

Evaluating Temporal Robustness of Mobile Networks

Salvatore Scellato*, Ilias Leontiadis*, Cecilia Mascolo*, Prithwish Basu†, Murtaza Zafer‡

*Computer Laboratory, University of Cambridge, UK

†Raytheon BBN Technologies, MA, USA

‡IBM T. J. Watson Research, NY, USA

Abstract—The application of complex network models to communication systems has led to several important results: nonetheless, previous research has often neglected to take into account their temporal properties, which in many real scenarios play a pivotal role. At the same time, network robustness has come extensively under scrutiny. Understanding whether networked systems can undergo structural damage and yet perform efficiently is crucial to both their protection against failures and to the design of new applications. In spite of this, it is still unclear what type of resilience we may expect in a network which continuously changes over time. In this work we present the first attempt to define the concept of temporal network robustness: we describe a measure of network robustness for time-varying networks and we show how it performs on different classes of random models by means of analytical and numerical evaluation. Finally, we report a case study on a real-world scenario, an opportunistic vehicular system of about 500 taxicabs, highlighting the importance of time in the evaluation of robustness. Particularly, we show how static approximation can wrongly indicate high robustness of fragile networks when adopted in mobile time-varying networks, while a temporal approach captures more accurately the system performance.

Index Terms—Mobile networks, temporal networks, network robustness.



1 INTRODUCTION

The study of real-world communication systems by means of complex network models has provided insightful results and has vastly expanded our knowledge on how single entities create connections and how these connections are used for communication or, more generally, interaction [1]. Statistical network models such as the Erdős Rényi random graph [2] and the scale-free networks generated by preferential-attachment mechanisms [3] have been largely exploited to study and understand real-world systems. While the former model embodies networks whose nodes are statistically indistinguishable, resulting in fairly homogenous properties, scale-free networks exhibit heterogeneous node properties, with the vast majority of nodes having a few connections and a small minority of hubs having a disproportionately large number of links. The different properties present in these two families of models are able to capture a vast range of real-world systems. In particular, technological networks such as the Internet and the World Wide Web have been under scrutiny in terms of structure and dynamic behavior [4], [5]. More recently, with the widespread adoption of mobile and opportunistic networks, it has become important to develop new analytical tools to keep into account network dynamics over time [6], [7] and how this affects phenomena such as information propagation [8], [9].

Results have shown that time correlations and relative temporal ordering of connection events among nodes cannot be neglected, otherwise the performance of a given system can be greatly overestimated [10], [11].

At the same time, the problem of understanding whether real systems can sustain substantial damage and still maintain acceptable performance has been extensively addressed [12], [13]. Various measures of network robustness have been defined and investigated for several classes of networks, evaluating how different system can be more or less resilient against random errors or targeted attacks thanks to their underlying structural properties [14], [15].

Nonetheless, it is still unclear how to approach the study of robustness of networks by taking into account their time-varying nature: by adopting a static representation of a temporal network, important features which impact the actual performance might be missed. Thus, it becomes important to develop a robustness metric which takes into account the temporal dimension and gives insights on how a mobile network is affected by damage or change. Particularly, the fact that links are not always active means that information spreading can be delayed or even stopped and that *relative ordering in time of connection events may affect the creation of temporal paths in mobile networks*.

Our main goal is to design a novel framework for the analysis of robustness in mobile time-varying networks. We adopt *temporal network metrics* [10] to quantify network performance and define a measure of robustness against generic network damages. At first, we study its

An earlier initial version of this work appeared in Proceedings of IEEE Infocom 2011, Miniconference track, with the title "Understanding Robustness of Mobile Networks through Temporal Network Measures".

performance on random network models to understand its properties; then we apply our method to study a real mobile network, describing how temporal robustness gives a more accurate evaluation of system resilience than static approaches.

Our contributions can be summarized as follows:

- We describe the use of temporal network metrics such as temporal distance to estimate the current network connectivity taking temporal variability into account. We define *temporal network robustness* (Section 2), a novel measure which quantifies how the communication of a given time-varying network is affected by damage.
- We provide an analytical computation of temporal metrics on a temporal version of the Erdős Rényi (ER) random graph model (Section 3) and we evaluate through numerical simulation both a Markov-based link connectivity model, which provides time-correlations, and two random mobility models, which introduce space-correlations (Section 4).
- We show how, *unlike what has been demonstrated with static measures, temporal networks do not exhibit sharp breakdowns but instead fail gracefully when they are subjected to failures*. The temporal dimension is able to capture the evolution dynamics, exposing the fact that time allows to create temporal paths across otherwise disconnected portions of the network.
- We investigate the usefulness of our approach to characterize a real communication system, an opportunistic vehicular network simulated with real mobility traces of about 500 taxis in the San Francisco Area (Section 5): we show that *the robustness evaluation based on static network analysis is too optimistic in real scenarios*, since it greatly overestimates how resilient a mobile system really is, without giving reliable measures.

We discuss some implications of our findings for the design of new systems and applications (Section 6). Finally, we review related results on this topic (Section 8) and we conclude the paper (Section 9).

2 TEMPORAL ROBUSTNESS

In this section we will review some basic metrics for temporal networks and describe how these measures can be adopted to quantitatively define temporal network robustness.

2.1 Network Robustness

The study of robustness of complex networks has mainly focused on describing how a given performance metric of the network is affected when nodes are removed according to a certain rule. The underlying assumption is that the absence or malfunctioning of some nodes will cause the removal of their edges and, thus, some paths will become longer, increasing the distances between the remaining nodes, or completely disappear, resulting in

the loss of connectivity in the whole system. The performance measures previously adopted include the network diameter [12], the size of the giant component [12], [14] and the average inverse geodesic length [14], [15]. Moreover, the strategies used to choose which nodes are to be removed can be divided in two broad categories: random failures, where every node has the same probability of being removed, and targeted attacks, where nodes are ranked according to a performance metric and then accordingly removed [14].

In this work we will study the problem of defining and analyzing robustness in evolving networks: as a consequence, *we need to use a performance metric which includes the temporal dimension in its definition*. At the same time, we focus on the first strategy of node removal: we consider random and independent failures for every node and we evaluate how the system tolerates increasing level of malfunctioning nodes.

2.2 Temporal Network Metrics

In networks where nodes can connect with a large number of other nodes over time (e.g., mobile and peer to peer networks) time ordering of events is important. Traditional approaches that aggregate links over time always overestimate network connectivity, which means that some paths that seem to be present are in reality not present due to the actual ordering of the links over time. Hence, in this section we describe the notion of *temporal graph* we use to model mobile networks [10], [11].

Temporal graph: A temporal graph is a graph $G(t) = (V(t), E(t))$ where $V(t)$ is the set of nodes at time t and $E(t)$ is the set of edges at time t . We assume that $|V(t)| = N$, thus nodes cannot be added from the graph, while nodes can be removed by disconnecting them from the rest of the network. Furthermore, we treat time as a discrete entity and we create a sequence of graphs $G(t_1), \dots, G(t_2)$ by adopting an appropriate resolution time.

Temporal distance: Then, given two nodes i and j we can define a temporal path $p_{ij}(t_1, t_2)$ between them in the time window $[t_1, t_2]$. The length of a temporal path is defined as the amount of time steps it takes to spread information from node i to node j on that particular path. This value is always a positive integer. As a consequence, we can define the *shortest temporal distance* $d_{ij}(t_1, t_2)$ as the smallest length among all the temporal paths between nodes i and j in time window $[t_1, t_2]$. For example, if a message sent by node A is received by node B at time t_d , then $d_{AB} = t_d - t_1$. If there is no temporal path between i and j in $[t_1, t_2]$, their distance can be considered infinite, i.e. $d_{ij}(t_1, t_2) = \infty$. The *average temporal distance* $L(t_1, t_2)$ of a given temporal graph G during a time window $[t_1, t_2]$ is defined as:

$$L(t_1, t_2) = \frac{1}{N(N-1)} \sum_{i,j} d_{ij}(t_1, t_2) \quad (1)$$

This quantity is not well defined when some pairs of nodes are not temporally connected. As a consequence, another metric has been introduced.

