gamma_1 \log n + 2d\gamma_2] \\ \Delta'_1 &= (2\gamma_1 \log n + 2d\gamma_2, 3\gamma_1 \log n + 2d\gamma_2]\end{aligned}$$

Note that the length of the total time interval we consider, described by the span of the union of all the above intervals, is $\Theta(\log n)$. Also, note that any directed path $(s, v_1, v_2, \dots, v_{d+1}, v^*, v'_{d+1}, \dots, v'_2, v'_1, t)$ with consecutive labels -one per edge- $\lambda_i \in \Delta_i, \lambda^* \in \Delta^*, \lambda'_i \in \Delta'_i, i = 1, 2, \dots, d+1$, is a journey and its arrival time is $\Theta(\log n)$. Algorithm 1 demonstrates the algorithmic construction of journeys like that¹. In Algorithm 1, $\Gamma_i(s)$ denotes the set of vertices v such that there is a vertex in the previous Γ set, namely $\Gamma_{i-1}(s)$, that has an outgoing edge to v with label in the interval Δ_i ; by $\Gamma_0(s)$ we denote the set that contains exactly s . So, roughly speaking, $\Gamma_i(s)$ is the set of vertices in *depth* i from s in the expansion process that Algorithm 1 demonstrates. Respectively, $\Gamma'_i(t)$ denotes the set of vertices v such that there is a vertex in the previous Γ' set, namely $\Gamma'_{i-1}(t)$, that has an incoming edge from v with label in the interval Δ'_i ; by $\Gamma'_0(t)$ we denote the set that contains exactly t . So, roughly speaking, $\Gamma'_i(t)$ is the set of vertices in *depth* i starting from t and going backwards towards s in the expansion process that Algorithm 1 demonstrates.

Algorithm 1 The Expansion Process algorithm

Input: An instance, $N(1, n, 1)$, of an ephemeral sparse random temporal network, the underlying digraph, $D = (V, E)$, of which is an instance of $D_{n, \varepsilon}$, and vertices $s, t, s \neq t$

- 1: $\Gamma_0(s) = \{s\}$;
 - 2: **for** $i = 1, \dots, d + 1$ **do**
 - 3: $\Gamma_i(s) = \{v \in V : \exists w \in \Gamma_{i-1}(s) \text{ such that } l_{wv} \in \Delta_i\}$;
 - 4: **end for**
 - 5: $\Gamma'_0(t) = \{t\}$;
 - 6: **for** $i = 1, \dots, d + 1$ **do**
 - 7: $\Gamma'_i(t) = \{v \in V : \exists w \in \Gamma'_{i-1}(t) \text{ such that } l_{vw} \in \Delta'_i\}$;
 - 8: **end for**
 - 9: **if** $\exists u \in \Gamma_{d+1}(s), v \in \Gamma'_{d+1}(t)$ such that $l_{uv} \in \Delta^*$ **then**
 - 10: An (s, t) -journey has been created on the concatenation of the directed path from s to u , the arc (u, v) and the directed path from v to t
-

¹in the algorithm, the label assigned to the arc (u, v) is denoted by l_{uv}

Algorithm 1 The Expansion Process algorithm (continued)

11: **return** success and the journey;
 12: **else**
 13: **return** failure;
 14: **end if**

Note 1. *The above construction of a fast journey may also fail if any of the Γ_i, Γ'_i , $i = 1, 2, \dots, d + 1$ is empty.*

We proceed to calculate the probability of success of Algorithm 1. We will denote by p_1 the probability that a particular arc out of s exists and is assigned a label in Δ_1 . So,

$$\begin{aligned} p_1 &= \varepsilon \cdot \frac{\gamma_1 \log n}{n} \\ &= \frac{c_1 \log n}{n}, \text{ where } c_1 = \varepsilon \gamma_1 \end{aligned} \quad (2)$$

In fact, p_1 is also the probability that a particular arc (u, v) , with $u \in \Gamma_{d+1}(s), v \in \Gamma'_{d+1}(t)$, exists and is assigned a label in Δ'_1 .

In addition, denote by p_2 the probability that a vertex v in $\Gamma_{i-1}(s)$, $i \geq 2$ (resp. a v in $\Gamma'_{i-1}(t)$, $i \geq 2$) has a particular outgoing arc (resp. incoming arc) with a label in the interval Δ_i (resp. Δ'_i). So,

$$\begin{aligned} p_2 &= \varepsilon \cdot \frac{\gamma_2}{n} \\ &= \frac{c_2}{n}, \text{ where } c_2 = \varepsilon \gamma_2 \end{aligned} \quad (3)$$

Note 2. *In the following analysis, we reveal each possible arc and the 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.*

Algorithm 1 describes a (limited) 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 in 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 in Δ_2 , etc. We show that, with high probability, there is a journey from s to t through vertices in the consecutive Γ_i sets.




