

Temporal Reachability Graphs

John Whitbeck^{†‡}
john.whitbeck@lip6.fr

Marcelo Dias de Amorim[†]
marcelo.amorim@lip6.fr

Vania Conan[‡]
vania.conan@fr.thalesgroup.com

Jean-Loup Guillaume[†]
jean-loup.guillaume@lip6.fr

[†]UPMC Sorbonne Universités
4, place Jussieu
75005 Paris
France

[‡]Thales Communications & Security
160 Boulevard de Valmy
92700 Colombes
France

ABSTRACT

While a natural fit for modeling and understanding mobile networks, time-varying graphs remain poorly understood. Indeed, many of the usual concepts of static graphs have no obvious counterpart in time-varying ones. In this paper, we introduce the notion of *temporal reachability graphs*. A (τ, δ) -reachability graph is a time-varying directed graph derived from an existing connectivity graph. An edge exists from one node to another in the reachability graph at time t if there exists a journey (i.e., a spatiotemporal path) in the connectivity graph from the first node to the second, leaving after t , with a positive edge traversal time τ , and arriving within a maximum delay δ . We make three contributions. First, we develop the theoretical framework around temporal reachability graphs. Second, we harness our theoretical findings to propose an algorithm for their efficient computation. Finally, we demonstrate the analytic power of the temporal reachability graph concept by applying it to synthetic and real-life datasets. On top of defining clear upper bounds on communication capabilities, reachability graphs highlight asymmetric communication opportunities and offloading potential.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Design and Architecture—Store and forward networks; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems; G.2.2 [Discrete Mathematics]: Graph Theory

General Terms

Algorithms, Theory, Performance

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'12, August 22–26, 2012, Istanbul, Turkey.

Copyright 2012 ACM 978-1-4503-1159-5/12/08 ...\$15.00.




Figure 1: Snapshots of a time-varying graph with five nodes.

Keywords

Time-varying graphs, Reachability, Opportunistic networks, Communication performance bounds

1. INTRODUCTION

In time-varying graphs (TVG), vertices and edges appear and disappear as a function of time. Alternatively called temporal networks [15] or evolving graphs [29], time-varying graphs have emerged over the past few years as a key model for a variety of complex systems [10, 13, 16, 26, 27]. Time-varying graphs can also serve as a solid theoretical framework for investigating fundamental properties of mobile networks [3, 9]. In Fig. 1, we show an example of a time-varying graph representing a five-node mobile network.

Many of the typical concepts of static graphs such as paths, distance, diameter, or node degree have no obvious counterpart in time-varying graphs. Theorems that are true on static graphs may not hold in time-varying ones [15] and dynamic equivalents of well-known problems on static graphs, such as finding strongly connected components, turn out to be intractable [2]. Naturally, it is always possible to track the evolution of static metrics such as node degree or clustering on snapshots of the time-varying graph at certain time intervals [6] but, while instructive, this approach often fails to capture the correlations between successive snapshots [14]. Recently a number of concepts specific to time-varying graphs have been studied such as *journeys* [29], *temporal diameter* [5], or *reachability time* [14].

In the context of opportunistic mobile networking, much attention has been focused on aggregate inter-contact time distributions. These have been found to fit a power-law (with an α parameter smaller than 1) followed by an exponential cutoff in a number of real-life datasets [4]. Should this power-law also hold for pairwise inter-contact times, then




Figure 2: Temporal reachability graphs obtained from the time-varying graph of Fig. 1 for $\delta=1s$ and $2s$. Arrows indicate uni-directional communication opportunities, while their absence indicates bi-directional ones. We show snapshots for different times t . We assume in this example that one-hop communications require $\tau=1s$.

the expected values of these would be infinite, a pessimistic result indeed. However, further work has shown that aggregate power-law distributions can emerge from a diversity of pairwise inter-contact laws (including exponential) [7, 22]. Moreover, all information on multi-hop connectivity is ignored. Inter-contact time measurements do not therefore suffice to characterize connectivity in mobile networks.

Temporal reachability graphs (TRG), our contribution, on the other hand, offer an immediate view of communication possibilities in a dynamic network. Given a time-varying graph \mathcal{G} , an arc exists from vertex a to vertex b at time t in its derived (τ, δ) -reachability graph if a journey exists in \mathcal{G} from a to b leaving a after time t and arriving at b before $t + \delta$, given that each single-hop communication takes time τ . Here, τ is the edge traversal time and δ is the upper bound on journey times in \mathcal{G} . By definition, temporal reachability are *directed* time-varying graphs. For example, in the context of studying information dissemination in delay-tolerant networks (DTN) [11], τ would be the one-hop transmission delay and δ is the maximum tolerated delay.

In Fig. 2, we show the instances of the temporal-reachability graph extracted from the opportunistic mobile networks of Fig. 1 at different times (assuming that the graph does not change between the snapshots shown in Fig. 1). We consider in this example that the link traversal time is $\tau=1s$. When $\delta = 1s$, the only possible journeys in \mathcal{G} at $t \in \{1s, 2s, 3s\}$ are the direct links between nodes – communications tolerate a delay of $1s$, which is the time to traverse one single link. When $\delta = 2s$, several more journeys appear. Consider for instance the reachability graph at $t = 0s$. There is a directed link between a and d as a can wait one time unit for the link (b, d) to appear ($a \rightsquigarrow b$ and $b \rightsquigarrow d$); thus, a two-hop journey $a \rightsquigarrow d$ through b becomes possible.

We make three main contributions:

Temporal reachability graphs. We formalize the concept of reachability graphs. In particular, under conditions that hold for all empirical datasets, we prove an additive property on the delay of reachability graphs. Roughly speaking, knowledge of the reachability graphs for delays δ and μ is enough to derive the reachability graph for delay $\delta + \mu$.

Table 1: Notations used in this paper

Notation	Meaning
N	Number of vertices
\mathcal{G}	A time-varying graph (TVG)
τ	Edge traversal time of a TVG
η	Time step of an η -regular TVG
\mathcal{R}_δ	Reachability graph with maximum delay δ derived from a TVG \mathcal{G}
$\ \mathcal{R}_\delta(t)\ $	Number of arcs in \mathcal{R}_δ at time t
$\mathcal{U}_\delta, \mathcal{L}_\delta$	Upper and lower bounds of a reachability graph \mathcal{R}_δ . $\mathcal{L}_\delta \subseteq \mathcal{R}_\delta \subseteq \mathcal{U}_\delta$.
\mathcal{D}_δ	A time-varying out-dominating set of a reachability graph \mathcal{R}_δ
$\text{sym}(\mathcal{R}_\delta)$	Symmetric subset of \mathcal{R}_δ : $(u, v) \in \text{sym}(\mathcal{R}_\delta)(t) \Leftrightarrow (v, u) \in \text{sym}(\mathcal{R}_\delta)(t)$
$\text{asym}(\mathcal{R}_\delta)$	Asymmetric subset of \mathcal{R}_δ : $(u, v) \in \text{asym}(\mathcal{R}_\delta)(t) \Leftrightarrow (v, u) \notin \text{asym}(\mathcal{R}_\delta)(t)$

Algorithm for efficient computation of TRG. We translate this additive property into an efficient single-pass streaming algorithm [1] to calculate entire families of reachability graphs.

Insights into real-world mobility traces. Once calculated, reachability graphs yield many original insights on the temporal structure of the time-varying graphs they derive from. In this paper, we calculate the reachability graphs of several synthetic and real-life high-resolution connectivity traces [18, 20, 23, 25]. These highlight the concepts of *temporal connectivity*, *temporal asymmetry*, and *temporal dominating set*. From the point of view of opportunistic communications, reachability graphs immediately provide, at all times, the maximum delivery ratio and the minimum set of users required to broadcast a message to the entire network, thereby offering a complete view of opportunistic communication capabilities and offloading potential.

The rest of the paper is structured as follows. Sections 2 and 3 introduce the concepts of time-varying and reachability graphs. Their theoretical properties justify their computation by a streaming algorithm detailed in Section 4. Section 5 then analyzes the reachability graphs of synthetic and real-life connectivity graphs. Related work is presented in Section 6. Finally, Section 7 concludes.

2. TIME-VARYING GRAPHS

Time-varying graphs are a very useful high-level abstraction for studying connectivity over time in mobile network. Terms, notations, and definitions around time-varying graphs vary considerably, but recently, Casteigts et al. have proposed a unified framework for TVGs and, wherever applicable, we will use their definitions and notations [3].

Definition 1. (Time-varying graph) Let V be a set of vertices, and $E \subseteq V \times V$ the set of possible edges between vertices in V . Events occur over a time span $\mathcal{T} \subseteq \mathbb{T}$, where \mathbb{T} is the temporal domain (\mathbb{N} for discrete-time systems or \mathbb{R}^+ for continuous-time ones). In the general case, a TVG is a tuple $\mathcal{G} = (V, E, \mathcal{T}, \rho, \zeta)$ where $\rho : E \times \mathcal{T} \rightarrow \{0, 1\}$, called *presence function*, indicates whether a given edge exists at a given time and $\zeta : E \times \mathcal{T} \rightarrow \mathbb{T}$, called *latency function*, indicates the time it takes to cross a given edge if starting at a given date (the latency of an edge could vary in time).

In this paper, we will mostly consider continuous-time TVGs (i.e., $\mathbb{T} = \mathbb{R}^+$). Furthermore, we will always assume a constant ζ function such that $\forall (e, t) \in E \times \mathcal{T}, \zeta(e, t) = \tau$ where $\tau \geq 0$ is our uniform edge traversal time. We also note $\mathcal{G}(t) \subseteq E$ the set of edges in the snapshot of \mathcal{G} at time t . Hence $e \in \mathcal{G}(t) \Leftrightarrow \rho(e, t) = 1$. Finally, given an edge $e = (u, v) \in E$, we define $\text{from}(e) = u$ and $\text{to}(e) = v$.

Definition 2. (Inclusion) Let \mathcal{G} and \mathcal{G}' be two TVGs that differ only by their presence functions ρ and ρ' . We write $\mathcal{G} \subseteq \mathcal{G}'$ if and only if $\forall t \in \mathcal{T}, \forall e \in E, \rho(e, t) \leq \rho'(e, t)$, or equivalently if and only if $\forall t \in \mathcal{T}, \mathcal{G}(t) \subseteq \mathcal{G}'(t)$.

Definition 3. (Union) Let \mathcal{G}, \mathcal{H} , and \mathcal{I} be three TVGs that differ only by their presence functions. We write $\mathcal{G} = \mathcal{H} \cup \mathcal{I}$ if and only if $\forall t \in \mathcal{T}, \mathcal{G}(t) = \mathcal{H}(t) \cup \mathcal{I}(t)$.