Temporal efficiency: We can define the *temporal global efficiency* $E(t_1, t_2)$ of a given temporal graph G

$$E(t_1, t_2) = \frac{1}{N(N-1)} \sum_{i,j} \frac{1}{d_{ij}(t_1, t_2)} \quad (2)$$

This value is not affected by disconnected pairs of nodes, because their contribution to the efficiency is computed as zero. Network efficiency is normalized between 0 and 1 and it does not depend on the size of network, hence, it can be adopted to compare graphs with different sizes. The concept of network efficiency was originally introduced for static graphs explicitly to cope with disconnected couples and it measures how well nodes in can communicate [16].

Since a temporal graph is continuously evolving, we can evaluate how temporal efficiency changes over time by considering a value τ and evaluating $E_G(t)$ as the relative temporal efficiency of the temporal graph in the time window $[t - \tau, t]$. The effect of τ is to effectively impose an upper bound on the temporal distances, as all paths longer than τ simply do not exist. As a consequence, τ should be chosen so that any communication whose delay is longer than τ itself can be ignored. Similarly, $L_G(t)$ is computed as the average temporal distance in $[t - \tau, t]$. If the properties of the temporal graph are stationary, we expect these values to reach a steady value as time progresses. For instance, a temporal graph whose links are evolving according to an ergodic random process will exhibit stationary behavior, while a temporal graph whose evolution is ruled by a time-dependent process might not reach stationarity at all. In practice, a temporal graph with stationary properties might reach an equilibrium point where temporal efficiency is fluctuating around its average value, while this might not be valid for non-stationary temporal graphs.

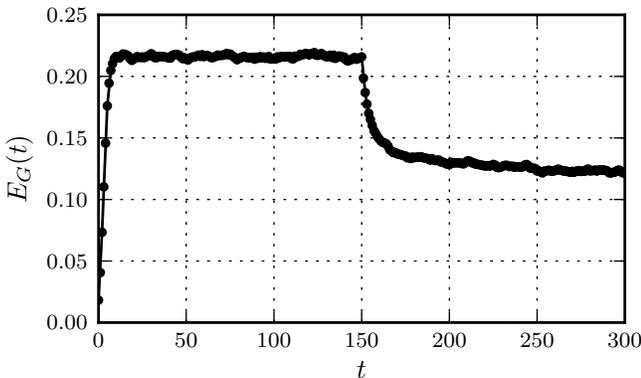


Fig. 1. Temporal efficiency in a time-varying network: at $t = 150$ 20% of the nodes are removed. The relative drop in efficiency after the damage can be used to quantify the robustness of the network.

2.3 Temporal Robustness Metric

Given a temporal graph G , we define a damage D as any structural modification on it and we define G_D as the graph resulting by the effect of the damage D on G . A damage may be the deactivation of some nodes or the removal of some edges at a particular time t_D . Because of damage D , some temporal shortest paths will be longer or will not exist any more, thus, we expect that the temporal efficiency will eventually reach a new steady value $E_{G_D} \leq E_G$. It is important to evaluate the new value of the temporal efficiency on a new temporal graph which still contains the deactivated nodes, in order to obtain a decrease in efficiency. Otherwise, we might obtain a smaller temporal graph which is more efficient than the original graph, although it has lost much of its structure. Hence, we do not consider highly dynamic systems where nodes can be constantly added or removed. Instead, we focus on evaluating the service degradation in a more controlled environment where only a number of existing nodes could fail.

We define the loss in efficiency $\Delta E(G, D)$ caused by the damage D on the temporal graph G as $\Delta E(G, D) = E_G - E_{G_D}$. Finally, we define the **temporal robustness** $R_G(D)$ of the temporal graph against the damage D as

$$R_G(D) = 1 - \frac{\Delta E(G, D)}{E_G} = \frac{E_{G_D}}{E_G} \quad (3)$$

This value is normalized between 0 and 1 and it measures the relative loss of efficiency caused by the damage: if the damage does not impact the efficiency of the graph ($E_{G_D} = E_G$) then its robustness is 1, while if the damage destroys the efficiency of the graph ($E_{G_D} = 0$) the robustness drops to 0. Temporal efficiency is a particularly suitable metric to study temporal network robustness as it denotes both longer temporal paths and the lack of paths among temporally disconnected nodes at the same time. Nonetheless, other metrics have been used to assess robustness in static systems: for instance, there could be scenarios where fast communication with small delays can be more important than global connectivity, thus other measures can be adopted. Provided that these measures can be extended to the temporal case, they can be easily integrated in our framework.

For example, as we see in Figure 1 for a random temporal network (see Section 3 for more details), a damage which affects a fraction of the nodes in a temporal graph will result in a drop of the temporal efficiency. By measuring the amount of lost efficiency we can quantify the robustness of the network to that particular damage. Moreover, by normalizing the efficiency drop with respect to the original efficiency of the network, we can compare the robustness of systems with different properties and sizes.

3 THEORETICAL ANALYSIS

In this section we analytically investigate how temporal metrics and temporal robustness evaluate on a random

model of a temporal network, namely a generalization of the Erdős Rényi random graph [2].

3.1 Temporal Random Network Model

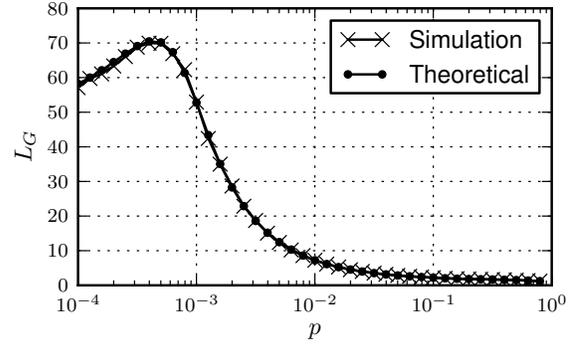
An Erdős Rényi (ER) random graph with N nodes and parameter p is created by independently including each possible edge in the graph with probability p , thus resulting in an expected number of edges $K = \frac{pN(N-1)}{2}$, and it is denoted as $G(N, p)$ [2]. This model has been successfully adopted to represent a simple random network, where only the density of links is fixed and all the rest is randomized. We generalize this model to the temporal case by creating a sequence of T ER random graphs $G(N, p)$ and we denote the resulting temporal graph as $G(N, p, T)$.

The simplicity of this model allows a theoretical analysis and the computation of its temporal metrics as a function of its parameters N , p and T . More precisely, we are interested in this question: if the flooding of a message initiates from a random node at time $t = 0$, at which time step will other nodes receive it? This problem has already been addressed by researchers studying epidemic dissemination in mobile systems [17]: nonetheless, their temporal analysis of message spreading usually exploits either mean-field approximation to study their behavior in the limit of large systems [18]. Some attempts have also focused on providing tight bounds on the speed of information dissemination on Markovian time-varying graphs [19]. These solutions provide good estimation of steady state behavior of the system, but can be misleading when they are used to analyze the transient behavior of the first steps of the dissemination. Instead, we are interested in the *exact* solution of this model in the case of a finite system with N nodes and with discrete treatment of time, since we need to focus on the exact temporal values of message delay also for the transient phase. As a consequence, we cannot use continuous mean-field formulation and we instead solve the problem by modeling it as a discrete state random process.

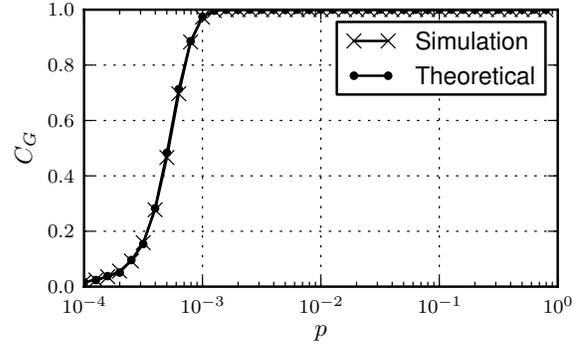
Let N_t be the random discrete process that counts the number of nodes which have received the message at time t : our aim is to compute the probability distribution of this random process for any given time $0 \leq t \leq T$. In the following we use $P(E)$ to denote the probability that event E occurs. In addition, we will treat the case of an arbitrary given node, without specifying any index, as in our case different nodes are statistically indistinguishable. We will make use of the following lemma to compute the probability distribution of N_t for $0 \leq t \leq T$:

Lemma 1. *The probability p_m that a single node without the message receives it in a time step when m other nodes already have it is given by $p_m = 1 - (1 - p)^m$.*