Figure 1: The Expansion Process.

(a) **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 label from s to each one of them is in Δ_1 . Note that the probability of an arc $(s, u), u \in V$ existing and having a label in Δ_1 is exactly $p_1 = \frac{c_1 \log n}{n}$.

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

Lemma 2. *It holds that:*

$$\Pr(\mathcal{E}_1) = \Pr\left(|\Gamma_1(s)| \in \left(\frac{1}{2}, \frac{3}{2}\right)(c_1 \log n(1 - \frac{1}{n}))\right) \geq 1 - \frac{1}{n^6}$$

Proof. 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 $B(N_0, p_1)$, where $N_0 = n-1$, $\forall \beta \in (0, 1)$, it holds:

$$\Pr(\#successes \in (1 \pm \beta)N_0p_1) \geq 1 - e^{-\frac{\beta^2}{2}N_0p_1}$$

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

$$\begin{aligned}
Pr(\#successes \in (\frac{1}{2}, \frac{3}{2})N_0p_1) &\geq 1 - e^{-\frac{1}{8}N_0p_1} \\
&\geq 1 - e^{-\frac{1}{8}(c_1 \log n - \frac{c_1 \log n}{n})} \\
&\geq 1 - e^{-\frac{1}{8}(c_1-1) \log n} \\
&\geq 1 - \frac{1}{n^{\frac{c_1-1}{8}}}
\end{aligned}$$

We can choose $c_1 \geq 49$, and thus have $\frac{c_1-1}{8} \geq 6$. This completes the proof of Lemma 2. \square

(b) **The expansion phase until reaching $\Theta(\sqrt{n})$ vertices via journeys.**

We now show that given:

- $|\Gamma_1(s)| \in \Theta(\log n)$, and
- the probability of a potential edge e having a label in a particular interval $\Delta_i, i = 2, \dots, d+1$, provided that e exists, is exactly $p_2 = \varepsilon \cdot \frac{|\Delta_i|}{n} = \frac{c_2}{n}$

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 $\Gamma_j(s), j = 1, 2, \dots, i$ (and the fact that we look for directed edges), i.e.,

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

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 arc $(v, u), v \in \Gamma_i(s)$, is in the interval Δ_{i+1} , provided that (v, u) exists, i.e., equal to:

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

$$= 1 - \left(1 - \frac{c_2}{n}\right)^{|\Gamma_i(s)|}$$

We need the following fact:

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

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

$$(1 - p)^k \leq 1 - kp + \binom{k}{2} p^2 \quad (4)$$

We will show that:

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

and, thus, by relation 4 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, so we have now proven Fact 1. \square

So, we have:

$$\begin{aligned} q &\geq 1 - \left(1 - \frac{c_2|\Gamma_i(s)|}{2n}\right) \\ &= \frac{c_2|\Gamma_i(s)|}{2n} \end{aligned}$$

$$\geq \frac{c_1 c_2 \log n}{16n} = q'$$

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:

$$\Pr(|\Gamma_{i+1}(s)| \in (\frac{1}{2}n_i q, \frac{3}{2}n_i q)) \geq 1 - e^{-\frac{1}{8}n_i q'} \quad (5)$$

But, $n_i \geq n - (\lambda\sqrt{n})d \geq \frac{n}{2}$. So, relation 5 becomes:

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

We will select c_2 so that $\frac{c_1 c_2}{256} \geq 6$. So, with probability at least $1 - \frac{1}{n^6}$, it is:

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

We have proved that the event

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

holds with probability at least $1 - \frac{1}{n^6}$, 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:

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

and also

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

Choose d so that:

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

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

and also:

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

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

The above choice is always possible.

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

$$d \frac{1}{n^6} \leq c' \log n \frac{1}{n^6} \leq \frac{1}{n^5}, \text{ for some } c' > 0$$

Thus, we have shown the following:

Corollary 1. *With probability at least $1 - \frac{1}{n^5}$, the expansion process out of s arrives at $\Theta(\sqrt{n})$ vertices with temporal paths of length $d+1 \in \Theta(\log n)$, consistently labelled in the intervals Δ_i , $i = 1, 2, \dots, d+1$, in time at most $c_1 \log n + dc_2 \in \Theta(\log n)$.*

(c) **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 Corollary 1:

Corollary 2. *The expansion process out of t arrives at $\Theta(\sqrt{n})$ vertices with temporal paths (reverse direction) of length $d + 1 \in \Theta(\log n)$, consistently labelled in the intervals Δ'_i , $i = 1, 2, \dots, d + 1$. Thus, it arrives to each of these vertices in time at most $c_1 \log n + dc_2 \in \Theta(\log n)$ with probability at least $1 - \frac{1}{n^5}$.*

(d) **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 $1 - 2\frac{1}{n^5}$. Note that we just need one arc (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 pair of vertices* (v_1, v_2) , $v_1 \in \Gamma_{d+1}(s)$, $v_2 \in \Gamma'_{d+1}(t)$, the arc appears and has a label in Δ^* with probability exactly:

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

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 Corollaries 1 and 2, it is:

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

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

$$p \geq 1 - \frac{1}{n^5}$$

The probability of any of the events of Corollaries 1 and 2 or event A failing is at most $3\frac{1}{n^5}$.