Definition 4. (Journey) A journey in \mathcal{G} is a sequence of couples $\mathcal{J} = \{(e_1, t_1), (e_2, t_2), \dots, (e_k, t_k)\}$ such that $\forall i < k$: (i) $\text{to}(e_i) = \text{from}(e_{i+1})$; (ii) $\forall t$ s.t. $t_i \leq t < t_i + \tau, \rho(e_i, t) = 1$; and (iii) $t_{i+1} \geq t_i + \tau$. Here, $|\mathcal{J}| = k$ is this journey's *topological length* (i.e., the number of hops). Furthermore, $\text{departure}(\mathcal{J})$ and $\text{arrival}(\mathcal{J})$ denote the starting date t_1 and the last date $t_k + \tau$, respectively. Finally $\delta_{\mathcal{J}} = \text{arrival}(\mathcal{J}) - \text{departure}(\mathcal{J})$ is the journey's *temporal length*. A journey may represent, for example, the sequence of hops that a message follows through an opportunistic network.

3. TEMPORAL REACHABILITY GRAPHS

Now that we have the necessary background, let us formally define the notion of *temporal reachability graphs* (or simply *reachability graphs* in the remainder of this paper, for the sake of readability). Furthermore, we write \mathbb{N}^* for $\mathbb{N} \setminus \{0\}$. To the best of our knowledge, the definitions and results presented in this section are completely novel.

3.1 Reachability Graphs

Definition 5. (Reachability graph) For $\delta \in \mathbb{R}^+$, let \mathcal{R}_δ be the reachability graph with maximum delay δ derived from \mathcal{G} (with edge traversal time τ). Formally, $\forall t, (u, v) \in \mathcal{R}_\delta(t)$ if and only if $u \neq v$ and there exists at time t in \mathcal{G} a journey \mathcal{J} from u to v such that $\text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J}) \leq t + \delta$.

Note that if $\delta < \tau$ then \mathcal{R}_δ is empty, as even one-hop journeys do not have enough time to arrive before the maximum delay.

PROPOSITION 1 (Growth). Let \mathcal{R}_δ and \mathcal{R}_μ be two reachability graphs of \mathcal{G} . Then $\delta \leq \mu \implies \mathcal{R}_\delta \subseteq \mathcal{R}_\mu$.

This follows naturally from the definition of a reachability graph. Note that the reverse is not true. For example let \mathcal{G} be a TVG with edge traversal time $\tau > 0$. Its derived reachability graphs $\mathcal{R}_{\tau/2}$ and $\mathcal{R}_{\tau/3}$ are both empty as their maximum delays are smaller than the time it takes to cross one edge. Therefore, as per definition 2, $\mathcal{R}_{\tau/2} \subseteq \mathcal{R}_{\tau/3}$ even though $\frac{\tau}{2} > \frac{\tau}{3}$.

Definition 6. (Composition) Let \mathcal{R}_δ and \mathcal{R}_μ be two reachability graphs of \mathcal{G} . We define their *composition* $\mathcal{R}_\delta \otimes \mathcal{R}_\mu$ as the TVG such that, at all times t ,

$$(u, v) \in (\mathcal{R}_\delta \otimes \mathcal{R}_\mu)(t) \Leftrightarrow \begin{cases} (u, v) \in \mathcal{R}_\delta(t), \text{ or} \\ (u, v) \in \mathcal{R}_\mu(t + \delta), \text{ or} \\ \exists w \in V, (u, w) \in \mathcal{R}_\delta(t) \text{ and} \\ (w, v) \in \mathcal{R}_\mu(t + \delta) \end{cases}$$

THEOREM 1 (Decomposition). Let \mathcal{G} be a TVG with edge traversal time τ and $\mathcal{R} = \{\mathcal{R}_\delta\}_{\delta \geq 0}$ the set of all its reachability graphs. Let $\delta \geq 0$ and $\mu \geq \tau$. Then:

$$\mathcal{R}_{\delta+\mu} = \bigcup_{0 \leq \epsilon < 1} \mathcal{R}_{\delta+\epsilon\tau} \otimes \mathcal{R}_{\mu-\epsilon\tau}.$$

While not yet in a calculable form, this theorem is at the heart of our approach. Intuitively, it states that if all the reachability graphs with delays close to δ and μ are known, then one can calculate the reachability graph with delay $\delta + \mu$. A full formal proof is provided in the appendix (Section A.1), but intuitively, it relies on the following simple idea. Any arc in $\mathcal{R}_{\delta+\mu}$ corresponds to a journey in \mathcal{G} starting after time t and arriving before $t + \delta + \mu$. Either it can be neatly divided into a (possibly empty) journey in \mathcal{R}_δ starting after t and arriving before $t + \delta$ and another (possibly empty) journey in \mathcal{R}_μ starting after $t + \delta$ and arriving before $t + \delta + \mu$, or it starts crossing an edge before $t + \delta$ but after $t + \delta - \tau$ that straddles both time intervals. In the latter case, since the edge traversal time is τ , one can incorporate the straddling edge into a journey arriving before $t + \delta + \epsilon\tau$ and the rest into another journey leaving after $t + \delta + \epsilon\tau$ of temporal length less than $\mu - \epsilon\tau$.

3.2 Regular reachability graphs

Real-life datasets all have a *maximum resolution* (e.g., a second or a millisecond) that corresponds to the precision with which they were measured. While it is tempting to map their time domain \mathcal{T} to \mathbb{N} , in reality one must account for 0-second edge durations. For example, if the edge traversal time τ is null, then even ephemeral edges that last 0 seconds may be part of a journey. In the more common situation where τ is strictly positive, one may still encounter 0-duration arcs in a reachability graph. For example, consider an edge (u, v) that is present for only one second. If τ and δ are both also equal to one second, then a one-hop journey from u to v using that edge only exists at precisely the instant t that the edge appears. This, in turn, corresponds to two 0-second arcs in $\mathcal{R}_\delta(t)$: (u, v) and (v, u) . Note that both of these arcs will become one-second-long arcs in the reachability graph when $\delta = 2s$. These observations lead to the definition of *regular time-varying graphs*.

Definition 7. (Regular TVG) A TVG \mathcal{G} is an η -regular TVG if there exists $\eta > 0$ such that $\forall k \in \mathbb{N}, k\eta < t_1 \leq t_2 < (k+1)\eta \implies \mathcal{G}(t_1) = \mathcal{G}(t_2) \subseteq \mathcal{G}(k\eta)$. Here, η is called \mathcal{G} 's resolution. The time interval $[k\eta, (k+1)\eta[$ is the k^{th} epoch of \mathcal{G} with starting time $k\eta$ and ending time $(k+1)\eta$.

Without loss of generality, this definition assumes that the first epoch starts at $t = 0$. Intuitively, an η -regular TVG is one whose instantaneous graph topology cannot change arbitrarily quickly, as, during each epoch, the graph topology remains constant. Ephemeral 0-second edges or arcs may exist at the start of an epoch, but the TVG then remains constant until the start of the next epoch. Not only are regular TVGs a natural fit for real-life traces, but, as we will see, they also represent a class of TVGs whose reachability graphs are calculable.

THEOREM 2 (Regular reachability graphs). Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau \in \eta\mathbb{N}^*$. For $\delta \in \eta\mathbb{N}$, let \mathcal{R}_δ be a reachability graph of \mathcal{G} . Then \mathcal{R}_δ is an η -regular TVG and $\forall k \in \mathbb{N}, k\eta < t < (k+1)\eta \implies \mathcal{R}_\delta(t) \subseteq \mathcal{R}_\delta(k\eta) \cap \mathcal{R}_\delta((k+1)\eta)$.

This theorem ensures that, as long as δ is a “multiple” of the resolution η , \mathcal{R}_δ carries on the η -regularity of \mathcal{G} . The full formal proof is a little technical (Section A.2 in the appendix), but it relies on the following idea. Let us consider a one-hop journey in an η -regular TVG \mathcal{G} leaving in the middle of an epoch at time t . Let k be the integer such that $k\eta < t < (k+1)\eta$. It arrives in the next epoch at time $(k+1)\eta < t + \tau < (k+2)\eta$ (the theorem assumes that $\tau \geq \eta$). Because \mathcal{G} is η -regular, the departure time of this one-journey can be nudged forwards or backwards as long as it remains with the same epoch. Hence all one-hop journeys with a departure time t' such that $k\eta \leq t' \leq (k+1)\eta$ are also valid journeys in \mathcal{G} . This “nudging” can be extended, though not trivially, to multi-hop journeys which proves the theorem.

Furthermore, the $\tau \geq \eta$ hypothesis guarantees that an arc present in a reachability graph during an epoch is present not only at the start of the epoch like in any regular graphs both also at the start of the next epoch. This property will be leveraged later in this section. However $\mathcal{R}_\delta(t) \neq \mathcal{R}_\delta(k\eta) \cap \mathcal{R}_\delta((k+1)\eta)$. An example of such a situation is discussed in Section 3.3.

THEOREM 3 (Sampling). *Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau = n\eta$ with $n \in \mathbb{N}^*$. Let $\bar{\mathcal{R}} = \{\mathcal{R}_{i\eta}\}_{i \in \mathbb{N}}$ be the set of all its derived η -regular reachability graphs. For $(d, m) \in \mathbb{N} \times \mathbb{N}$ such that $m \geq n$, we have, $\forall a \in \mathbb{N}$:*

$$\mathcal{R}_{(d+m)\eta}(a\eta) = \bigcup_{0 \leq k < n} (\mathcal{R}_{(d+k)\eta} \otimes \mathcal{R}_{(m-k)\eta})(a\eta).$$

This theorem is the adaptation of Theorem 1 to regular TVGs. The formal proof is in the appendix (Section A.3) but, as previously, it relies on the idea that within an epoch, there exists a little freedom to “nudge” journey departure and arrival times backwards or forwards. In a sense, if both the minimum departure time and maximal arrival time of a journey fall exactly at the start of a time epoch, then this journey can be divided into two sub-journeys whose departure and arrival times also map exactly to the start of epochs. Under this form, the value of a reachability graph at the start of each epoch can be *exactly* calculated from the values, at the start of each epoch, of the proper set of τ pairs of reachability graphs.

Unfortunately, this exact formula does not extend to the time spent strictly within an epoch. In a way, the state of the reachability graph during epochs is more important than its state at their starting times. Indeed, the time spent at *exactly* the start of epochs is infinitesimally small. Therefore any metric averaged over time (e.g. average density) will depend only on the states strictly within epochs. In the next section, we propose an upper and a lower bound on the reachability graph during epochs. As we will later see in Section 5, these upper and lower bounds are in fact nearly always equal and therefore achieve an excellent approximation of the real reachability graph.