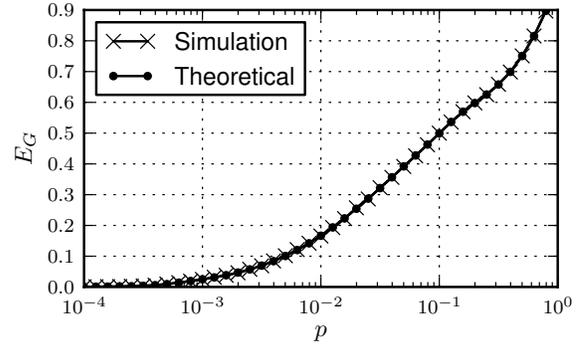
Proof: A node without the message will receive it in a given time step if at least one of the other m nodes with the message activates a link with it. Since every link



(a) Average temporal length



(b) Temporal connected couples



(c) Temporal efficiency

Fig. 2. Average temporal path length L_G (a), fraction of connected couples of nodes C_G (b) and temporal efficiency E_G (c) as a function of link probability P_{ON} for a temporal ER graph with $N = 100$ nodes with $T = 100$. Theoretical values and simulation results averaged over 100 runs are shown.

is activated with probability p , the node will not receive the message only if all m links will not be activated, which happens with probability $(1 - p)^m$. Thus, $p_m = 1 - (1 - p)^m$. \square

The following theorem then holds:

Theorem 1. *The probability distribution of the random pro-*

cess N_t is:

$$P(N_{t+1} = k) = \sum_{m=1}^k \binom{N-m}{k-m} p_m^{k-m} (1-p_m)^{N-k} P(N_t = m) \quad (4)$$

with the initial condition $P(N_0 = 1) = 1$.

Proof: We have an initial condition with only 1 node with the message at time $t = 0$, the sender, so $P(N_0 = 1) = 1$. In addition, the number of nodes reached by the message can only increase with time. Then, the probability distribution of N_{t+1} is defined as a function of the probability distribution of N_t :

$$\begin{aligned} P(N_{t+1} = k) &= \sum_{m=1}^k P(N_{t+1} = k | N_t = m) P(N_t = m) \\ &= \sum_{m=1}^k P(M_{k-m} | N_t = m) P(N_t = m) \end{aligned}$$

where M_{k-m} is the event that $k - m$ new nodes are reached by the message when there are already m nodes with the message in the system. The probability that this event happens does not depend on time t but only on m . When there are $N - m$ nodes that can be reached by the message, each of them will get the message with probability p_m . Then, Eq.(5) can be expressed in terms of a binomial distribution with parameter p_m for every value of m , giving Eq. (4).

$$\begin{aligned} P(N_{t+1} = k) &= \sum_{m=1}^k P(M_{k-m} | N_t = m) P(N_t = m) = \\ P(N_{t+1} = k) &= \sum_{m=1}^k \binom{N-m}{k-m} p_m^{k-m} (1-p_m)^{N-k} P(N_t = m) \end{aligned}$$

□

As a consequence, we have the following result on the probability of a certain message delay:

Corollary 1. *The probability R_t that a node has received the message after t steps is given by*

$$R_t = \sum_{k=1}^N \frac{k-1}{N-1} P(N_t = k) \quad (5)$$

Proof: The result comes from the fact that when $N_t = k$ any node different than the source, and chosen at random, has the message with uniform probability $\frac{k-1}{N-1}$. □

Since $P(N_0 = 1) = 1$, then $R_0 = 0$. Instead, as t increases the probability of receiving the message increases as well. In particular, if l is the message delay of a given node, then $R_t = P(l \leq t)$,

Corollary 2. *The probability that a node is reached by the message exactly at time t is*

$$d_t = P(l = t) = P(l \leq t) - P(l \leq t-1) = R_t - R_{t-1} \quad (6)$$

Proof: If a node is reached by time t then either it is reached exactly at time t or it is was already reached by time $t-1$, giving $R_t = d_t + R_{t-1}$. □

This derivation gives us the exact probability distribution d_t of temporal distances in $G(N, p, T)$, which enables us to compute temporal metrics. The final fraction C_G of nodes reached by a temporal path before $t = T$ is given by R_t , whereas the expected average temporal length L_G and temporal efficiency E_G can be computed using the probability distribution d_t :

$$C_G = R_T \quad (7)$$

$$L_G = \frac{1}{C_G} \sum_{t=1}^T t d_t \quad (8)$$

$$E_G = \sum_{t=1}^T \frac{1}{t} d_t \quad (9)$$

3.2 Temporal Metrics

In order to validate our results, we simulate this model for various values of p and we compute temporal efficiency $E_G(t)$ with $N = 100$ nodes for 500 time steps with a time window of $\tau = 100$, computing the average efficiency value over the last τ steps. We also evaluate the average temporal length L_G and the fraction of connected couples of nodes C_G . We then compare these results with the analytical solution of $G(N, p, \tau)$ for the same values of p . As we see in Figure 2, theoretical derivation and numerical simulation are in perfect agreement. We also note that as the probability p increases both temporal efficiency and the number of connected couples increases. Instead, the average temporal distance L_G reaches a maximum value and then decreases as we increase p (Figure 2(a)). This is due to the fact that for low values of p not all the couples of nodes can be connected by a temporal path in the sequence of T graphs: thus, by increasing p , more and more couples can be reached but with longer paths. As soon as total connectivity is reached (Figure 2(b)), so that there exists a temporal path between any pair of nodes, larger values of p give a decreasing temporal distance L_G . This demonstrates how temporal efficiency E_G is a better measure than temporal distance: since it is not affected by partial temporal connectivity it gives a clear indication of the temporal performance of the network (Figure 2(c)).

Furthermore, we note that there is no evidence of sharp transition from a disconnected to a connected temporal graph, since as we increase p the fraction of connected couples C_G smoothly increases from 0 to 1. This effect is due to the temporal dimension: in this model, no matter how small p is, the temporal graph will be eventually connected as $T \rightarrow \infty$ as long as there is a non-zero probability that each link is present. On the other hand, in the static case a ER random graph will experience a

1. This is computed only on the connected couples, so C_G is the conditional probability that a node is reached by a temporal path.

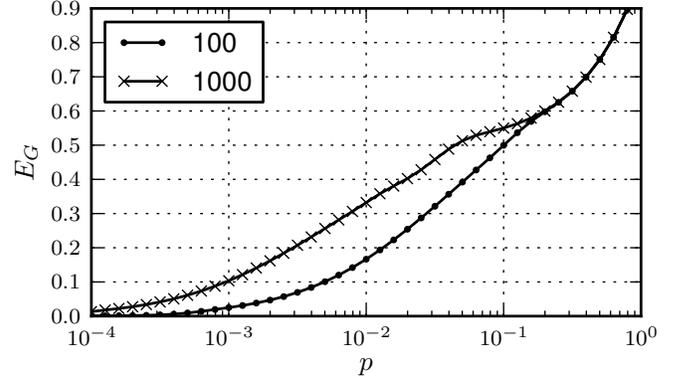
sharp transition and will be connected, on average, only when p approaches $\frac{\ln N}{N}$ [2].

Finally, we investigate the impact of parameters N and T on the temporal metrics by analyzing temporal efficiency E_G . Figure 3(a) reports the relation between E_G and probability p for two models with size $N = 100$ and $N = 1000$. A larger number of nodes results in higher efficiency for the same value of p , since the number of potential links in the model increases as N^2 and thus there are many more potential paths to be used to connect nodes. Instead, we see in Figure 3(b) that the length of the interval T is not affecting the results: as long as T is large enough to allow the creation of the temporal paths, increasing it will not make any difference. The influence of T on temporal efficiency decreases as T increases: all potential paths that are longer than T are effectively removed, but longer paths contribute less and less to temporal efficiency. As a result, when temporal efficiency is evaluated empirically by using a sliding window τ a value large enough to let temporal paths be created needs to be used. Larger values will not make any difference and, instead, they might slow down the computation.