This completes the proof of Lemma 1. □

For the temporal networks N considered, i.e., any $N(k, n, 1)$ on an instance of $D_{n, \varepsilon}$, we have in fact shown that:

$$Pr[TD(N) \leq \gamma \log n] \geq 1 - \frac{3}{n^3}, \text{ for some constant } \gamma > 1$$

since the probability that there exists a pair of vertices s, t for which the construction fails is at most $n^2 \frac{3}{n^5} = \frac{3}{n^3}$. We will now show the following:

Lemma 3. *$Pr[TD(N) \leq \gamma \log n, \text{ for } N \text{ over strongly connected instances of } D_{n, \varepsilon}] \geq 1 - \frac{4}{n^3}, \text{ for some constant } \gamma > 1.$*

Proof. Consider the event, A_1 , that the temporal diameter of N is $TD(N) \leq \gamma \log n$, for some constant $\gamma > 1$, and the event, A_2 , that the instance D of $D_{n, p}$ is strongly connected. For an instance D of $D_{n, p}$ to be strongly connected, it suffices for the undirected version of D to be connected and for any arc (u, v) that exists in D , that the arc (v, u) also exists. So,

$$Pr[A_2] = Pr[D \text{ of } D_{n, p} \text{ is strongly connected}] \geq Pr[G \text{ of } G_{n, p^2} \text{ is connected}]$$

The connectivity threshold of $G_{n, p}$ is $p_0 = 2 \frac{\log n}{n}$ [6]. In our case, the probability of existence of any undirected edge is $p^2 = \varepsilon^2$, which, for sufficiently large n , is greater than this threshold. Therefore:

$$Pr[A_2] \geq Pr[G \text{ of } G_{n, p^2} \text{ is connected}] \geq 1 - o(1)$$

It is:

$$Pr[A_1] = Pr[A_1/A_2]Pr[A_2] + Pr[A_1/\neg A_2]Pr[\neg A_2]$$

But, the negation of A_2 implies that $Pr[A_1] = 0$. So, the probability that $TD(N) \leq \gamma \log n$, given that N is over a strongly connected instance of $D_{n, \varepsilon}$, is:

$$\begin{aligned} Pr[A_1/A_2] &= \frac{Pr[A_1]}{Pr[A_2]} \\ &\geq \frac{1 - \frac{3}{n^3}}{1 - o(1)} \\ &\geq 1 - \frac{4}{n^3} \end{aligned}$$

This completes the proof of Lemma 3. □

However, for any strongly connected instance D of $D_{n,\varepsilon}$, its diameter is at most $n - 1$ and thus its Temporal Diameter is at most $n(n - 1) < n^2$. So, it holds that:

$$\begin{aligned} ETD(N) &\leq \gamma \log n \Pr[TD(N) \leq \gamma \log n] + n^2(1 - \Pr[TD(N) \leq \gamma \log n]) \\ &\leq \gamma \log n \left(1 - \frac{4}{n^3}\right) + n^2 \frac{4}{n^3} \\ &\leq \gamma \log n + \Theta\left(\frac{1}{n}\right), \text{ for some constant } \gamma > 1 \end{aligned}$$

This completes the proof of Theorem 1. \square

Our analysis holds for any $\varepsilon > 0$ so it also holds for $\varepsilon = 1$. Thus,

Corollary 3. *If $N(k, n, 1)$ is defined on the directed clique then $ETD(N) \leq \gamma \log n$, for some constant $\gamma > 1$.*

Note also that Lemma 1 and Theorem 1 extend trivially to the *undirected* $G_{n,\varepsilon}$ cases; an edge $\{u, v\}$ corresponds to both arcs (u, v) and (v, u) and the analysis is not significantly affected.