3.3 Upper and lower bounds

We define for every regular reachability graph \mathcal{R}_δ an upper and lower bound, i.e., two TVGs \mathcal{U}_δ and \mathcal{L}_δ such that $\mathcal{L}_\delta \subseteq \mathcal{R}_\delta \subseteq \mathcal{U}_\delta$. Both of these bounds are equal to the reachability graph at the start of each epoch but only differ slightly during the epoch. The upper bound is a straightforward application of Theorem 2.

PROPOSITION 2 (Upper Epoch approximation). *Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau \in \eta\mathbb{N}^*$. For $\delta \in \eta\mathbb{N}$, let \mathcal{R}_δ be a reachability graph of \mathcal{G} . We define \mathcal{R}_δ 's upper approximation, the TVG \mathcal{U}_δ , as follows ($\forall a\eta \leq t < (a+1)\eta$ with $a \in \mathbb{N}$):*

- if $t = a\eta$, then $\mathcal{U}_\delta(a\eta) = \mathcal{R}_\delta(a\eta)$;
- if $t > a\eta$, then $\mathcal{U}_\delta(t) = \mathcal{R}_\delta(a\eta) \cap \mathcal{R}_\delta((a+1)\eta)$.

As defined, $\forall \delta, \mathcal{R}_\delta \subseteq \mathcal{U}_\delta$.

The lower bound is a little more complicated and involves a modified version of the composition operator (Definition 6).

Definition 8. (Approximate composition) Let \mathcal{R}_δ and \mathcal{R}_μ be two reachability graphs of an η -regular TVG \mathcal{G} with δ and μ in $\eta\mathbb{N}^*$. Let \mathcal{L}_δ and \mathcal{L}_μ be two TVGs such that $\mathcal{L}_\delta \subseteq \mathcal{R}_\delta$ and $\mathcal{L}_\mu \subseteq \mathcal{R}_\mu$. For all times $a\eta \leq t < (a+1)\eta$ with $a \in \mathbb{N}$, we define their *approximate composition* $\mathcal{L}_\delta \odot \mathcal{L}_\mu$ as follows:

- if $t = a\eta$, then $(\mathcal{L}_\delta \odot \mathcal{L}_\mu)(a\eta) = (\mathcal{R}_\delta \otimes \mathcal{R}_\mu)(a\eta)$;
- if $t > a\eta$, then

$$(u, v) \in (\mathcal{L}_\delta \odot \mathcal{L}_\mu)(t) \Leftrightarrow \begin{cases} (u, v) \in \mathcal{R}_\delta((a+1)\eta), \text{ or} \\ (u, v) \in \mathcal{R}_\mu(a\eta + \delta), \text{ or} \\ \exists w \in V, (u, w) \in \mathcal{L}_\delta(t) \text{ and} \\ (w, v) \in \mathcal{L}_\mu(t + \delta) \end{cases}$$

By definition, at the start of each epoch the approximate composition is equal to the regular composition of the reachability graphs. During an epoch, the condition $(u, v) \in \mathcal{R}_\delta((a+1)\eta)$ is *easier* to meet than the more intuitive $(u, v) \in \mathcal{L}_\delta(t)$. Indeed, with Theorem 2, $\mathcal{L}_\delta(t) \subseteq \mathcal{R}_\delta(t) \subseteq \mathcal{R}_\delta((a+1)\eta)$, and this composition operator will catch the arcs (including the 0-second ones) occurring at the start of the next epoch. The same is true for the $(u, v) \in \mathcal{R}_\mu(a\eta + \delta)$ condition for arcs ending at the start of the current (δ -shifted) epoch. Combined, these more inclusive conditions make for a *tighter* lower bound.

PROPOSITION 3 (Lower Epoch approximation). *Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau = n\eta$ with $n \in \mathbb{N}^*$. Let $\bar{\mathcal{R}} = \{\mathcal{R}_\delta\}_{\delta \in \eta\mathbb{N}}$ be the set of all its derived reachability graphs. We recursively define \mathcal{R}_δ 's lower approximation (i.e., the TVG \mathcal{L}_δ) as follows. For $d < 2n$, $\mathcal{L}_{d\eta} = \mathcal{R}_{d\eta}$ and for all (d, m) in $\mathbb{N} \times \mathbb{N}$ such that $d \geq n$ and $m \geq n$,*

$$\mathcal{L}_{(d+m)\eta} = \bigcup_{0 \leq k < n} \mathcal{L}_{(d+k)\eta} \odot \mathcal{L}_{(m-k)\eta}.$$

As defined, $\forall d \in \mathbb{N}, \mathcal{L}_{d\eta} \subseteq \mathcal{R}_{d\eta}$.

The proof, detailed in Section A.4 of the appendix, is fairly straightforward and involves setting the value of ϵ in Theorem 1 as a function of η , t , k , and n .

This recursive definition means that the exact value of \mathcal{L}_δ during a time epoch *depends on the sequence of compositions that was used to calculate it*. For example, in the simple situation where $\tau = \eta = 1$, the lower bounds $\mathcal{L}_4 = \mathcal{L}_2 \odot \mathcal{L}_2$ and $\mathcal{L}'_4 = \mathcal{L}_3 \odot \mathcal{L}_1$ may be different. Proposition 3 only guarantees that they are both included in the real reachability graph \mathcal{R}_4 . However, as we will see, this is not a concern




Figure 3: Upper epoch approximation with $\tau = \eta = 1$. Arrows point in the direction of time. Intervals denote the presence of an edge or arc during that time. Dots are intervals reduced to a single point in time. In this example the lower bound \mathcal{L}_2 is equal to the reachability graph \mathcal{R}_2 but the upper bound \mathcal{U}_2 is not.




Figure 4: Lower epoch approximation with $\tau = \eta = 1$. Arrows point in the direction of time. Intervals denote the presence of an edge or arc during that time. Dots are intervals reduced to a single point in time. In this example the upper bound \mathcal{U}_4 is equal to the reachability graph \mathcal{R}_4 but the lower bound \mathcal{L}_4 is not.

because this lower bound is tight regardless of how it is calculated.

Fig. 3 is an example of a situation with four vertices where the upper bound has an arc during an epoch that is not in the reachability graph. Here $\tau = \eta = 1$ s. \mathcal{R}_1 is exactly derived from \mathcal{G} by shortening the end time of each edge by τ . For example, this leads to an ephemeral arc from a to b and a two-second-long arc from b to d (down from a three-second-long edge (b, d) in \mathcal{G}). Obviously, not all arcs are represented on Fig. 3. In \mathcal{R}_2 , there are two ephemeral arcs from a to d , corresponding to the $a \rightsquigarrow b \rightsquigarrow d$ journey, and the $a \rightsquigarrow c \rightsquigarrow d$ journey an epoch later. In between those instants, it is both too late to use the first journey and too early to use the second. Since this arc exists at the start of two successive epochs, \mathcal{U}_2 errs in considering that it exists during the entire epoch. In this example, the lower bound \mathcal{L}_2 matches the reachability graph.

Conversely, Fig. 4 is an example of a situation with three vertices where the lower bound misses an arc during an

epoch that is present in the reachability graph. Here again $\tau = \eta = 1$. The one-second-long edges in \mathcal{G} all become ephemeral 0-second arcs in \mathcal{R}_1 (not shown on the figure), that in turn become one-second-long arcs in \mathcal{R}_2 . Furthermore, \mathcal{R}_2 also contains two ephemeral arcs (a, c) and (b, d) that correspond to the $a \rightsquigarrow b \rightsquigarrow c$ and $b \rightsquigarrow c \rightsquigarrow d$ journeys respectively. Here $\mathcal{R}_2 = \mathcal{L}_2 = \mathcal{U}_2$. Finally, \mathcal{R}_4 contains an arc (a, d) that corresponds to the existence of the three-hop journey $a \rightsquigarrow b \rightsquigarrow c \rightsquigarrow d$. \mathcal{L}_4 correctly identifies this journey at two instants k and $k + 1$. At time k , it composes arc (a, b) from $\mathcal{R}_2(k)$ with the ephemeral arc (b, d) in $\mathcal{R}_2(k + 2)$. At time $k + 1$ it composes the ephemeral arc (a, c) in $\mathcal{R}_2(k + 1)$ with the arc (c, d) in $\mathcal{R}_2(k + 3)$. However, for $k < t < k + 1$, \mathcal{L}_4 has no way of finding the journey from a to b . Note that in this case, \mathcal{U}_4 is equal to \mathcal{R}_4 .

The example of Fig. 4 also helps us understand the absence of error propagation when calculating successive lower approximations from previous lower approximations. In this example, thanks to the conditions in Definition 8 that refer to the starting times of epochs, \mathcal{L}_8 will “bridge” over the missing (a, d) arc in \mathcal{L}_4 . However it may miss composed arcs that use the (a, d) arc in \mathcal{R}_4 . These missed arcs will be, in-turn, “bridged” in, for instance, a \mathcal{L}_{16} approximation. Roughly speaking, errors during epochs do not propagate beyond two compositions. Combined with the fact that the lower approximation matches the reachability graph at the start of each epoch, this leads, as we will verify empirically, to a tight lower approximation or the reachability graph.

3.4 A few case studies

In this section, we examine how the theoretical results above apply to some simple situations.

Zero edge traversal time. The reachability graphs of η -regular TVGs with zero edge traversal time ($\tau = 0$) are very easy to calculate. In particular $(u, v) \in \mathcal{R}_0(t)$ if and only if u and v are in the same connected component at time t . Similarly, $\mathcal{R}_\eta(t)$ can be calculated from the connected components at time t and $t + \eta$. Furthermore, Theorem 1 becomes: $\mathcal{R}_{\delta+\mu} = \mathcal{R}_\delta \otimes \mathcal{R}_\mu$. This relation can then be applied repeatedly (e.g., in a binary-exponentiation) to obtain any \mathcal{R}_δ with $\delta \in \eta\mathbb{N}^*$. Furthermore, the reachability graphs thus calculated are *exact*. However, in this specific case where the edge traversal time is null, more efficient algorithms for calculating reachability graphs exist. For example the algorithm proposed by Chaintreau et al. for computing delay-optimal paths could be easily adapted to this purpose [5].