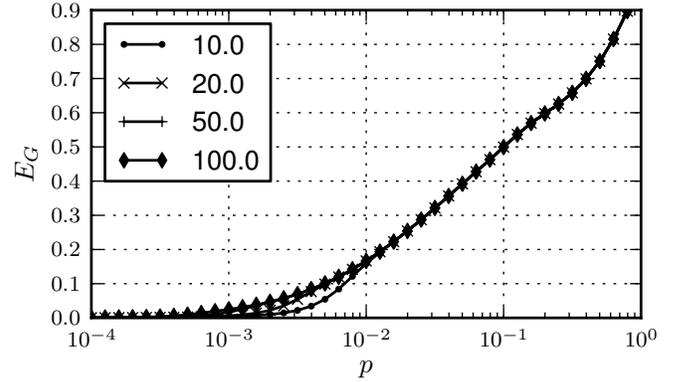
3.3 Temporal Robustness

We evaluate the temporal robustness of the temporal ER model by directly computing the temporal efficiency for a normal network and for a damaged network. We model a damage as a random failure of the nodes in the system by independently deactivating with probability P_{ERR} each node. As a result, an expected number of $N_D = NP_{ERR}$ nodes will be deactivated, thus not participating in any more communication. The damage can be then entirely quantified by the probability of error P_{ERR} : in this work we only consider independent and identically distributed error probabilities, even though our framework can be extended to address also more complex cases with correlated errors. Thus, we investigate this model to understand how the temporal robustness measure behaves across systems with different size and different efficiency.

As reported in Figure 4, temporal robustness shows that this random model fails smoothly as we increase the fraction of removed nodes, without any sudden disruption for any value of P_{ERR} . This is a main difference with respect to what happens in the static case: for a static ER random graph be a critical value of P_{ERR} which causes a breakdown of the network in several disconnected components may exist [12]. This is not true for temporal robustness, as new paths can still appear after the damage as the network rearranges its connections. Time provides more redundancy and, hence, more resilience. Moreover, we also note that temporal robustness does not depend on system size: since it is normalized with respect to the value of temporal efficiency before the damage, it depends only on the relative drop in efficiency, not on the absolute values of the metric.



(a) Effect of N



(b) Effect of τ

Fig. 3. Temporal efficiency E_G with $\tau = 100$ and for two systems with $N = 100$ and $N = 1000$ nodes (Figure 3(a)): larger systems result in higher efficiency values for the same p , except for higher values of p . Temporal efficiency E_G as a function of probability of link presence p for different values of τ for a system with $N = 100$ (Figure 3(b)): different values of τ result in similar values of temporal efficiency.

4 NUMERICAL EVALUATION

After the theoretical investigation of temporal robustness on a simplistic random model, we proceed to the evaluation of this metric on more complex models which introduce correlations in the temporal graph, with the aim of investigating how these correlations impact temporal network robustness. In this section we present a numerical analysis of temporal metrics and robustness for two different classes of random temporal networks, evaluating their resilience against random errors.

4.1 Temporal Network Models

4.1.1 Markov-based Temporal Network Model

The temporal ER network model does not provide temporal correlations between consecutive graphs in the sequence. We now consider a model where link evolution is described by a Markov process, thus enabling memory effects in network dynamics [20].

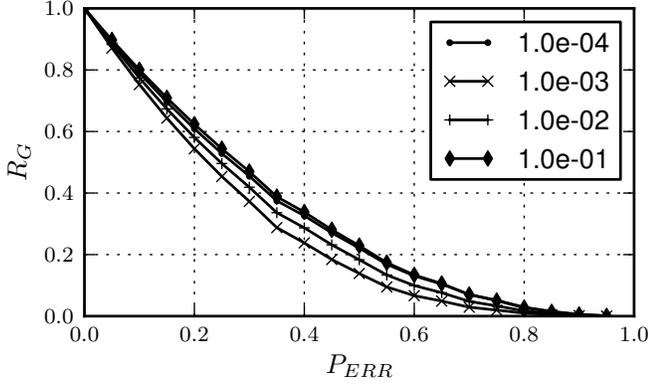
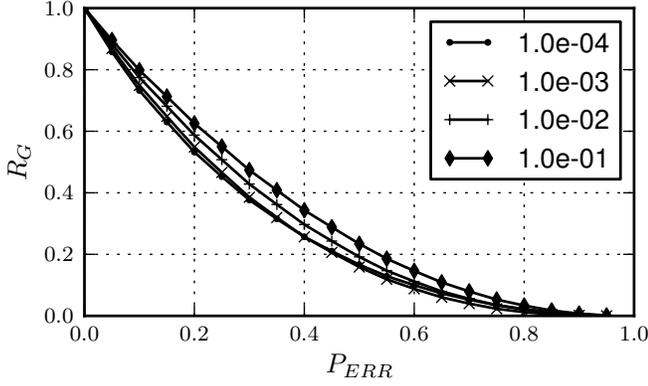
(a) $N = 100$ (b) $N = 1000$

Fig. 4. Temporal robustness R_G as a function of probability of error P_{ERR} in the ER random model for different link probability p . The size of the system has no impact on temporal robustness: furthermore, the system fails smoothly as the probability of error increases.

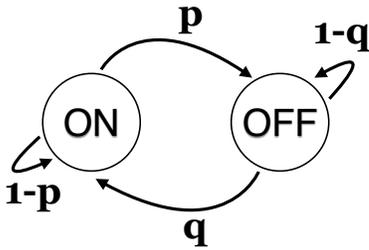


Fig. 5. Markov model of temporal link dynamics.

We consider a complete graph G with N nodes. At every discrete timestep t each link may or may not be present: we derive a temporal graph where the existence of each link evolves according to a 2-state discrete Markov process.

Each link is either ON or OFF and it transitions between these states at each time step according to constant probabilities: ON \rightarrow OFF with probability p and OFF \rightarrow ON with probability q , as sketched in Figure 5. Therefore

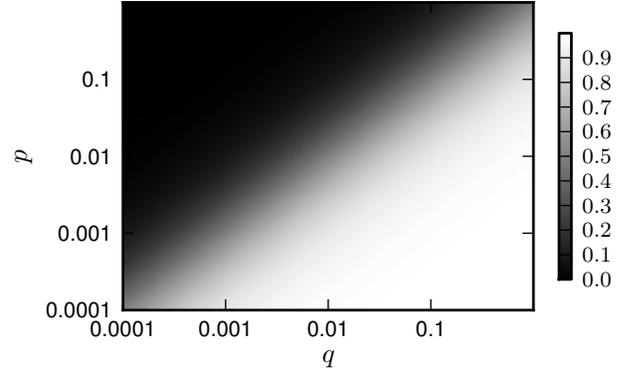
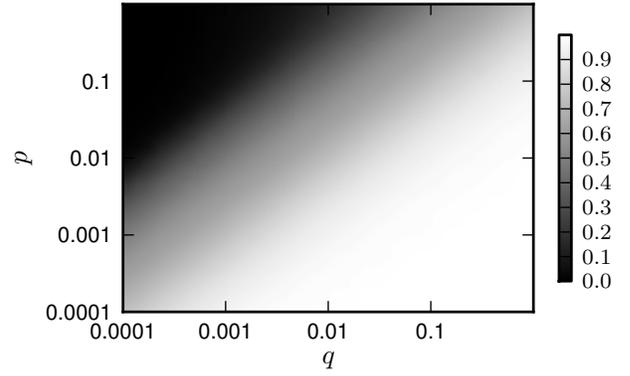
(a) P_{ON} (b) E_G

Fig. 6. Probability of link presence P_{ON} (a) and temporal efficiency E_G (b) as a function of parameters p and q in the Markov-based temporal network model. The two quantities exhibit similar trends in the parameter space.

the model exhibits the classic Markovian property, where the state at time t is only dependent upon the state at $t - 1$. As with any Markov process, each link will approach a steady-state probability distribution over sufficient time. The steady probability of link presence is $P_{ON} = \frac{q}{p+q}$: as a consequence, each observation of the temporal graph appears as an ER random graph with each edge present with probability P_{ON} . However, there will be temporal correlations between consecutive graphs because of the Markov dynamics and for a given value of P_{ON} there will be different combinations of p and q which give that value, albeit with different dynamical properties. In particular, the average number of consecutive steps a link persists in state ON is $D_{ON} = \frac{1}{p}$, while the average number of consecutive steps spent in state OFF is $D_{OFF} = \frac{1}{q}$. These two values can be also considered as the average contact and inter-contact time of the temporal network and they are geometrically distributed.

4.1.2 Mobility-based Temporal Network Models

We can also create a random model of a temporal network by using mobility models. In this case we are

introducing not only memory effects but also topological constraints: indeed, a key difference with the previous temporal models is that each node is not equally likely to connect with all the other nodes, due to the effect of spatial distance.