Remark 2. *The result of Theorem 1 is a threshold; The network $N(k, n, 1)$ on instances of $D_{n,p}$ or $G_{n,p}$ with $p \geq \varepsilon$ for ε some constant is temporally connected with probability at most equal to the probability that $G_{n,\hat{p}}$ (or $D_{n,\hat{p}}$) is (strongly) connected, where $\hat{p} = \varepsilon \cdot \frac{|\Delta_{max}|}{n}$, where Δ_{max} is the interval of desired labels. If $\Delta_{max} = o(\log n)$, then $G_{n,\hat{p}}$ (and $D_{n,\hat{p}}$) become disconnected almost surely. This implies then that $\Pr[TD(N) = o(\log n)] \xrightarrow{n \rightarrow +\infty} 0$ for the considered temporal networks.*

2.2. Spreading a message in random networks

Consider the following general protocol for broadcasting a message from any vertex s in the network.

for any vertex $u \in V(G)$, $u \neq s$, and any moment $t = 1, 2, \dots$ in time **do**
 if u has received the message from s before t **and** an edge (arc) out of u becomes available at time t **then**
 Send the message through that edge (arc) at time t ;
 end if
end for

Clearly, this protocol spreads the message from s to any vertex, for any temporal ephemeral network $N(k, \alpha, \rho)$ in time at most $ETD(N)$.

2.3. A case of slow networks

In contrast to the fast networks we studied in the previous section, there are also networks that are slow².

Lemma 4. *There exists a slow ephemeral random temporal network $N(k, n, 1)$ on a connected underlying graph G of n vertices.*

Proof. Consider the line graph G of n vertices, that is the graph which itself is a path ($e_1 = \{v_1, v_2\}, e_2 = \{v_2, v_3\}, \dots, e_{n-1} = \{v_{n-1}, v_n\}$). Let $N(k, n, 1)$, $k \geq \text{diam}(G) = n - 1$, be a random temporal network on G and suppose that v_1 wishes to send a message to v_n in the network. Now, let us consider the progress made within the first period, Π_1 , meaning the length of the longest journey starting from v_1 . Denote this progress by x_1 and let λ be an integer from 1 to $n - 1$. Also, denote by l_i , $i = 1, 2, \dots, n - 1$ the label assigned to the edge e_i on the first period. It holds that:

$$\begin{aligned} Pr[x_1 \geq \lambda] &= Pr[l_1 < l_2 < \dots < l_{n-\lambda} / \forall i, j = 1, 2, \dots, \lambda, i \neq j \Rightarrow l_i \neq l_j] \\ &= \frac{1}{\binom{n}{\lambda} \cdot \lambda!} \\ &= \frac{(n - \lambda)!}{n!} \\ &= \frac{1}{n \cdot (n - 1) \cdot \dots \cdot (n - \lambda + 1)} \end{aligned}$$

So, the expected value of the progress of the first period is:

$$\begin{aligned} E[x_1] &= \sum_{\lambda=1}^{n-1} \lambda \cdot Pr[x_1 = \lambda] \\ &= 1 \cdot Pr[x_1 = 1] + 2 \cdot Pr[x_1 = 2] + \dots + (n - 1) \cdot Pr[x_1 = n - 1] \\ &= \sum_{\lambda=1}^{n-1} Pr[x_1 \geq \lambda] \end{aligned}$$

²By definition, an ephemeral random temporal network $N(k, \alpha, \rho)$ on a connected (di)graph G is slow if $ETD(N) \in \Theta(\alpha \cdot \text{diam}(G))$.

$$\begin{aligned}
&= \sum_{\lambda=1}^{n-1} \frac{1}{n \cdot (n-1) \cdot \dots \cdot (n-\lambda+1)} \\
&< \sum_{\lambda=1}^{n-1} \left(\frac{1}{n-\lambda}\right)^\lambda \\
&= \sum_{\lambda=1}^{n-2} \frac{1}{(n-\lambda)^\lambda} + \frac{1}{(n-(n-1))^{n-1}} \\
&\leq \sum_{\lambda=1}^{n-2} \left(\frac{1}{2}\right)^\lambda + 1 \\
&= 3
\end{aligned}$$

Also,

$$\begin{aligned}
E[x_1] &= \sum_{\lambda=1}^{n-1} \frac{1}{n \cdot (n-1) \cdot \dots \cdot (n-\lambda+1)} \\
&> \sum_{\lambda=1}^{n-1} \left(\frac{1}{n}\right)^\lambda \\
&= \frac{1 - \left(\frac{1}{n}\right)^n}{1 - \frac{1}{n}} \\
&= \frac{\left(1 - \frac{1}{n}\right) \left(1 + \frac{1}{n} + \frac{1}{n^2} + \dots + \frac{1}{n^{n-1}}\right)}{1 - \frac{1}{n}} \\
&> 1
\end{aligned}$$

Therefore, the expected length of the longest journey starting from v_1 within the first period, Π_1 , is more than 1 and less than 3. Let $u = v_\mu$, for some $\mu = 2, 3, \dots, n$, be the last vertex of this journey. For the “message” from v_1 to continue its way towards v_n , we would then need to use labels from the next period, Π_2 .