Several hops per epoch. While all of the results on regular reachability graphs in Sections 3.2 and 3.3 assume that the edge traversal time τ is in $\eta\mathbb{N}^*$, they can be simply adapted to the case where $\tau = \eta/n$ with $n \in \mathbb{N}^*$, i.e., when up to n edges may be crossed during a single epoch. Indeed, in this case, any η -regular graph would also be τ -regular, and we can apply all of our results as if $\tau = \eta$. For example, Proposition 3 is reduced to $\mathcal{L}_{(d+m)\eta} = \mathcal{L}_{d\eta} \odot \mathcal{L}_{m\eta}$.

Unit delay. For $\tau > 0$, the reachability graph \mathcal{R}_τ is trivially calculable from \mathcal{G} . Indeed, for any edge (u, v) in \mathcal{G} that appears at time t_1 and disappears at time $t_2 \geq t_1 + \tau$, the arcs (u, v) and (v, u) appear in \mathcal{R}_τ at time t_1 and disappear at time $t_2 - \tau$. This derivation is simple but essential for bootstrapping iterations of lower bound compositions (e.g., repeated applications of Proposition 3).




Figure 5: Adding $\bar{\mathcal{L}}_4$ and $\bar{\mathcal{L}}_8$ families to obtain an $\bar{\mathcal{L}}_4 \oplus \bar{\mathcal{L}}_8 = \bar{\mathcal{L}}_{12}$ family. In this example $\tau = 2$ and $\eta = 1$.

4. EFFICIENT COMPUTATION OF REACHABILITY GRAPHS

4.1 Families of reachability graphs

Another interesting property of lower approximations of reachability graphs is that an upper approximation may be derived from it. Indeed, since it contains the exact values of the reachability graph at the start of each epoch, it is trivial to calculate their intersection during each epoch. While the method described in this section focuses on the efficient computation of the lower approximations of reachability graphs, it simultaneously computes the upper approximations.

At a high level, the algorithm presented in this section is a *binary exponentiation* on families of lower bounds of reachability graphs using a special additive operator. To simplify notations, we will consider in this section that $\eta = 1$ and $\tau \in \mathbb{N}^*$.

Definition 9. (Lower bound family) Let \mathcal{G} be a 1-regular TVG whose edge traversal time is $\tau \in \mathbb{N}^*$. For $d \in \tau\mathbb{N}$, let \mathcal{L}_d be the lower approximation of a reachability graph \mathcal{R}_d of \mathcal{G} . For $d \geq \tau$, we define $\bar{\mathcal{L}}_d$, the *family* of \mathcal{L}_d such that $\bar{\mathcal{L}}_d = \{\mathcal{L}_{d+i}\}_{-\tau < i < \tau}$.

PROPOSITION 4 (Family additivity). Let \mathcal{G} be a 1-regular TVG whose edge traversal time is $\tau \in \mathbb{N}^*$. Let \mathcal{R}_d and \mathcal{R}_m be two reachability graphs of \mathcal{G} such that $d \in \tau\mathbb{N}^*$ and $m \in \tau\mathbb{N}^*$, and $\bar{\mathcal{L}}_d$ and $\bar{\mathcal{L}}_m$ the respective families of their lower approximations. We define $\bar{\mathcal{L}}_d \oplus \bar{\mathcal{L}}_m$ the set of TVGs such that

$$\bar{\mathcal{L}}_d \oplus \bar{\mathcal{L}}_m = \left\{ \bigcup_{0 \leq k < \tau} \mathcal{L}_{d+k} \odot \mathcal{L}_{m+i-k} \right\}_{0 \leq i < \tau} \cup \left\{ \bigcup_{0 \leq k < \tau} \mathcal{L}_{d+i+k} \odot \mathcal{L}_{m-k} \right\}_{-\tau < i < 0}$$

Then (i) any element $\mathcal{L}_{d+m+i} \in \bar{\mathcal{L}}_d \oplus \bar{\mathcal{L}}_m$ with $-\tau < i < \tau$ is a lower approximation of \mathcal{R}_{d+m+i} , and (ii) any TVG in $\bar{\mathcal{L}}_{d+m}$ can be calculated from τ compositions of pairs of TVGs in $\bar{\mathcal{L}}_d \times \bar{\mathcal{L}}_m$.

PROOF. Consider any lower approximation $\mathcal{L}_{d+m+i} \in \bar{\mathcal{L}}_{d+m}$. If $-\tau < i < 0$, $\mathcal{L}_{d+m+i} = \mathcal{L}_{(d+i)+m} = \bigcup_{0 \leq k < \tau} \mathcal{L}_{d+i+k} \odot \mathcal{L}_{m-k}$ (Theorem 3). If $0 \leq i < \tau$, $\mathcal{L}_{d+m+i} = \mathcal{L}_{d+(m+i)} = \bigcup_{0 \leq k < \tau} \mathcal{L}_{d+k} \odot \mathcal{L}_{m+i-k}$ (Theorem 3). \square

We now have our self-sufficient elements, the lower bound families, and an additive operation \oplus between them. For

illustration, Fig. 5 details the process of adding $\bar{\mathcal{L}}_4$ and $\bar{\mathcal{L}}_8$ families to obtain an $\bar{\mathcal{L}}_4 \oplus \bar{\mathcal{L}}_8 = \bar{\mathcal{L}}_{12}$ family. The \mathcal{L}_3 , \mathcal{L}_4 , \mathcal{L}_5 , \mathcal{L}_7 , \mathcal{L}_8 , and \mathcal{L}_9 lower bounds are combined as inputs for three applications of Proposition 3 that yield \mathcal{L}_{11} , \mathcal{L}_{12} , and \mathcal{L}_{13} . Viewed as a black box, this operation combines the $\bar{\mathcal{L}}_4$ and $\bar{\mathcal{L}}_8$ families into an $\bar{\mathcal{L}}_{12}$ family.

The inner workings of the \oplus operator are embarrassingly parallel. Indeed, each application of Proposition 3 can be run completely independently of the others. This opens the way for highly distributed implementations, whose speed will be determined by that of the composition of τ pairs of lower bounds of reachability graphs. Accordingly, the next section proposes an efficient algorithm for this composition operation.

4.2 Composing reachability graphs

A time-varying graph \mathcal{G} may be stored as a time-indexed sequence of edge UP and DOWN events. For example, if at time t an event $\{(u, v), \text{DOWN}\}$ occurs, then the edge (u, v) disappears at time t in $\mathcal{G}(t)$. Such a representation is well suited for algorithms that sequentially examine all states of the TVG.

We present a streaming algorithm for composing τ reachability graph lower bounds as in Proposition 3. Streaming algorithms are well suited to TVGs as their memory requirements do not depend on the duration of the trace but only on the number of vertices [1]. Memory is indeed a limited resource, as reachability graphs can become fully connected cliques and data structures such as adjacency matrices cannot therefore be considered sparse. Furthermore, in our case, a streaming algorithm facilitates a parallel implementation of the \oplus operator as one process can read the input families, duplicate their events, and dispatch these to various workers, each calculating one composition operation.

Algorithm 1 reads its input from τ pairs of event streams $\{\mathcal{L}_{d+k}, \mathcal{L}_{m-k}\}$ and writes $\mathcal{L}_{d+m} = \bigcup_{0 \leq k < \tau} \mathcal{L}_{d+k} \odot \mathcal{L}_{m-k}$ to its output stream. It makes use of an arc counter that tracks how many of the conditions in the definitions of $\mathcal{L}_{d+k} \odot \mathcal{L}_{m-k}$ (see Definition 8) are verified by each arc. When a previously down link fulfills one of these conditions, its counter is initialized to 1 and an UP event is written to the output stream (Line 5). When this counter goes to 0 the arc is removed from \mathcal{L}_{d+m} and a DOWN event is written to the output stream (Line 7). Note that an arc may be brought up and down at the same time i if its counter goes to 0 right after it appears (i.e., an ephemeral arc). In this case, the arc triggers both lines 5 and 7.

In more detail, this algorithm directly maps to the three conditions in the definition of the approximate composition \odot (Definition 8):

Condition 1. $(u, v) \in \mathcal{R}_{d+k}(i+1)$. This condition is handled on line 3, *before* making a decision on bringing arcs up or down at the *previous* epoch. Indeed, the output stream at time i can only be written to after reading the input streams up to time $i+1$.

Condition 2. $(u, v) \in \mathcal{R}_{m-k}(i+d+k)$. This condition means that down events in \mathcal{L}_{m-k} must be delayed for one epoch before lowering an arc's counter. This accounts for the local *delayed_k* lists that are processed on lines 16 to 19.

Condition 3. $\exists w, (u, w) \in \mathcal{L}_{d+k}(t)$ and $(w, v) \in \mathcal{L}_{m-k}(t+d)$. This condition is checked by maintaining for each k two adjacency lists adj_{d+k} and adj_{m-k} and checking upon

ALGORITHM 1: Lower bound composition

Require: $\{\mathcal{L}_{d+k}, \mathcal{L}_{m-k}\}$ // τ pairs of input streams
Ensure: \mathcal{L}_{d+m} // the output stream
Local: $\{adj_{d+k}, adj_{m-k}\}$, // τ pairs of adjacency matrices
Local: $counter[]$ // arc counter
Local: $delayed_k[]$ // τ lists of delayed down events

```

1 for  $i \leftarrow 0$  to  $T$  do
2   for  $k \leftarrow 0$  to  $\tau - 1$  do
3     for all the  $\{a, UP\} \in \mathcal{L}_{d+k}[i]$  do  $counter[a]++$ 
4     for all the new arcs  $a$  do
5       append  $\{a, UP\}$  to  $\mathcal{L}_{d+m}[i - 1]$ 
6     for all the arcs  $a$  s.t.  $counter[a] = 0$  do
7       append  $\{a, DOWN\}$  to  $\mathcal{L}_{d+m}[i - 1]$ 
8     for  $k \leftarrow 0$  to  $\tau - 1$  do
9       for all the  $\{(u, v), UP\} \in \mathcal{L}_{d+k}[i]$  do
10        add  $(u, v)$  to  $adj_{d+k}$ 
11        for all the  $(v, w) \in adj_{m-k}$  do  $counter[(u, w)]++$ 
12        for all the  $\{(u, v), UP\} \in \mathcal{L}_{m-k}[i + d + k]$  do
13           $counter[(u, v)]++$ 
14          add  $(u, v)$  to  $adj_{m-k}$ 
15          for all the  $(w, u) \in adj_{d+k}$  do  $counter[(w, v)]++$ 
16        for all the  $a \in delayed_k$  do  $counter[k]--$ 
17         $delayed_k \leftarrow \emptyset$ 
18        for all the  $\{(u, v), DOWN\} \in \mathcal{L}_{m-k}[i + d + k]$  do
19          add  $(u, v)$  to  $delayed_k$ 
20          remove  $(u, v)$  from  $adj_{m-k}$ 
21          for all the  $(w, u) \in adj_{d+k}$  do  $counter[(w, v)]--$ 
22        for all the  $\{(u, v), DOWN\} \in \mathcal{L}_{d+k}[i]$  do
23           $counter[(u, v)]--$ 
24          remove  $(u, v)$  from  $adj_{d+k}$ 
25          for all the  $(v, w) \in adj_{m-k}$  do  $counter[(u, w)]--$ 

```

UP/DOWN events whether the end of an arc in adj_{d+k} corresponds to the origin of an arc in adj_{m-k} (lines 10-11, 14-15, 20-21, and 24-25).