We consider $N = 100$ nodes moving in a square area 1000×1000 meters and we define a communication range r : at every time step we create a random geometric graph where two nodes are connected with a link if their Euclidean distance is shorter than r . Thus, we can change the probability of link presence P_{ON} by varying either the communication range or the density of nodes and in our simulations we only vary the former. Then, a temporal graph can be defined as the sequence of graphs extracted at each time step while the nodes move.

We investigate two different mobility models that are implemented using the Universal Mobility Model Framework [21]: Random Waypoint Model (RWP) and Random Waypoint Group Model (RWPG): both these models provide a simple yet inherently randomized temporal network that can be directly compared with the previous models.

In RWP each node selects uniformly at random a location and it moves towards this location with speed uniformly distributed in a fixed range $[5, 40]$ mph. As soon as the node reaches its destination, it waits for a randomly distributed time in $[0, 120]$ seconds and repeats the above steps until the end of the simulation. We chose these values as they resemble the real traces that we will study in the following Section 5. RWP is a very simplistic random model that may not capture the complexity of real mobility, however it provides a certain degree of homogeneous spatial mixing among nodes.

In RWPG nodes are divided into two classes: there are M group leaders and $N - M$ group followers. Every group follower has its own leader so that the N nodes are divided into equally-sized groups. Each group leader selects a random target and moves towards it, similarly to the RWP mobility model. Group members do not select any target; instead they just follow their group leader according to the *pursuit force* [21] which is set to give a group span of 200 meters: this force acts as an attraction force by predicting the geographical point where the group members (pursuers) can catch the evader (group leader). Once that location is determined a steering force is applied on the pursuers to randomize their path towards the leader. The attraction force is influenced by the *group span* parameter that in our case is set to be 200 meters, e.g., group members will try to stay within 200 meters from the group leader.

4.2 Simulation Strategy

We numerically evaluate temporal efficiency $E_G(t)$ over time, adopting a time window of $\tau = 100$, for a graph with $N = 100$ nodes: after an initial phase, the random temporal graph reaches an equilibrium state and we compute the steady value of temporal efficiency. We run

each simulation for 2τ steps and we compute the average value of temporal efficiency over the last τ steps. All results have been averaged over 100 different runs.

We also evaluated numerically temporal robustness by adopting the same failure strategy as in the previous case, removing each node independently with probability P_{ERR} . We measure temporal efficiency before and after the failure, when the network reaches a new equilibrium state.

In order to quantify temporal robustness, we compute temporal efficiency $E_G(t)$ over time, adopting a time window of $\tau = 100$: after an initial phase, the temporal graph will reach an equilibrium and we can compute the steady value of temporal efficiency before the damage E_G as the average temporal efficiency between $t = 100$ and $t = 150$. At $t_{ERR} = 150$ we randomly remove nodes with probability P_{ERR} and we wait until the network reaches a new steady state. Then, we compute the average value of temporal efficiency after the damage $E_{G_{ERR}}$ as the average value between $t = 250$ and $t = 300$. Hence, a simulation runs for 300 time steps.

4.3 Temporal Metrics

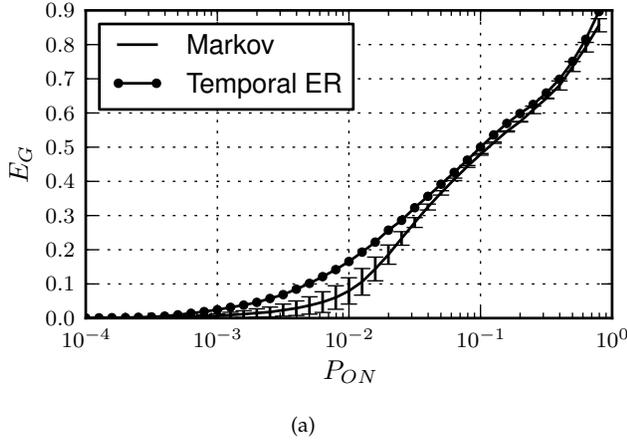
4.3.1 Markov-based model

In Figure 6(a) we report the probability of link presence P_{ON} as a function of the two parameters of the Markov process p and q : we see how the parameter space appears divided in two regions according to which parameter is larger than the other. As we see in Figure 6(b), temporal efficiency shows a similar behavior as P_{ON} in the parameter space. This is an indication that in this model the most important parameters is the probability of link presence, while other parameters such as D_{ON} or D_{OFF} are less important.

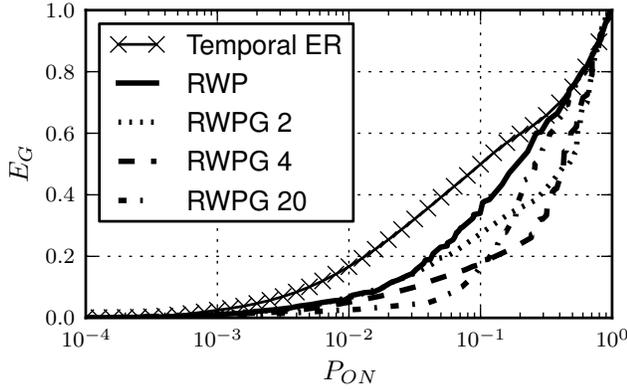
However, this intuition is only partially confirmed in Figure 7(a), where temporal efficiency is a function of the probability of link presence both for the temporal ER model and for the Markov model. Similar values of P_{ON} results in similar values of efficiency, regardless the actual values of p and q . Yet, the same value of P_{ON} results, on average, in higher efficiency in the temporal ER case, since avoiding time-correlations allows the creation of new edges at higher rate: thus, given an equal time interval, single nodes enjoy more opportunities to communicate directly with new nodes in the uncorrelated case than in the correlated model, where they might maintain the same set of links for a longer period due to memory effects. Instead, for higher values of P_{ON} the two models behave in a similar way as they reach almost complete connectivity: this is because with high P_{ON} at each time step every node is connected by a direct link to a large fraction of the other nodes and so efficiency is already close to the maximum value regardless how fast edge dynamics is.

4.3.2 Mobility-based models

Figure 7(b) depicts temporal efficiency E_G as a function of probability of link presence P_{ON} in the RWP



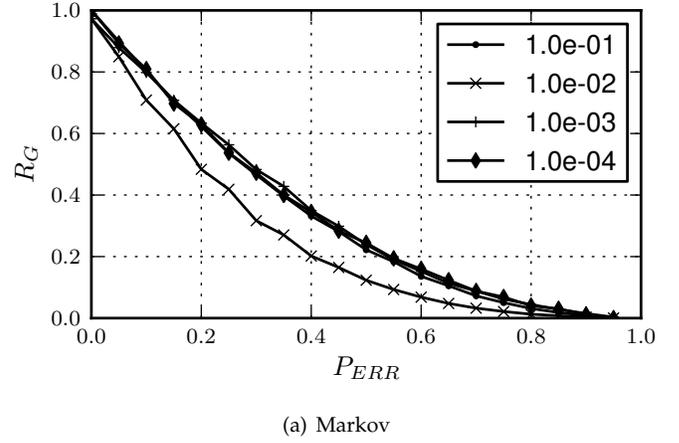
(a)



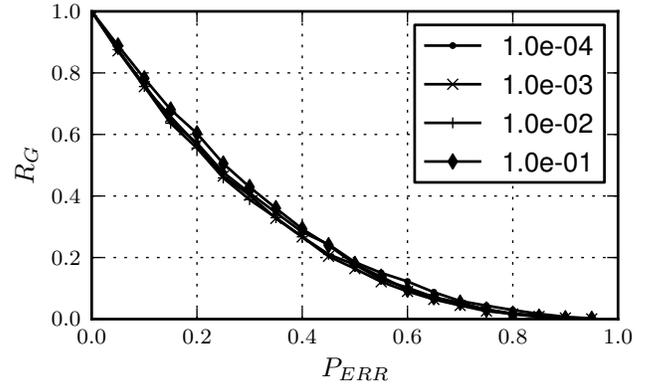
(b)

Fig. 7. Temporal efficiency E_G as a function of probability of link presence P_{ON} in the Markov model and compared to the temporal ER model (Figure 7(a)). The Markov model has error bars which show standard deviation of E_G for different parameter combinations which hold approximately the same P_{ON} (logarithmic binning has been adopted). The two models behave similarly, with some differences only for lower values of efficiency. Temporal efficiency E_G as a function of probability of link presence P_{ON} for different mobility-based network models (Figure 7(b)): Random Waypoint Model (RWP) and Random Waypoint Group Model (RWPG) with different number of groups. Temporal ER model is shown for comparison. The introduction of groups decreases the overall temporal efficiency of the system.