Notice that the assignment of labels for any period is independent of the assignment of labels for any other period. Therefore, the analysis for the progress, x_i , within any period Π_i , $i = 2, 3, \dots, k$, would be the same as the analysis for the progress within Π_1 . Consequently, the expected value of the progress made within each of the k periods is the same as $E[x_1]$, that is more than 1 and less than 3. So, for the message starting from v_1 to reach

v_2 , we would need to use a number of periods between $\frac{n-1}{3}$ and $n-1$, by linearity of expectation. Since we use at least the first $\frac{n-1}{3}$ periods and since each period has a total of n available labels, the expected arrival time of the (v_1, v_2) -journey under consideration will be at least $\frac{n-1}{3} \cdot n$. It will also be at most $n \cdot (n-1)$, since we will not need more than $n-1$ total periods.

Since v_1 and v_n are the two vertices with the largest distance between them in G and because of the fact that G is the line graph, it holds that $ETD(N) = E[\delta(v_1, v_2)] \in \Theta(n^2) \equiv \Theta(n \cdot \text{diam}(G))$. □

3. On efficient random temporal networks

3.1. Introduction

Let $N(k, \alpha, \rho)$ be the ephemeral random temporal network defined on an underlying undirected, connected graph $G = (V, E)$. (Similar considerations hold for strongly connected digraphs). We know that for $k \geq \text{diam}(G)$, it is $ETD(N) \leq \alpha \cdot \text{diam}(G)$. Intuitively, when we assign enough random labels per edge per period, i.e., the density $\frac{\rho}{\alpha}$ increases, then we may reduce the expected temporal diameter. One would even hope, for suitable density, that $ETD(N) \leq \alpha$, i.e., any vertex can reach any other vertex via a journey within the first period with high probability.

Definition 14 (Critical availability, critical density). *Consider the normalized ephemeral random network $N(k, n, \rho)$ on an underlying connected graph $G = (V, E)$, with $k \geq \text{diam}(G)$. Let ρ^* be a positive integer such that:*

- (a) *when $\rho \geq \rho^*$, then $ETD(N) \leq n$, and*
- (b) *if $\rho = o(\rho^*)$, then $\Pr[TD(N) \leq n] = o(1)$.*

We call the density $\frac{\rho^}{n}$ the critical density corresponding to G , and the value ρ^* the critical availability.*

3.2. ρ^* is bounded below by $\log n$.

We will now prove the following theorem:

Theorem 2. *There are connected graphs, G , (even with diameter $\text{diam}(G) = 2$) for which the critical availability is $\rho^* = \Theta(\log n)$.*

Proof. We consider the star graph of n vertices, denoted here by G_n , i.e., the tree of n vertices with one root and $n - 1$ leaves. Let $N = N(k, n, \rho)$ on G_n .

- (a) We first establish that $\rho(n) = \Theta(\log n)$ random labels per edge per period suffice to have $ETD(N_{G_n}) \leq n$. Let $\rho(n) = r \log n$, for some $r > 1$. Denote by c the center vertex of G_n . Now consider two fixed leaves, u_1, u_2 , of G_n .




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

Each of the edges $e_1 = \{u_1, c\}$ and $e_2 = \{c, u_2\}$ is assigned $\rho(n)$ random labels in the first period, Π_1 . 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 $\Delta_1 = (0, \frac{n}{4})$ and the second temporal edge has a label within the interval $\Delta_2 = (\frac{n}{4}, \frac{n}{2})$ (see Figure 2).

The probability that an element of s_1 falls within the interval Δ_1 is $\frac{1}{4}$. So, the probability that no element of s_1 falls within this interval is:

$$\begin{aligned} Pr[\text{no element of } s_1 \text{ falls within the interval } \Delta_1] &= \left(1 - \frac{1}{4}\right)^{r \log n} \\ &= \left(\frac{3}{4}\right)^{r \log n} \\ &= \left(\frac{1}{4}\right)^{r \log n \log 3} \\ &= \left(\frac{1}{n^2}\right)^{r \log 3} \end{aligned}$$

Similarly, the probability that an element of s_2 falls into Δ_2 is $\frac{1}{4}$. So, the probability that no element of s_2 falls within this interval is:

$$\begin{aligned} Pr[\text{no element of } s_2 \text{ falls within the interval } \Delta_2] &= \left(1 - \frac{1}{4}\right)^{r \log n} \\ &= \left(\frac{1}{n^2}\right)^{r \log 3} \end{aligned}$$

Let E be the event that a particular label, l_1 , of e_1 happens to be in s_1 and a particular label, l_2 , of e_2 happens to be in s_2 . Then, by independence of label assignments:

$$Pr[E] = (1 - Pr[l_1 \in \Delta_1]) \cdot (1 - Pr[l_2 \in \Delta_2])$$

So,

$$\begin{aligned} Pr[E] &\geq \left(1 - \left(\frac{1}{n^2}\right)^{r \log 3}\right)^2 \\ &\geq 1 - \frac{2}{n^{2r \log 3}} \end{aligned}$$