Algorithm 1 is then used as a building block to implement the \oplus addition of lower bound families. Starting from \mathcal{L}_τ 's family, for any $n \in \mathbb{N}^*$, $\mathcal{L}_{n\tau}$ is obtained in $\log(n)$ applications of the \oplus operation using a binary exponentiation process.

An implementation of Algorithm 1, fully integrated into a binary exponentiation algorithm over lower bound families, is available as a part of our dynamic trace library (DiTL [28]). This package also contains the code for transforming a lower bound TVG into an upper bound TVG, as well as the time-varying dominating set computation used in Section 5.

4.3 Complexity analysis

The worst-case memory requirements for this one-pass streaming algorithm are straightforward. Its local memory contains a non-sparse arc counter that requires up to $O(N^2)$ space. Furthermore, for each $0 \leq k < \tau$, it maintains two non-sparse adjacency matrices and an arc event list thereby requiring $O(N^2)$ space. Adding everything together yields a worst-case space complexity of $O(\tau N^2)$ that is *independent* of the duration of the trace.

Before examining the worst-case time complexity, a word must be said about the implementation of the adjacency matrices (the arc counter is backed by an adjacency matrix). In our implementation, these are backed by per-vertex hash tables. Insertion and removal are therefore constant time operations but this approach may not scale to TVGs with much greater number of vertices than those considered in

Table 2: Dataset characteristics. N is the number of vertices, M the number of edge UP/DOWN events, T the duration, P the beaconing period, and η the time resolution.

Name	N	M	T	P	η
Rollernet	62	99k	3h	15s	1s
Stanford	782	704k	8h	20s	20s
Random Waypoint	50	46k	8h	1s	1s
Community	50	824k	8h	1s	1s

this paper. In this analysis we will consider that insertion and removal cost $O(\log(N))$ (e.g., by using binary trees).

At each epoch, arcs are brought up and/or down. For each $0 \leq k < \tau$, we note M_{d+k} and M_{m-k} the number of UP and DOWN events in the entire TVGs \mathcal{L}_{d+k} and \mathcal{L}_{m-k} , respectively. Let $M = \max_k \{\max\{M_{d+k}, M_{m-k}\}\}$ be the maximum number of events in all the involved TVGs. In the worst case scenario where all possible arcs are updated at each epoch, we have $M = O(TN^2)$ events, where T is the number of epochs. Each processed event costs a modification of the arc counter ($O(\log(N))$), an adjacency matrix ($O(\log(N))$), and up to N counters for the composed paths ($O(N \log(N))$). All of these must be performed for the τ pairs of input TRGs. Adding everything together yields a worst-case time complexity of $O(\tau MN \log(N))$, or $O(\tau TN^3 \log(N))$ in terms of N exclusively. In practice however $M \ll TN^2$, so the first formulation is more accurate.

5. APPLICATIONS

Having formalized reachability graphs and detailed a method for efficiently computing upper and lower approximations of them in the previous sections, we now study the reachability graphs of several synthetic and real-life traces for a variety of edge traversal times τ and maximum delays δ . In particular, we show how the properties of these reachability graphs place bounds on communication capabilities and highlight the asymmetric nature of dynamic networks.

5.1 Datasets and metrics

While we have calculated reachability graphs on many publicly available datasets, in this paper we present results based on two real-life contact traces selected for their short beaconing periods:

Stanford [23]. As part of an epidemiology study, this trace captures face-to-face contacts among all students, teachers, and staff in a US high school between 7 a.m. and 4 p.m. The 782 ZigBee motes (TelosB Crossbow) sent beacons every 20 seconds (sending times are synchronized in the published trace).

Rollernet [25]. Opportunistic sighting of Bluetooth devices by groups of rollerbladers carrying Intel iMotes during a roller tour. The 62 iMotes performed neighborhood scans every 15 seconds.

For comparison purposes, we also provide results based on two well understood synthetic mobility models:

Random-Waypoint [18]. We simulated 50 nodes with speeds between 3 and 7 m/s from the stationary state in a $1,000 \times 500\text{m}^2$ rectangle sending beacons every second with a 20 m transmission range.

Table 3: Metrics on η -regular reachability graphs of duration T . \mathcal{X}_k is the value of \mathcal{X} during the k^{th} epoch, i.e., $\mathcal{X}_k = \mathcal{X}((k + \frac{1}{2})\eta)$.

Metric name	Definition
Avg. dominating set size	$\frac{\eta}{TN} \sum_{k=0}^{T/\eta} \ \mathcal{D}_k\ $
Avg. density	$\frac{\eta}{TN(N-1)} \sum_{k=0}^{T/\eta} \ \mathcal{R}_k\ $
Avg. asymmetry	$\frac{\eta}{T} \sum_{k=0}^{T/\eta} \frac{\ \text{asym}(\mathcal{R}_k)\ }{\ \text{asym}(\mathcal{R}_k)\ + \ \text{sym}(\mathcal{R}_k)\ /2}$

Community Model [19]. We simulated 50 nodes using the same parameters as above with 8 communities. In the community model, nodes with stronger social ties are more likely to be in geographic proximity.

The characteristics of these four η -regular traces are summarized in Table 2.

For each reachability graph, we compute its time-varying dominating set (TVDS), i.e., a mutable set such that, at all times, there exists an incoming arc to any vertex in the TVG from a member of the TVDS. Thus, at any given time t , the TVDS is a traditional dominating set on the directed graph at time t . We note this TVDS \mathcal{D}_δ if it is derived from a reachability graph \mathcal{R}_δ .

Even on static graphs, calculating a minimal dominating set is a classic NP-complete problem. Here we adapt the well-known greedy algorithm for choosing multipoint relays for broadcasting in a wireless network to the time-varying context [17]. The size of the dominating set calculated by this algorithm is within a factor $\log(N)$ of the optimal, where N is the number of nodes. Our TVDS calculation algorithm reacts to arc UP/DOWN events as follows.

1. If, after the arc event is processed, the previous dominating set is no longer a dominating set, then proceed to step 2. Otherwise do nothing.
2. Iteratively build the new dominating set, starting from an empty set, by greedily adding the node with the greatest outgoing degree to nodes not yet covered by the new dominating set. If two nodes have equal outgoing degrees, pick the one that is in the previous dominating set.

The time-varying dominating set thus calculated attempts to be both reasonably stable over time and close the optimal.

The rest of this section focuses on the following metrics. Their formulas and average values are found in Table 3.

- **Density.** The ratio of the number of arcs in a reachability graph at time t over the total number of possible arcs.
- **Dominating set size.** The normalized number of vertices in the dominating set at time t .
- **Asymmetry.** Here we no longer consider directed arcs but undirected pairs of vertices. At time t , the asymmetry is the ratio of asymmetric pairs among pairs that have at least one edge between them.

5.2 Algorithm accuracy

The approximation algorithm detailed in Section 4 is *extremely accurate*. Indeed, thanks to Theorem 3 and Definition 8, the approximation calculates the *exact* value for all




Figure 6: The upper and lower bounds give nearly identical values of density and dominating set size. Example taken from the Rollernet trace with $\tau=5\text{s}$ and $\delta=1\text{min}$.

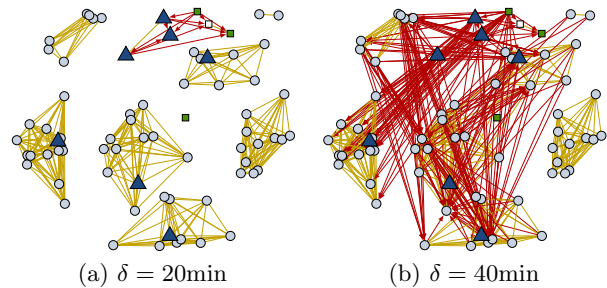
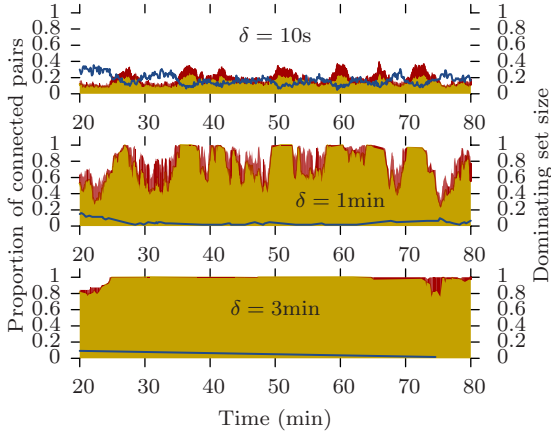


Figure 7: Subset of Stanford's 20 and 40-minute reachability graphs during morning classes ($\tau = 20\text{s}$). Dark blue triangles are teachers; circles are students. The dark red arrows represent asymmetric arcs. The classroom structure is clearly visible.


times $t \in \eta\mathbb{N}$. Furthermore, the upper and lower bounds are nearly identical during epochs ($t \notin \eta\mathbb{N}$). Fig. 6 plots the density over time for both the upper and lower bounds of the Rollernet reachability graph for $\tau = 5\text{s}$ and $\delta = 1\text{min}$. The plot has been zoomed in to show only a small span of the y -axis for 30 seconds. The upper and lower bounds on density are nearly equal at all times. So are the values of the dominating set size computed from the upper and lower approximations that only disagree in the circled area on Fig. 6. These observations hold for all the reachability graphs computed in this paper. Indeed, looking at the values of average density and average dominating set of the over 5,000 pairs of upper/lower bound TVGs calculated in this paper, the maximum disagreement is 8.10^{-3} for the former and 4.10^{-2} for the latter. In practice, the difference is smaller than the width of the line in plots. Therefore, in the rest of this paper, we only plot the values based on the lower approximation.