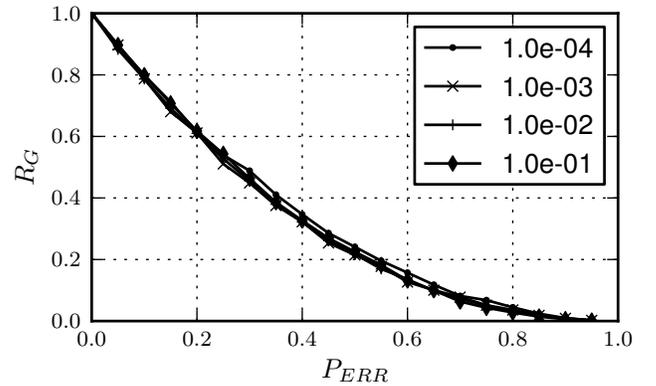
case: we vary the communication range to obtain an expected given value of P_{ON} . There is a trend similar to the previous models: however, for the same value of P_{ON} , the resulting efficiency is always smaller in the RWP case than in the temporal ER case. This can be attributed to the fact that nodes move in a geographically restricted manner and, thus, they do not exhibit the same probability to connect with any other node and therefore the efficiency of RWP is much lower. This is an important observation as it shows that in more realistic mobile scenarios efficiency might be affected by spatial correlations among links.



(a) Markov



(b) RWP



(c) RWPG 20

Fig. 8. Temporal robustness R_G as a function of probability of error P_{ERR} and for different values of P_{ON} for the different random models.

We also investigated RWPG with various group sizes and we present here three extreme situations: i) 20 groups of 5 nodes (RWPG 20), ii) 4 groups of 25 nodes (RWPG 4) and iii) 2 groups of 50 nodes (RWPG 2). Figure 7(b) presents the value of temporal efficiency E_G obtained for the RWPG case: in this model the spatial distribution of the groups appears to have a major impact on the efficiency. Group mobility is actually less

efficient than RWP: group members have high efficiency between them but much smaller efficiency with nodes that belong to other groups. Moreover, every RWPG scenario undergoes a transition in the trend of E_G : as P_{ON} increases, we see a particular value when the efficiency starts increasing more quickly. This behavior is due to the different groups finally being in direct temporal connection with each other, rather than connected through longer temporal paths. Before this transition, communication mainly occurs within single groups, so scenarios with larger groups have higher temporal efficiency. After this transition, scenarios with smaller groups become more efficient, since they enjoy very fast communication both within a single group and among different groups (because there are more groups evenly distributed in the simulation area). Also, this transition happens at higher values of P_{ON} for scenarios with larger groups since they are more clustered together, thus requiring longer ranges to inter-connect different groups.

4.4 Temporal Robustness

4.4.1 Markov-based Model

As reported in Figure 8(a), temporal robustness is affected by probability of error P_{ERR} in the same way as in the temporal ER model: the system fails gradually as we remove more nodes. However, we note that for intermediate values of P_{ON} robustness has lower values. Exactly in the same range of P_{ON} we see in Figure 7(a) that the Markov-model deviates from the temporal ER: this indicates how in that range of values memory effects result in a network which is not well connected nor highly dynamic, with consequently lower values of temporal efficiency. At the same time, high and low values of P_{ON} provide the same robustness, even if in these two cases the absolute value of temporal efficiency can be very different, thanks to the normalization of temporal robustness.

4.4.2 Mobility-based Models

In the case of mobility-based temporal networks, reported in Figure 8(b)-8(c), both RWP and RWPG exhibit a similar behavior: again, the network loses efficiency in a smooth way and temporal robustness is not affected by P_{ON} in this case as the spatial characteristics of the network are mainly affecting the resulting robustness.

5 CASE STUDY

We have seen that temporal networks do not exhibit sudden breakdowns when nodes are being removed and that various temporal network models exhibit analogies in their resilience. We now shift our attention to real time-varying networks: our aim is to understand whether temporal robustness gives us more information than static robustness in a real case and to investigate whether random models can offer a good approximation to real networks.

5.1 Dataset

This case study is based on Cabspotting, a publicly available dataset of mobility traces: the Cabspotting project tracked taxi cabs in San Francisco traveling through all the Bay Area for about two years with the aim of gathering data about city life [22]. The vehicles were equipped with GPS sensors and every device was periodically updating its position and uploading it to a central server to be stored, along with the timestamp of the record. Thus, it is possible to reconstruct each taxi's trajectory over space.

We have selected an area of about 20 km x 20 km around the city of San Francisco and we have extracted 24 consecutive hours of mobility traces, corresponding to Wednesday, 21 May 2008. After this, we have generated an artificial contact trace by defining a communication range of 200 m for the vehicles, which roughly corresponds to WiFi connectivity range in similar scenarios [23]: whenever two cars are within this distance they can communicate to each other. Time granularity is in seconds, so we have a sequence of 86,400 graphs with 488 nodes and more than 350,000 contacts among them. The average contact duration is about 2 minutes while the average inter-contact time is more than 2.5 hours.

5.2 Analysis

We study the reaction of the Cabspotting temporal network to random failures and compare it to our findings on random models. We adopt numerical simulation, but since the temporal dynamics of this network is not stationary, we can not compare the temporal efficiency E_G before and after a certain error, because the two temporal window will likely have already different properties. Instead, we fail nodes according to P_{ERR} at the very first time step of the temporal sequence of graphs: in this way, we can compare the average temporal efficiency over all the time for the original network and for the damaged one. We adopt a value of $\tau = 3600$, which allows us to consider temporal paths up to 1 hour, even if such longer paths can not contribute much to temporal efficiency.

The first comparison that we show in Figure 9(a) is between static robustness and temporal robustness for the Cabspotting temporal network. In this case static robustness is computed on the static graph obtained by aggregating all the contacts in the trace and adopting static global efficiency as performance measure. Since the resulting static graph contains more than 100,000 edges it is clearly an overestimation of the communication properties of the real system, as not all these links are continuously available over time and some paths can not be used due to temporal ordering constraints. Indeed, static robustness appears much larger than the temporal counterparts: only temporal robustness is able to capture the realistic communication capabilities of the system and how they are affected by random failures.

Then, we attempt to understand if the various random temporal network models we have studied can be used

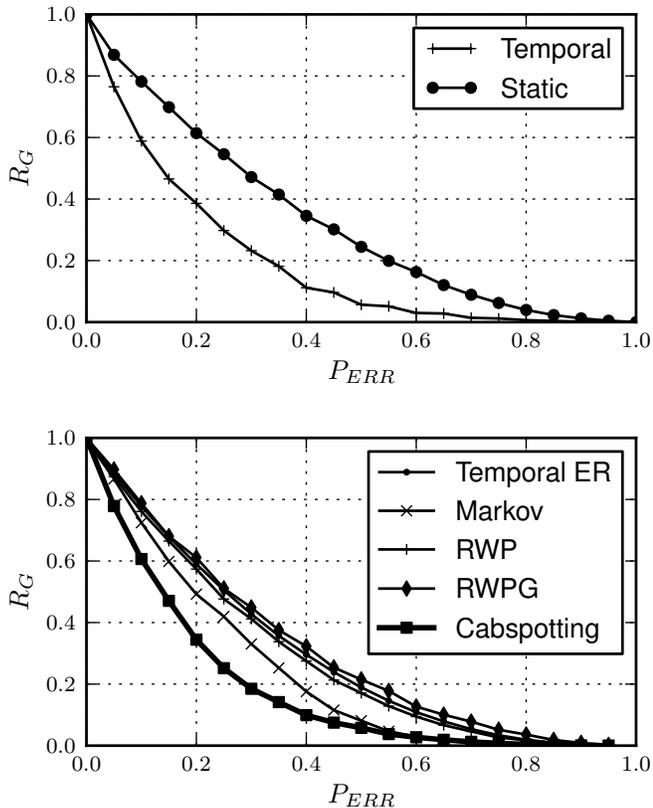


Fig. 9. Comparison between temporal robustness R_G and static robustness as a function of probability of error P_{ERR} for the Cabspotting dataset (Figure 9(a)). The static approach overestimates system robustness. Right: Comparison between the temporal robustness R_G of the dataset and random null models with the same number of nodes N and P_{ON} of the Cabspotting temporal network (Figure 9(b)).