Therefore, for $r \log 3 > 2$, i.e., $r > \frac{2}{\log 3}$:

$$\begin{aligned} Pr[\exists s, t \in V(G_n), s \neq t : \exists 2\text{-}split(s, t)\text{-}journey] &\leq \\ &\leq n(n-1) \frac{2}{n^{2r \log 3}} < \\ &< \frac{2}{n^2} \end{aligned}$$

But, for $r > \frac{2}{\log 3}$, it holds that:

$$\begin{aligned} Pr[TD(N) \leq \frac{n}{2}] &\geq Pr[\forall s, t \in V(G_n), s \neq t, \exists(s, t)\text{-}journey] \\ &\geq 1 - \frac{2}{n^2} \end{aligned}$$

Therefore,

$$\begin{aligned} ETD(N) &\leq \frac{n}{2} \left(1 - \frac{2}{n^2}\right) + 2n \frac{2}{n^2} \\ &\leq \frac{n}{2} + \frac{3}{n} \\ &< n \end{aligned}$$

- (b) We now prove condition (b) of Definition 14 for the star graph. In particular, we show that if $\rho = o(\log n)$, then $Pr[TD(N) \leq n] \leq \frac{1}{n}$. Since we look for the TD to be at most n , let us henceforward consider only the first period, Π_1 . Suppose that, through an assignment L , each edge of G_n now receives $k = \frac{\log n}{\beta(n)}$ random labels (from the set $\{1, 2, \dots, n\}$), where $\beta(n) \rightarrow +\infty$ as $n \rightarrow +\infty$. Consider two fixed leaves $u_1, u_2 \in V(G)$ and let $e_1 = \{u_1, c\}, e_2 = \{c, u_2\}$ and E_{u_1, u_2} be the following event:

There exists no (u_1, u_2) -journey in (G_n, L)

$$\equiv \exists a \in \{2, 3, \dots, n-2\} : \text{all of } e_1 \text{'s labels fall within } (a, n] \text{ and all of } e_2 \text{'s labels within } (0, a]$$

Given a specific $a \in \{2, 3, \dots, n-2\}$, 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:

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

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 $a \in \{2, 3, \dots, n-2\}$, e.g., for $a = \frac{n}{2}$. So:

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

The probability that no a exists, such that all of e_1 's labels fall within $(a, n]$ and all of e_2 's labels fall within $(0, a]$, 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 $\{u_1, u_2\}, \{u_3, u_4\}, \dots, \{u_{n'-1}, u_{n'}\}$ defining the paths (*start, center, end*) $P_1 = (u_1, c, u_2), P_2 = (u_3, c, u_4), \dots, P_{n'} = (u_{n'-1}, c, u_{n'})$. These paths receive *independent labels since no edges of P_i overlap with any edge of P_j , $i, j = 1, 2, \dots, n', i \neq j$* . So:

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

i.e.,

$$\Pr[TD(N) \leq n] \leq e^{-\frac{n'}{2^{2k}}}$$

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

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

So:

$$\frac{n'}{2^{2k}} \geq \frac{n}{3} \left(\frac{1}{n^2}\right)^{\frac{1}{\beta(n)}} > \log n \quad (6)$$

Relation (6) 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 (6), we have:

$$\begin{aligned} -\frac{n'}{2^{2k}} < -\log n &\Rightarrow e^{-\frac{n'}{2^{2k}}} < e^{-\log n} = \frac{1}{n} \\ \Rightarrow \Pr[TD(N) \leq n] &\leq \frac{1}{n} \end{aligned}$$

□

3.3. An upper bound on ρ^*

Consider $N(k, n, \rho)$ on a connected undirected $G = (V, E)$, where $|V| = n$ and $k \geq \text{diam}(G)$. We introduce here a structure that will help us prove an upper bound on ρ^* . This structure is associated with the set of labels in the first period, namely $\{1, \dots, n\}$, and consists of a sequence of “boxes”; we partition the set $\{1, \dots, n\}$ into consecutive subsets and associate each one of them with a box in the structure.

Figure 3: Structure $s(e)$.

More specifically, for each edge e of G , consider a structure $s(e)$ being a sequence of boxes $B_1(e), B_2(e), \dots, B_{\text{diam}(G)}(e)$ (see Figure 3).