5.3 Revealing temporal structural properties

The dynamics of reachability graphs highlight temporal structural properties of the original connectivity trace that are not otherwise accessible. Fig. 7 shows snapshots of a subset of the reachability graph in the high school network captured in the Stanford dataset. When the delay is 20 minutes, the classroom structure of the trace is clearly visible. In the original contact trace, edges within a classroom are unstable and lead to merges and splits of small connected



(a) Rollernet ($\tau = 5s$)



(b) Stanford ($\tau = 1s$)

Figure 9: Proportion of connected pairs of vertices over time for two real-life datasets. Connected pairs are divided into symmetric pairs (yellow) and asymmetric pairs (dark red).




Figure 8: Snapshots of rollernet’s 10 and 60-seconds reachability graphs after 20 minutes ($\tau = 5s$). The rollerblading tour is moving from left to right. The dark red arrows represent asymmetric arcs. The asymmetry is caused by the acceleration phase in the accordion phenomenon [25].

components. In the reachability graph, a classroom is a stable complete subgraph with one teacher. A group of interacting teachers are also visible at the top of Fig. 7a. When some teachers and students later change classrooms, they create strong asymmetries in higher-delay reachability graphs such as the 40-minute one depicted in Fig. 7b.

Fig. 8 shows snapshots of the reachability graph captured in the Rollernet dataset. These correspond to an acceleration phase, where the head of the tour (right on Fig. 8) pulls ahead of the rest of the rollerbladers. Due to the accordion phenomenon, the tail of the tour does not react immediately [25]. This completely prevents short delay communications between these two groups ($\delta = 10s$ on Fig. 8a). With a longer delay (e.g., $\delta = 60s$ on Fig. 8b), *backward* communications towards the rear of the tour become possible (in particular through organizers who stop on the side of the road and let the tour pass them), thereby creating strong asymmetry in the reachability graph.

A more systematic study of dynamic properties over time is shown on Fig. 9. It plots both the proportion of connected pairs of vertices (left axis) and the size of the dominating set (right axis). The pairs of vertices are further divided into symmetric pairs (yellow) and asymmetric pairs (red). At a

given time, if the red histogram reaches 1 then a journey exists in at least one direction between all pairs of nodes. If the yellow histogram reaches 1, then a journey exist in both direction between all pairs of nodes.

For a 1-minute delay, the Rollernet reachability graph alternates, sometimes very rapidly, between fully connected states and highly asymmetric partially connected states (Fig. 9a). Therefore, any opportunistic communication system aiming for latencies under a minute will be strongly impacted by the accordion phenomenon. However, if a communication system can tolerate up to three minute delays, then it should be possible to smooth out the dynamic mobility. Indeed, the 3-minute reachability graph is fully connected for the entire duration of the trace, and the size of its dominating set is almost always equal to one.

The Stanford trace alternates static phases in classroom (the valleys on Fig. 9b) and dynamic phases moving between classrooms or around the food court (the peaks on Fig. 9b). The progressive “shrinking” of the valleys illustrates how reachability graphs “grow backwards” with increasing delays. Indeed, if there exist a journey from a to b within 20 minutes at time t , then there exists a journey from a to b within 40 minutes at time t minus 20 minutes. Of course, after the increase in delays exceeds the width of the valleys, the reachability graph eventually reaches it maximum density. Note that in this case, an incompressible amount of asymmetry subsists throughout the trace.

Due to space constraints, we have only shown, in this section, results for specific values of τ but, as a general rule, the smaller the value of τ , the faster the TRG becomes a clique. When analysing a network with a specific application in mind, the value of τ can be set to a realistic value (e.g. message size over bit-rate), and the plots in this section give an immediate visual understanding of the communication possibilities. The next section examines quantitatively the importance the delay (δ) and the edge traversal time (τ) parameters.

5.4 Bounds on communication capabilities

Reachability graphs give straightforward bounds on communication capabilities. Indeed, the density at a given mo-

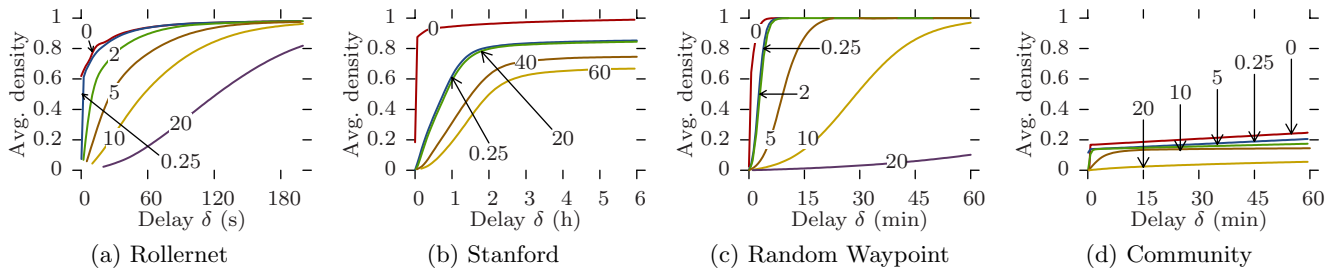


Figure 10: Average density vs. maximum delay δ for different edge traversal times τ (in seconds).

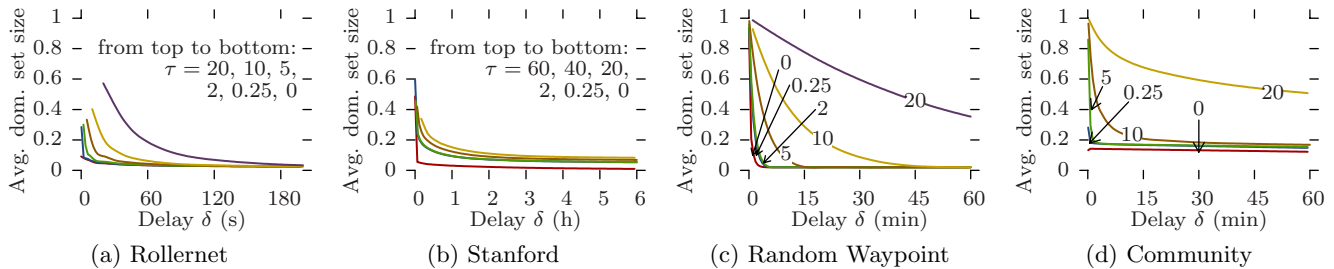


Figure 11: Average dominating set size vs. maximum delay δ for different values of τ (in seconds).

ment is exactly equal to the *maximum delivery ratio expectancy* of a perfect opportunistic routing protocol whose delay-tolerance is equal to δ and whose message size over bit-rate ratio is equal to τ . Furthermore, the size of the dominating set indicates the achievable *offload ratio* in a scenario where the opportunistic network is assisting an infrastructure (e.g., 3G) for disseminating content to all nodes in the network [12].

Given real-world system requirements, i.e., wireless bit-rate estimates, messages sizes, target delivery ratio, and delay-tolerance, reachability graphs provide an immediate answer to the following question: can an opportunistic network support this service? If not, can it effectively supplement an infrastructure in an offloading scheme?

Fig. 10 plots the average density against the delay tolerance for all datasets and for increasing edge traversal times, while Fig. 11 plots the average dominating set size. Rollernet and Random Waypoint share similar characteristics. Both are very sensitive to increasing edge traversal times. When these are close to one second, near 100% density is achievable with a couple of minutes of delay tolerance (Figs. 10a, 10c, 11a, and 11c). In this case, they can support pure opportunistic communications. However, when edge traversal times are longer, tight delay constraints are impossible unless as part of an offloading scenario (e.g., $\tau = 10$ s and $\delta = 60$ s for Rollernet, Fig 10a). For $\tau = 20$ s, Random Waypoint cannot even provide offloading for reasonable delays (Fig. 11c).

Similarly, Stanford and Community share similar features. Regardless of the delay-tolerance and message size, no pure opportunistic routing protocol can provide anything near 100% delivery ratio¹ (Figs. 10b and 10d). Despite this, they are both good offloading scenarios as the size of their dominating sets is consistently below 20% of the total number

of nodes, thereby offering potential offload ratios of around 80%. Indeed, in order to disseminate content to the entire network, pushing one copy per classroom in the Stanford case, or one copy per community in the Community case, plus copies to single nodes is an obvious strategy. However, for larger values of τ (e.g., 20), Community is no longer able to offload content with reasonable delays (Fig. 11d).

5.5 Asymmetry

Asymmetric communications are a fundamental aspect of opportunistic networks. Too strong focus on inter-contact times may lead to overlook asymmetry, but reachability graphs provide a natural way of studying and quantifying it.

Fig. 12 plots for all datasets the average asymmetry against the maximum delay for increasing values of the edge traversal time. The asymmetry for the \mathcal{R}_τ graph directly derived from the original contact trace is always 0 as it only contains symmetric arcs (see Section 3.4). Asymmetry in Random Waypoint follows a regular pattern: a bell-shaped curve of constant max value, whose width and center increase with τ (Fig. 12c). In this scenario, depending on the delay constraint, up to 80% on average of connected pairs of vertices can only communicate in one direction. This is due to nodes traveling long straight distances creating asymmetric reachability to nodes they meet from nodes they had met earlier. As the maximum delay increases, return journeys using different intermediate nodes appear, the reachability becomes complete, and the asymmetry returns to 0.

As previously, Rollernet and Random Waypoint show similar behavior. However, for Rollernet the maximum asymmetry also increases with τ (Fig. 12a). Indeed, as fewer edges are traversable, journeys become less likely between the front and the rear of the rollerblading tour. When such a journey is possible, it creates a longer-lasting asymmetry. In the Stanford trace, asymmetry follows the same increasing bell shape as in Rollernet, but never returns to 0 (Fig. 12b). In

¹The $\tau = 0$ results for Stanford are an artifact of the 20-second resolution.

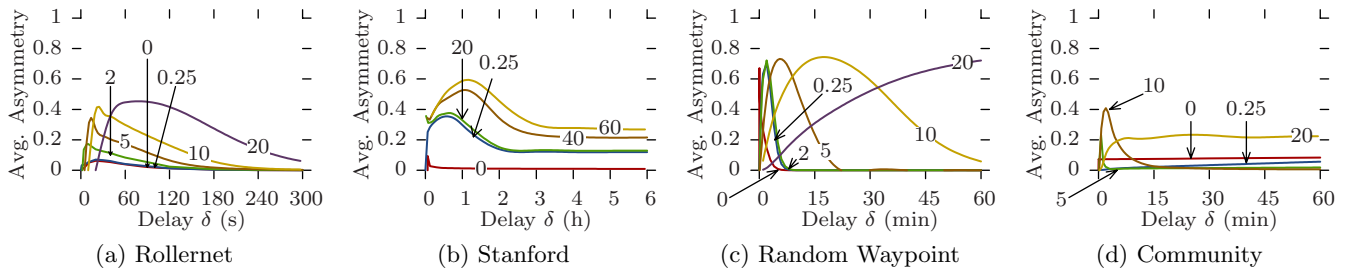


Figure 12: Average asymmetry vs. maximum delay δ for different values of τ (in seconds).

fact, this incompressible *minimum* asymmetry also increases with τ .