to approximate the robustness properties of the real scenario. For each model, we compute the temporal robustness as a function of P_{ERR} for a network with the same number of nodes N and the same P_{ON} measured in the Cabspotting temporal network (about 0.005), using the same simulation parameters as in the real scenario. As reported in Figure 9(b), all temporal networks present the same trend in network robustness, albeit random models have higher values of temporal robustness than the real network. Interestingly, the closest match is the Markov-based temporal model, while the mobility-based models are closer to the ER model than to the Cabspotting network, even if this is actually a mobility-based contact network. However, the assumptions used in mobility models require homogeneity of space and absolute freedom to move continuously and independently in a boundless area, while in reality taxis are usually constrained to move on streets and bridges and they often move together along the same direction or stop together in a particular place to wait for customers (i.e., airport or stations). The Markov model, instead,

introduces the type of time correlations that appear to better mimic the real scenario. In fact, the most important aspect that needs to be captured is time ordering of events: in random mobility models connections do not follow particular time patterns, whereas real traces do (rush hour, working hours, human sleeping cycles). Only temporal robustness can take into account these unique characteristics.

These two results provide evidence that *temporal robustness is a more accurate measure to be used on mobile networks instead of standard static approaches*. Therefore, when testing protocols and applications to be deployed in mobile networks, a temporal study is more meaningful and should not be substituted by a static approximation.

6 IMPLICATIONS

In the previous sections we have seen how static robustness is not adequate to capture all aspects of mobile networks, especially real-world (i.e., non random) scenarios. Instead, a temporal approach allows for a better understanding of the robustness, since it takes into account time-dependent connections. This work presents many implications for the study of mobile networks and for the design of systems and applications in this domain.

First of all, a key advantage of our approach is that temporal robustness accurately models connectivity disruption in mobile networks: random models fail in a controlled way as we increase the fraction of removed nodes, without any sudden network disconnection. However, in a static ER random graph there could be a critical value of P_{ERR} which causes a complete breakdown of the system. This happens because the static analysis is unable to consider that time provides redundancy as new paths can still appear.

Moreover, our temporal robustness measure is not affected by the network size and the selection of τ . Yet the most interesting aspect is that although temporal efficiency is affected by P_{ON} , temporal robustness is not: it depends only on the relative efficiency drop caused by a network damage.

Another important property of the approach is that it does not overestimate connectivity. Time ordering and the temporal connectivity threshold τ exclude a number of connection paths that the static analysis would include. Therefore, the temporal model is able to correctly identify network connectivity disruptions, especially in real networks, where time ordering is important. For example, as we illustrated in Section 5, the static robustness analysis results in an almost linear relationship between removed nodes and loss of network connectivity. Instead, when temporal aspects are taken into account, the fact that temporal paths become longer due to less connection opportunities causes a sudden performance drop. This implication further supports our claim that static metrics cannot encapsulate the complexity of mobile networks.

Finally, although this temporal model may appear to be solely a theoretical tool to calculate the robustness of a network in simulation, there are practical implications too. This approach can be implemented in a real mobile network: each node i may maintain a table of *Lamport clocks* [24] that contains the shortest temporal distances to all known nodes that can contact i within τ time steps. Clocks are then updated whenever two nodes are in contact with each other. Apart from providing real-time connectivity information, this table can be used by each node to individually estimate temporal efficiency and robustness. It is worth mentioning that since temporal efficiency is based on the computation on shortest temporal distances, it represents an upper bound on the speed of information dissemination that a practical routing protocol can aspire to achieve. In addition, each node might detect anomalies and failures, even if they happen in remote parts of the network, by monitoring temporal efficiency over time. Routing algorithms can further exploit this property to self-tune their communication parameters and maintain a certain level of service in challenged environments.

7 COMPUTATIONAL ISSUES

The evaluation of the temporal network robustness for a mobile communication system relies on the computation of temporal efficiency at every time step: as a consequence, it is important to understand how temporal distances can be computed on a sliding window for a given temporal network in a computationally efficient way. Furthermore, decentralized and distributed approaches that can be used in real-world scenarios become particularly important and attractive.

More formally, in a time window $[t_1, t_2]$ the temporal distance $d_{ij}(t_1, t_2)$ can be computed by considering a message which is sent by node i at time t_1 and then flooded on the temporal network. At every time step the message can only be sent by each node to its neighbors, until it reaches node j at time $t^* \leq t_2$: thus, $d_{ij}(t_1, t_2) = t^* - t_1 + 1$, that is the temporal distance is exactly the number of time steps the message has traveled after t_1 to reach the destination. When the time window is slid to $[t_1 + 1, t_2 + 1]$, a new message is injected in the network at time $t_1 + 1$ and the time when it arrives at node j is recomputed to get $d_{ij}(t_1 + 1, t_2 + 1)$.

However, in this case the temporal distance between i and j will be tied to instant t_1 in such a way that if another path appears later in the same time window $[t_1, t_2]$, this is not considered. In a real communication system what really matters is how *fresh* is the information that a node has received from other nodes. Thus, we define a **relative temporal distance** by taking into account how a node i sends a message in every time steps within a time window $[t_1, t_2]$, while a node j will receive up to $k \leq t_2 - t_1$ among these messages, each one sent at time $t_{s_1}, t_{s_2}, \dots, t_{s_k}$. If we consider all the received messages and we define $t_s = \max(t_{s_1}, \dots, t_{s_k})$,

the relative temporal distance between the two nodes is given by $d_{ij}^r(t_1, t_2) = t_2 - t_s + 1$. This value is always relative to the current time step, differently to the original definition which is not affected by the progress of time.

Hence, relative temporal distance focuses on both how quickly a node can reach other nodes and on how frequently this happens over a certain period of time. With this definition, the most recent shortest path will always be considered and a short distance between two nodes needs to be continuously kept alive by new potential communication among them, otherwise it will gradually degrade with time. Hence, by adopting Lamport timestamps and message flooding mobile nodes can compute their relative temporal distances in a straightforward manner. Each node can then compute an estimation of global temporal efficiency by considering the distances between itself and the rest of the nodes.

8 RELATED WORKS

Previous related works lie in the two broad areas of temporal network analysis and network robustness studies.

One of the first attempts to generalize static network models to handle temporal information was to adopt time labels on edges to express temporal constraints on their presence [7]: this mainly algorithmic approach does not handle temporally disconnected nodes and, thus, is less suitable to investigate temporal networks arising from communication systems. Instead, some first attempts to investigate the properties of human contact networks reported on the temporal correlations and periodicities in these systems that arise at peculiar time scales [8], and on the impact of the frequency distribution of inter-contact time on delivery properties of opportunistic protocols [9]. More recently, the concept of network distance for temporal graphs has been formalized and explored [10], [11]. We build up on these results, by adopting these temporal measures in the definition of temporal network robustness.

The study of network robustness initially focused on how different classes of random networks exhibit different behavior when affected by random errors or targeted attacks [12]. In particular, exponential random graphs appear equally robust against both errors and attacks, while scale-free networks have higher resilience against random failures while being easily disrupted by intelligent attacks. Another approach to analytically study this problem is to exploit percolation models and explore how the network behaves as edges or nodes are removed [13]. Several extension on this topic have provided both extensive analysis on what type of attacks can be more disruptive in real networks [14] and dynamic models of failure which take into account how errors modify not only the structural properties of a network but also its dynamic properties [15]. Nonetheless, there have been no attempts to address the problem of network robustness in time-varying graphs. To our

knowledge, our work presents the first method to quantitatively evaluate network robustness by taking into account temporal properties.