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

$$\forall i = 1, 2, \dots, \text{diam}(G),$$

$$\text{Box}_i \text{ corresponds to } L_i(e) = \{(i-1)\lambda + 1, \dots, i\lambda\}$$

Claim 1. *If for all $e \in E$ and for all $\text{Box}_i(e)$, at least one label of $L_i(e)$ gets into $\text{Box}_i(e)$, then $TD(N) \leq n$, i.e., any vertex can reach any other vertex via a journey within the first period.*

Proof. For any s, t , any shortest path p from s to t will be of length $|p| \leq \text{diam}(G)$. Any edge e may be at any “position” in p (first, second, \dots , last) or not belong to p at all. The journey from s to t is the path $p = (e_{p_1} = \{s, u_1\}, e_{p_2} = \{u_1, u_2\}, \dots, e_{p_{\text{last}}} = \{u_{|p|-1}, t\})$ which, for each edge e_{p_i} , uses a label that is in the box $\text{Box}_{p_i}(e_{p_i})$. \square

Note now that when we assign a random label to edge e for the first period, Π_1 , the probability that this label falls in $\text{Box}_i(e)$ is exactly $\frac{\lambda}{n}$. For ρ random labels assigned to e for the first period, Π_1 , the probability that none of them falls in $\text{Box}_i(e)$ is $(1 - \frac{\lambda}{n})^\rho$. Thus, the probability of the event:

$$A(e) = \text{“there exists a box of } e \text{ without a label in the first period”}$$

is at most $\text{diam}(G) \left(1 - \frac{\lambda}{n}\right)^\rho$.

Clearly,

$$\left(1 - \frac{\lambda}{n}\right)^\rho \leq e^{-\frac{\lambda\rho}{n}} = e^{-\frac{\rho}{\text{diam}(G)}}$$

and since $diam(G) \leq n$, it is enough to have $\frac{\rho}{diam(G)} \geq 2 \log n$ to get $diam(G) \left(1 - \frac{\lambda}{n}\right)^\rho \leq n \frac{1}{n^2} = \frac{1}{n}$. But,

$$\frac{\rho}{diam(G)} \geq 2 \log n \Leftrightarrow \rho \geq 2 diam(G) \log n$$

Indeed, if we assign at least $2 diam(G) \log n$ per edge per period, we get:

$$\begin{aligned} ETD(N) &\leq n \left(1 - \frac{1}{n}\right) + n^2 \frac{1}{n} \\ &\leq n - 1 + \frac{1}{n} \\ &\leq n \end{aligned}$$

Therefore, it must be $\rho^* \leq 2 diam(G) \log n$.

4. Conclusions and further research

In this work, we extend the research on temporal networks, by introducing and studying the model of random ephemeral temporal networks. A further goal of our research is to fully characterise the fast such networks, in which information dissemination can be expected to be very quick. We also aim in establishing a tighter upper bound on the critical availability, ρ^* . Yet another goal for our future research is to study models of random temporal networks, where the random selection of availability labels for the edges of the underlying graph follows distributions, other than the uniform. The subject of our current research also includes designing the link availability of networks (under constraints) so that low *cost* is achieved, where the cost is the total number of labels used; this research involves the study of random labellings to the edges of the underlying graph, but also the deterministic assignment of labels to achieve optimal cost.

5. Acknowledgments

This work was supported in part by:

- (i) the School of EEE/CS of the University of Liverpool
- (ii) the NeST initiative of the School of EEE/CS of the University of Liverpool

- (iii) the FET EU IP Project MULTIPLEX under contract No. 317532, and
 (iv) the EPSRC Grant EP/K022660/1.

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper.

6. Glossary of symbols

Below is a table of symbols, meant to assist the reader with identifying core notions used and referred to throughout the paper. Unless otherwise stated, the following symbols have the respective meaning in the main text. For more details about certain symbols, the reader is referred to the definition or explanation of the respective symbol within the main text.

Symbol	Meaning
G	- (di)graph
V	- set of vertices
E	- set of edges
n	- number of vertices
e	- an edge
L_e	- the label set of edge e
L	- assignment of labels to the edges of a (di)graph
$N(G)$	- temporal network on G
$N(k, \alpha, \rho)$	- ephemeral random temporal network
k	- number of periods of an ephemeral random temporal network
α	- length of period in an ephemeral random temporal network
ρ	- availability of labels per period in an ephemeral random temporal network
Π_i	- i^{th} period of an ephemeral random temporal network
$L(j, e)$	- the set of random labels of e in the j^{th} period of an ephemeral random temporal network

$\frac{\rho}{\alpha}$	- density of periodic random availabilities
$l_{max} = \alpha k$	- maximum lifetime of $N(k, \alpha, \rho)$
(u, v, l)	- time edge
$\delta(u, v)$	- temporal distance
TD	- temporal diameter
$diam(G)$	- diameter of G
$G_{n,p}, D_{n,p}$	- Erdős-Renyi (di)graphs
$\Gamma_i(s)$	- set of vertices in depth i from vertex s (more details on page 11)
$\Gamma'_i(t)$	- set of vertices in depth i from vertex t in reverse order (more details on page 11)
$B(i, j)$	- binomial distribution with parameters i and j
c_1, c_2	- constants used in Lemma 1 (more details on page 12)
ρ^*	- critical availability
$\frac{\rho^*}{n}$	- critical density of a normalized ephemeral random temporal network

7. References

- [1] E. C. Akrida, L. Gasieniec, G. B. Mertzios, and P. G. Spirakis. Ephemeral networks with random availability of links: diameter and connectivity. In *26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '14, Prague, Czech Republic - June 23 - 25, 2014*, pages 267–276, 2014.
- [2] 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.
- [3] 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.