6. RELATED WORK

Adaptations of traditional static graph distance metrics and algorithms to time-varying graphs have yielded many different concepts. For example, Orda *et al.* propose a shortest path algorithm for TVGs based on different waiting policies (unrestricted, forbidden, and source waiting) [21]. Our work corresponds to the *unrestricted* policy, in which a message may wait for an unlimited amount of time anywhere along its path through the TVG. Bui Xuan *et al.* have proposed efficient algorithms for calculating shortest (in number of hops), fastest (in path traversal time), and foremost (i.e., earliest arrival) paths in TVGs [29]. All these algorithms are designed to compute the shortest paths to all destination from a source and a fixed starting time. Our algorithms, in contrast, computes the reachability graph by estimating the shortest paths for all possible starting time.

Several approaches to reachability in time-varying graphs exists. For strictly positive edge traversal times, a simple heuristic consists in dividing time into successive slots of length τ and keeping only edges that are persistently present during each slot [8]. This provides a good lower-bound approximation for small values of τ (i.e., less than η) whereas our approach can handle arbitrary edge traversal times. From a given starting time t , reachability among all pairs of nodes can be calculated by iterating over all edge UP/DOWN events [24]. This calculation can then be repeated for a sample of starting times [14]. This approach yields static reachability graphs for a discrete sequence of starting times, whereas the temporal reachability graphs defined in this paper calculate reachability in continuous time.

Chaintreau *et al.*, in their work on the diameter of opportunistic networks, calculate a *Last Departure / Earliest Arrival* structure for each pair of nodes [5]. This structure can tell for any pair of nodes (A, B) , at any time t , when is the earliest arrival for a message leaving A for B at time t . A reachability graph could be easily derived from these structures, but would unfortunately only cover the $\tau = 0$ case. Our approach is more general since it can also handle non-zero edge traversal times.

7. CONCLUSION AND FURTHER WORK

In this paper, we introduced the notion of *temporal reachability graphs*. Given an edge traversal time and a maximum journey delay, temporal reachability graphs capture the temporal connectivity of the time-varying graphs they derive

from. After formalizing the concept, we proved that regular reachability graphs, which encompass all experimental datasets, can be composed to compute reachability graphs of higher maximum delay with very high accuracy. Furthermore, we proposed a scalable highly-parallel streaming algorithm for their efficient computation. By applying this algorithm to synthetic and real-life contact traces, we showed how reachability graphs provide fresh new insights on temporal connectivity in time-varying graphs. In particular, they yield an immediate and intuitive picture of the communication capabilities and offloading potential of opportunistic networks.

This work on reachability graphs will be pursued in several directions. Firstly, our results could be extended to the more general case where the edge traversal time τ is not constant but may take values among multiples of a TVG's resolution η . Secondly, as seen in this paper, reachability graphs seem to reveal community structures that are not immediately apparent in the contact traces. They could therefore lead to new approaches for the difficult problem of community detection in time-varying graphs. Finally, much work remains to be done on the statistical analysis and modeling of reachability graphs. For example, what are the correlations between symmetric and asymmetric arcs from a given node? How do degree and inter-arc time distributions evolve with edge traversal time and delay? If reachability graphs turn out to be easier to model than their underlying time-varying graphs, could they be used as a first step for realistic synthetic connectivity models?

8. REFERENCES

- [1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. In *Proc. ACM STOC*, 1996.
- [2] S. Bhadra and A. Ferreira. Computing multicast trees in dynamic networks using evolving graphs. Technical Report RR-4531, INRIA, 2002.
- [3] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. In *Ad-hoc, Mobile, and Wireless Networks*, volume 6811 of *Lecture Notes in Computer Science*, pages 346–359. Springer Berlin / Heidelberg, 2011.
- [4] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, 6(6):606–620, 2007.
- [5] A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot.

- The diameter of opportunistic mobile networks. In *Proc. ACM CoNEXT*, 2007.
- [6] A. Clauset and N. Eagle. Persistence and Periodicity in a Dynamic Proximity Network. In *Proc. DIMACS Workshop on Computational Methods for Dynamic Interaction Network*, 2007.
- [7] V. Conan, J. Leguay, and T. Friedman. Characterizing pairwise inter-contact patterns in delay tolerant networks. In *Proc. ACM Autonomics*, 2007.
- [8] E. M. Daly and M. Haahr. Social network analysis for information flow in disconnected delay-tolerant manets. *IEEE Transactions on Mobile Computing*, 8(5):606–621, May 2009.
- [9] F. De Pellegrini, D. Miorandi, I. Carreras, and I. Chlamtac. A graph-based model for disconnected ad hoc networks. In *Proc. IEEE INFOCOM*, 2007.
- [10] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks. *Advances In Physics*, 51(1079-1187):4, 2002.
- [11] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. ACM SIGCOMM*, 2003.
- [12] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan. Mobile data offloading through opportunistic communications and social participation. *IEEE Transactions on Mobile Computing*, 11(5):821 – 834, 2011.
- [13] F. Harary and G. Gupta. Dynamic graph models. *Mathematical and Computer Modelling*, 25(7):79–88, 2005.
- [14] P. Holme. Network reachability of real-world contact sequences. *Phys. Rev. E*, 71:046119, Apr 2005.
- [15] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proc. ACM STOC*, 2000.
- [16] V. Kostakos. Temporal graphs. *Physica A Statistical Mechanics and its Applications*, 388:1007–1023, 2009.
- [17] A. Laouiti, A. Qayyum, and L. Viennot. Multipoint relaying: An efficient technique for flooding in mobile wireless networks. In *Proc. IEEE HICCS*, 2001.
- [18] J.-Y. Le Boudec and M. Vojnovic. Perfect simulation and stationarity of a class of mobility models. In *Proc. IEEE Infocom*, 2005.
- [19] M. Musolesi, S. Hailes, and C. Mascolo. Adaptive routing for intermittently connected mobile ad hoc networks. In *Proc. IEEE WoWMoM*, 2005.
- [20] M. Musolesi and C. Mascolo. A community based mobility model for ad hoc network research. In *Proc. ACM REALMAN*, 2006.
- [21] A. Orda and R. Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *J. ACM*, 37:607–625, 1990.
- [22] A. Passarella and M. Conti. Characterising aggregate inter-contact times in heterogeneous opportunistic networks. In *Proc. IFIP NETWORKING*, 2011.
- [23] M. Salathé, M. Kazandjieva, J. W. Lee, P. Levis, M. W. Feldman, and J. H. Jones. A high-resolution human contact network for infectious disease transmission. *PNAS*, 107(51):22020–22025, 2010.
- [24] J. Tang, M. Musolesi, C. Mascolo, and V. Latora. Characterising temporal distance and reachability in mobile and online social networks. *SIGCOMM Computer Communications Review*, 40(1):118–124, Jan. 2010.
- [25] P.-U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. Amorim, and J. Whitbeck. Density aware routing in highly dynamic dtns: the rollernet case. *IEEE Transactions on Mobile Computing*, 10(12):1755–1768, 2011.
- [26] A. Wagner. The yeast protein interaction network evolves rapidly and contains few redundant duplicate genes. *Molecular Biology and Evolution*, 18(7):1283–1292, 2001.
- [27] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.
- [28] J. Whitbeck. DynamIc Trace Library. <http://github.com/neush/ditl>.
- [29] B. B. 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.

APPENDIX

A. PROOFS

A.1 Proof of Theorem 1

LEMMA 1. Let \mathcal{R}_δ be a reachability graph of a TVG \mathcal{G} with edge traversal time τ . Let \mathcal{J} be a journey from u to v such that $\text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J}) < t + \delta + \tau$. Then $\exists 0 \leq \epsilon < 1$ such that $(u, v) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$.

PROOF. If $\text{arrival}(\mathcal{J}) \leq t + \delta$, then we set $\epsilon = 0$ and we have $(u, v) \in \mathcal{R}_\delta(t)$.

If $\text{arrival}(\mathcal{J}) > t + \delta$, let $0 \leq \epsilon < 1$ such that $\text{arrival}(\mathcal{J}) = t + \delta + \epsilon\tau$. In this case, $(u, v) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$. \square

LEMMA 2. Let \mathcal{R}_μ be a reachability graph of a TVG \mathcal{G} with edge traversal time τ . Let \mathcal{J} be a journey from u to v such that $\text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J}) \leq t + \mu$. Then $\mu \geq \tau$ and $\exists 0 \leq \epsilon < 1$ such that $(u, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \epsilon\tau)$.

PROOF. The temporal length of journey \mathcal{J} is $\delta_{\mathcal{J}} \leq \mu$. Since it is a valid journey from u to v , $\delta_{\mathcal{J}} \geq \tau$, hence $\mu \geq \tau$.

If $\text{departure}(\mathcal{J}) \geq t + \tau$, then $\delta_{\mathcal{J}} \leq \mu - \tau$, and $\forall 0 \leq \epsilon < 1$, $(u, v) \in \mathcal{R}_{\mu-\tau}(t + \tau) \subseteq \mathcal{R}_{\mu-\epsilon\tau}(t + \epsilon\tau)$.

If $\text{departure}(\mathcal{J}) < t + \tau$, then we set $0 \leq \epsilon < 1$ such that $\text{departure}(\mathcal{J}) = t + \epsilon\tau$. Then, since $\text{arrival}(\mathcal{J}) \leq t + \mu$, $\delta_{\mathcal{J}} \leq \mu - \epsilon\tau$ and $(u, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \epsilon\tau)$. \square

We can now prove Theorem 1.

PROOF. (Theorem 1) First, let us show that if an arc fits one of the decompositions then it belongs to $\mathcal{R}_{\delta+\mu}$. For all times t , let (u, v) be an arc in $(\mathcal{R}_{\delta+\epsilon\tau} \otimes \mathcal{R}_{\mu-\epsilon\tau})(t)$ ($0 \leq \epsilon < 1$). By definition, one of the following conditions holds.