9 CONCLUSIONS AND FUTURE WORK

This paper has presented a study of temporal robustness in time-varying network: we adopt temporal network metrics to assess network performance in presence of increasingly larger random failures. We have investigated the performance of our method both analytically on a random temporal network model and via simulations in a Markov-based and in two mobility-based models, exploring how the temporal dimension provides more redundancy to communication systems compared to static evaluation. Finally, we have shown how temporal robustness gives a more realistic estimation of the resilience of a real-world temporal network than standard static approaches. We plan to extend our work by taking into account attack strategies, where network damages might be correlated across different nodes, investigating whether random models and real networks react differently to this type of damage. Another interesting research direction would be to extend the theoretical analysis of the temporal network model to analyze more properties, such as the settling time of the system after a damage.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. Army Research Laboratory and the U.K. MOD under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. MOD or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D. Hwang, "Complex networks: Structure and dynamics," *Physics Reports*, vol. 424, no. 4-5, pp. 175-308, Feb. 2006.
- [2] P. Erdős and A. Rényi, "On the evolution of random graphs," *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, vol. 5, pp. 17-61, 1960.
- [3] R. Albert and A.-L. Barabási, "Statistical Mechanics of Complex Networks," *Review of Modern Physics*, vol. 74, pp. 47-97, 2002.
- [4] B. A. Huberman and L. A. Adamic, "Growth dynamics of the world-wide web," *Nature*, vol. 401, no. 6749, p. 131, 1999.
- [5] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proceedings of SIGCOMM '99*. New York, NY, USA: ACM, 1999, pp. 251-262.
- [6] V. Kostakos, "Temporal graphs," *Physica A*, vol. 388, no. 6, pp. 1007-1023, March 2009.
- [7] D. Kempe, J. Kleinberg, and A. Kumar, "Connectivity and inference problems for temporal networks," in *J. Comput. Syst. Sci.*, 2000, p. 2002.

- [8] A. Clauset and N. Eagle, "Persistence and Periodicity in a Dynamic Proximity Network," in *Proceedings of DIMACS Workshop on Computational Methods for Dynamic Interaction Networks*, September 2007.
- [9] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of Human Mobility on Opportunistic Forwarding Algorithms," *IEEE Trans. on Mobile Computing*, vol. 6, no. 6, pp. 606-620, 2007.
- [10] J. Tang, M. Musolesi, C. Mascolo, and V. Latora, "Temporal Distance Metrics for Social Network Analysis," in *Proceedings of WOSN '09*, Barcelona, Spain, August 2009.
- [11] J. Tang, S. Scellato, M. Musolesi, C. Mascolo, and V. Latora, "Small-world behavior in time-varying graphs," *Phys. Rev. E*, vol. 81, no. 5, p. 055101, May 2010.
- [12] R. Albert, H. Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378-382, July 2000.
- [13] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Network robustness and fragility: Percolation on random graphs," *Physical Review Letters*, vol. 85, no. 25, pp. 5468-5471, Dec 2000.
- [14] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han, "Attack vulnerability of complex networks," *Phys. Rev. E*, vol. 65, no. 5, May 2002.
- [15] P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda, "Error and attack tolerance of complex networks," *Physica A*, vol. 340, pp. 388-394, 2004.
- [16] V. Latora and M. Marchiori, "Efficient Behavior of Small-World Networks," *Physical Review Letters*, vol. 87, no. 19, October 2001.
- [17] P. Eugster, R. Guerraoui, A. M. Kermarrec, and L. Massoulié, "From Epidemics to Distributed Computing," *IEEE Computer*, vol. 37, no. 5, pp. 60-67, May 2004.
- [18] M. Musolesi and C. Mascolo, "Controlled Epidemic-style Dissemination Middleware for Mobile Ad Hoc Networks," in *Proceedings of MOBIQUITOUS '06*. ACM Press, July 2006.
- [19] A. E. F. Clementi and C. Macci and A. Monti and F. Pasquale and R. Silvestri, "Flooding Time in edge-Markovian Dynamic Graphs," in *Proceedings of PODC '08*, Toronto, Canada, Aug. 2008.
- [20] S. Hwang and D. Kim, "Markov Model of Link Connectivity in Mobile Ad Hoc Networks," *Telecomm. Sys.*, vol. 34, no. 1, pp. 51-58, 2007.
- [21] A. Medina, G. Gursun, P. Basu, and I. Matta, "On the Universal Generation of Mobility Models," in *Proceedings of IEEE/ACM MASCOTS '10*, Miami Beach, FL, US, August 2010.
- [22] M. Piorowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "CRAWDAD data set epfl/mobility (v. 2009-02-24)," Downloaded from <http://crawdad.cs.dartmouth.edu/epfl/mobility>, Feb. 2009.
- [23] A. Chaintreau, J. Y. Le Boudec, and N. Ristanovic, "The age of gossip: spatial mean field regime," in *Proceedings of SIGMETRICS '09*. New York, NY, USA: ACM, 2009, pp. 109-120.
- [24] L. Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Commun. ACM*, vol. 21, no. 7, pp. 558-565, 1978.



Salvatore Scellato is a PhD candidate in the Computer Laboratory, University of Cambridge. He holds a BSc in Computer Science (2006) and a MSc in Computer Science (2008) from the University of Catania, Italy. He has been working at University College London and at the University of Cambridge before joining the Computer Laboratory, and at Google during his PhD. His main research interests include online social networks, location-based services, human mobility modeling and related mobile and distributed systems. More information are available at <http://www.cl.cam.ac.uk/~ss824/>.



Dr. Ilias Leontiadis is a Research Fellow in the Computer Laboratory, University of Cambridge. He holds a PhD in Computer Science from University College London, United Kingdom (2009) and a MSc in Computer Systems from the University of Ioannina, Greece (2004). He also spent some time at University of California Los Angeles (UCLA) and Microsoft Research Cambridge. His research interests include vehicular networks, delay tolerant networking, mobile networking and systems, wireless sensor systems and mobility modeling. More information about his profile and his research work can be found at <http://www.cl.cam.ac.uk/~il235/>.

More information about his profile and his research work can be found at <http://www.cl.cam.ac.uk/~il235/>.



Dr Cecilia Mascolo is a Reader in Mobile Systems in the Computer Laboratory, University of Cambridge. She has published extensively in the areas of mobile computing, mobility and social network modelling. She is and has been PI of various national and international projects on social network analysis, sensing and mobility modelling. Dr Mascolo has served as a programme committee member of many mobile, sensor systems and networking conferences and workshops. She has taught various courses

on mobile and sensor systems and a Masters course on social and technological network analysis. More details of her profile are available at www.cl.cam.ac.uk/users/cm542.



Dr. Prithwish Basu is a Senior Scientist at Raytheon BBN Technologies. He has been playing leading roles in several DoD funded research and development programs at BBN. His research interests include network science, mobile ad hoc networks (MANET), and in general, theoretical aspects of networking. He is the technical lead on the US Army funded Network Science Collaborative Technology Alliance (NSCTA) program and the Principal Investigator for BBN on the US/UK International Technology Alliance (ITA) program. Prithwish received the MIT Technology Review's TR35 award (top 35 innovators under 35 years of age) in 2006 and is also a Senior Member of the IEEE. He has published over 50 papers on various topics in networking and also serves on the technical program committees of leading networking conferences such as IEEE INFOCOM and SECON. Prithwish holds a Ph.D. in Computer Engineering from Boston University and a B.Tech in Computer Science and Engineering from IIT Delhi.

Prithwish received the MIT Technology Review's TR35 award (top 35 innovators under 35 years of age) in 2006 and is also a Senior Member of the IEEE. He has published over 50 papers on various topics in networking and also serves on the technical program committees of leading networking conferences such as IEEE INFOCOM and SECON. Prithwish holds a Ph.D. in Computer Engineering from Boston University and a B.Tech in Computer Science and Engineering from IIT Delhi.



Dr. Murtaza Zafer (M'08) received the B.Tech degree in Electrical Engineering from the Indian Institute of Technology (IIT) Madras, India, in 2001, and the S. M. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (MIT), MA, USA, in 2003 and 2007 respectively. Currently, he is a Research Staff Member at IBM Thomas J. Watson Research Center, NY, USA, where his research focuses on the theory, performance analysis and design of computer and

communication networks including wireless, mobile, sensor and cloud computing networks. He spent the summer of 2004 at the Mathematical Sciences Research center, Bell Laboratories Alcatel-Lucent Inc., and the summer of 2003 at the R&D division of Qualcomm Inc. He co-authored the paper that received the Best Student Paper award at IEEE WiOpt'05 conference (2005). He is a recipient of the Siemens prize and Philips award in 2001 for academic excellence.