- [4] 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.
- [5] B. Bollobás. *Modern Graph Theory*. Springer, 1998.
- [6] B. Bollobás. *Random Graphs, 2nd Edition*. Cambridge University Press, 2001.
- [7] 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.
- [8] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems (IJPEDES)*, 27(5):387–408, 2012.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.

- [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] S. Kar and J. M. F. Moura. Distributed average consensus in sensor networks with random link failures. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2007, Honolulu, Hawaii, USA, April 15-20, 2007*, pages 1013–1016, 2007.
- [18] S. Kar and J. M. F. Moura. Sensor networks with random links: Topology design for distributed consensus. *IEEE Transactions on Signal Processing*, 56(7-2):3315–3326, 2008.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] R. Koch, E. Nasrabadi, and M. Skutella. Continuous and discrete flows over time - A general model based on measure theory. *Math. Meth. of OR*, 73(3):301–337, 2011.
- [23] 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.

- [24] 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.
- [25] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Mediated population protocols. *Theoretical Computer Science*, 412(22):2434–2450, 2011.
- [26] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.
- [27] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Causality, influence, and computation in possibly disconnected synchronous dynamic networks. In *OPODIS*, pages 269–283, 2012.
- [28] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer, 2002.
- [29] 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.
- [30] B. Pittel. Linear probing: The probable largest search time grows logarithmically with the number of records. *J. Algorithms*, 8(2):236–249, 1987.
- [31] C. Scheideler. Models and techniques for communication in dynamic networks. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2285, pages 27–49, 2002.



Eleni C. Akrida was born and raised in Patras, Greece. In 2010, she received her degree in Mathematics with distinction from the University of Patras. In 2013, she received her MSc. in Mathematics of Computation and Decision Making from the Department of Mathematics and the Computer Engineering and Informatics Department in the University of Patras. Since October 2013, she has been a PhD student in the department of Computer Science in the University of Liverpool, UK. Her work is supervised by Prof Paul G. Spirakis and Prof Leszek Gąsieniec. Her research interests are in the field of algorithms, with a focus on distributed computing and dynamic networks, and include, amongst others, probabilistic methods in algorithms, algorithms and complexity of graph theoretic problems and optimization problems in dynamic networks.



Leszek Gasieniec received MSc (1991) and PhD (1994) degrees in Computer Science from Warsaw University in Poland. He was a postdoctoral researcher in Quebec University ('94-'95) in Canada, and in Max-Planck Institute für Informatik ('95-'97) in Germany. Since 1997 he has been with University of Liverpool where he is now a full professor and the Head of Department of Computer Science. The main strand of his research explorations include parallel and distributed algorithms, network communication, distributed search problems as well as string and graph algorithmics with application to Bioinformatics.



George B. Mertzios received his Diplom (with Distinction) in Mathematics and Informatics from the Technische Universitaet Muenchen, Germany (2005) and his Ph.D. in Computer Science from the RWTH Aachen University, Germany (2009). From 2010 to 2011 he worked as a Postdoctorate Researcher at the Department of Computer Science at Technion - Israel Institute of Technology, Haifa, Israel, and in the Caesarea Rothschild Institute for Computer Science, University of Haifa, Israel. Since September 2011 he is a Lecturer in Computer Science at Durham University, UK. His research areas include efficient graph algorithms, computational complexity, parameterized complexity, and evolutionary graph theory.



Prof Paul G. Spirakis is Chair in Computer Science at the University of Liverpool, President of the Computer Technology Institute in Greece, and the leader of Networks Science and Technology (NeST) initiative. His research interests include algorithms, complexity and game theory. Prof Spirakis has contributed to more than 375 scientific publications, including papers on network formation protocols, evolutionary processes on graphs, and more recently temporal graphs. Including current activities he participated as the team leader in 19 European Union IST Projects, and he led 10 Greek Research Projects. Together with his colleagues, Prof Spirakis founded and further popularised several core Computer Science conferences including SPAA, DISC, WINE, ESA and SAGT. He also chaired other such as ICALP, STOC, and IPDPS. Prof. Spirakis is a member of Academia Europe, a member of the ACM Europe Council, the Vice-President of the European Association of Theoretical Computer Science (EATCS) as well as EATCS Fellow. He has been acknowledged as amongst the top 50 scientists worldwide in Computer Science with respect to The Best Nurturers in Computer Science Research (published by B. Kumar and Y.N. Srikant, SIAM Data Mining, 2005). He has also supervised more than 27 PhD students.