1. $(u, v) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$. Therefore $(u, v) \in \mathcal{R}_{\delta+\mu}(t)$ (Proposition 1).
2. $(u, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \delta + \epsilon\tau)$. In this case, there exists a journey \mathcal{J} in \mathcal{G} such that $\text{departure}(\mathcal{J}) \geq t + \delta + \epsilon\tau \geq t$ and $\text{arrival}(\mathcal{J}) \leq t + \delta + \epsilon\tau + \mu - \epsilon\tau \leq t + \delta + \mu$. Therefore, by definition, $(u, v) \in \mathcal{R}_{\delta+\mu}(t)$.
3. $\exists w \in V$, s.t. $(u, w) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$ and $(w, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \delta + \epsilon\tau)$. Let \mathcal{J} and \mathcal{J}' be the respective journeys from u to w and from w to v . $\mathcal{J} \cup \mathcal{J}'$ is a journey from u to v such that $\text{departure}(\mathcal{J} \cup \mathcal{J}') = \text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J} \cup \mathcal{J}') = \text{arrival}(\mathcal{J}') \leq t + \delta + \epsilon\tau + \mu - \epsilon\tau$, i.e. $\text{arrival}(\mathcal{J} \cup \mathcal{J}') \leq t + \delta + \mu$. Therefore $(u, v) \in \mathcal{R}_{\delta+\mu}(t)$.

Conversely, let us show that for any time t and any arc $(u, v) \in \mathcal{R}_{\delta+\mu}(t)$, there exists $0 \leq \epsilon < 1$, such that $(u, v) \in (\mathcal{R}_{\delta+\epsilon\tau} \otimes \mathcal{R}_{\mu-\epsilon\tau})(t)$. Let $\mathcal{J} = \{(e_1, t_1), \dots, (e_k, t_k)\}$ be the journey from u to v such that $\text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J}) \leq t + \delta + \mu$.

Let $i = \max\{1 \leq j \leq k \mid t_j < t + \delta\}$. There are three possible situations.

1. i is not defined, i.e., $\text{departure}(\mathcal{J}) \geq t + \delta$. Since $\text{arrival}(\mathcal{J}) \leq t + \delta + \mu$, $\exists 0 \leq \epsilon < 1$, such that $(u, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \delta + \epsilon\tau)$ (Lemma 2).
2. $t_i = t_k$. Then \mathcal{J} is a path from u to v such that $\text{departure}(\mathcal{J}) \geq t$ and $\text{arrival}(\mathcal{J}) < t + \delta + \tau$. Therefore $\exists 0 \leq \epsilon < 1$, such that $(u, v) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$ (Lemma 1).
3. $t_1 \leq t_i < t_{i+1} \leq t_k$. In this case, we can divide the journey \mathcal{J} from u to v at time t into the sub-journeys $\mathcal{J}_1 = \{(e_1, t_1), \dots, (e_i, t_i)\}$ and $\mathcal{J}_2 = \{(e_{i+1}, t_{i+1}), \dots, (e_k, t_k)\}$. We set $w = \text{to}(e_i) = \text{from}(e_{i+1})$. We have $\text{arrival}(\mathcal{J}_1) =$

$t_i + \tau < t + \delta + \tau$. Therefore, $\exists 0 \leq \epsilon < 1$ such that $(u, w) \in \mathcal{R}_{\delta+\epsilon\tau}(t)$ (Lemma 1). Since \mathcal{J} is a valid journey in \mathcal{G} , $t_{i+1} \geq t_i + \tau \geq t + \delta + \epsilon\tau$. Therefore $\text{departure}(\mathcal{J}_2) \geq t + \delta + \epsilon\tau$, and $(w, v) \in \mathcal{R}_{\mu-\epsilon\tau}(t + \delta + \epsilon\tau)$.

In all three cases, $(u, v) \in (\mathcal{R}_{\delta+\epsilon\tau} \otimes \mathcal{R}_{\mu-\epsilon\tau})(t)$. \square

A.2 Proof of Theorem 2

LEMMA 3 (ONE-HOP JOURNEYS). Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau \in \eta\mathbb{N}^*$. For $k \in \mathbb{N}$, if a one-hop journey $\{(e, t)\}$ exists in \mathcal{G} such that $k\eta < t < (k+1)\eta$, then all one-hop journeys $\{(e, t)\}$ with $k\eta \leq t \leq (k+1)\eta$ also exist in \mathcal{G} .

PROOF. Let $\tau = n\eta$ with $n \in \mathbb{N}^*$. If $\mathcal{J} = \{(e, t)\}$ is a valid journey in t , then by definition for all $t \leq \zeta < t + \tau$, $e \in \mathcal{G}(\zeta)$. Since $k\eta < t < (k+1)\eta$ and \mathcal{G} is η -regular, then for all $k\eta \leq \zeta < (k+n)\eta$, $e \in \mathcal{G}(\zeta)$ (Definition 7). Furthermore, since $\text{arrival}(\mathcal{J}) = t + \tau > (k+n)\eta$, there exists $0 \leq \epsilon < 1$ such that $(k+n)\eta < (k+n+\epsilon)\eta < (k+n+1)\eta$. In this case, $e \in \mathcal{G}((k+n+\epsilon)\eta)$. Therefore, since \mathcal{G} is η -regular, for all $(k+n)\eta \leq \zeta < (k+n+1)\eta$, $e \in \mathcal{G}(\zeta)$ (Definition 7). Finally let $\hat{\mathcal{J}} = \{(e, \hat{t})\}$ be a one-hop journey in \mathcal{G} with $k\eta \leq \hat{t} < (k+1)\eta$. We have $\text{departure}(\hat{\mathcal{J}}) \geq k\eta$ and $\text{arrival}(\hat{\mathcal{J}}) \leq (k+1+n)\eta$. Since for all $k\eta \leq \zeta < (k+1+n)\eta$, $e \in \mathcal{G}(\zeta)$, $\hat{\mathcal{J}}$ is a valid journey in \mathcal{G} . \square

LEMMA 4 (EPOCH INCLUSION). Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau \in \eta\mathbb{N}^*$. For $\delta \in \eta\mathbb{N}$, let \mathcal{R}_δ be a reachability graph of \mathcal{G} . For $k \in \mathbb{N}$, let times t_a and t_b be such that $k\eta < t_a < (k+1)\eta$ and $k\eta \leq t_b \leq (k+1)\eta$. Then $\mathcal{R}_\delta(t_a) \subseteq \mathcal{R}_\delta(t_b)$.

PROOF. For $\delta = 0$, \mathcal{R}_δ is always empty and the lemma is trivially true. Hereafter we write $\delta = d\eta$, with $d \in \mathbb{N}^*$. If an arc (u, v) is in $\mathcal{R}_\delta(t_a)$, then there exists a journey $\mathcal{J} = \{(e_1, t_1), \dots, (e_k, t_k)\}$ from u to v such that $t_1 \geq t_a$ and $t_k + \tau \leq t_a + \delta$.

First we assume that $t_a < t_b$.

Let $i = \min\{j \mid t_j \geq t_b + (j-1)\tau\}$.

If i is not defined then $\forall j, t_a + (j-1)\tau \leq t_j < t_b + (j-1)\tau$.

We consider the journey $\hat{\mathcal{J}} = \{(e_j, \hat{t}_j)\}_{1 \leq j \leq k}$, such that $\hat{t}_j = t_b + (j-1)\tau$. By construction, the departure times of all of its one-hop journeys verify $\hat{t}_j - \hat{t}_{j-1} \geq \tau$. Furthermore, since $k\eta + (j-1)n\eta < t_j < \hat{t}_j < (k+1)\eta + (j-1)n\eta$, all one-hop journeys $\{(e_j, \hat{t}_j)\}$ are valid (Lemma 3). Therefore $\hat{\mathcal{J}}$ is a valid journey from u to v such that $\text{departure}(\hat{\mathcal{J}}) \geq t_b$ and $\text{arrival}(\hat{\mathcal{J}}) \leq t_b + \delta$, i.e., $(u, v) \in \mathcal{R}_\delta(t_b)$.

If i is defined then there are two possibilities. If $i = 1$ then $\text{departure}(\mathcal{J}) \geq t_b$ and (u, v) is trivially in $\mathcal{R}_\delta(t_b)$. If $i > 1$, then we can divide journey \mathcal{J} into two subjourneys $\mathcal{J}_1 = \{(e_j, t_j)\}_{0 \leq j < i}$ and $\mathcal{J}_2 = \{(e_j, t_j)\}_{i \leq j \leq k}$. As above, we can transform \mathcal{J}_1 into a valid journey $\hat{\mathcal{J}}_1 = \{(e_j, \hat{t}_j)\}_{0 \leq j < i}$ with $\hat{t}_j = t_b + (j-1)\tau$ such that $\text{departure}(\hat{\mathcal{J}}_1) \geq t_b$ and $\text{arrival}(\hat{\mathcal{J}}_1) \leq t_b + (i-1)\tau$. Since $\text{departure}(\mathcal{J}_2) \geq t_b + (i-1)\tau$, $\hat{\mathcal{J}}_1 \cup \mathcal{J}_2$ is a valid journey from u to v that leaves after t_b and arrives before $t_b + \delta$. Therefore $(u, v) \in \mathcal{R}_\delta(t_b)$.

Conversely, let us assume that $t_a > t_b$. Let $i = \max\{j \mid t_j + \tau \leq t_b + \delta - ($

$t_a + \delta + (k-j)\tau$. We consider the journey $\hat{\mathcal{J}} = \{(e_j, \hat{t}_j)\}_{1 \leq j \leq k}$, such that $\hat{t}_j = t_a + \delta + (k-j-1)\tau$. By construction,

the departure times of all of its one-hop journeys verify $\hat{t}_j - \hat{t}_{j-1} \geq \tau$. Furthermore, since $(k+d)\eta + (k-j-1)n\eta < \hat{t}_j < t_j < (k+d+1)\eta + (k-j-1)n\eta$, all one-hop journeys $\{(e_j, t_j)\}$ are valid (Lemma 3). Therefore $\hat{\mathcal{J}}$ is a valid journey from u to v such that $\text{departure}(\hat{\mathcal{J}}) \geq t_b$ and $\text{arrival}(\hat{\mathcal{J}}) \leq t_b + \delta$, i.e., $(u, v) \in \mathcal{R}_\delta(t_b)$.

If i is defined then there are two possibilities. If $i = k$ then $\text{arrival}(\mathcal{J}) \leq t_b + \delta$ and (u, v) is trivially in $\mathcal{R}_\delta(t_b)$. If $i < k$, then we can divide journey \mathcal{J} into two subjourneys $\mathcal{J}_1 = \{(e_j, t_j)\}_{0 \leq j \leq i}$ and $\mathcal{J}_2 = \{(e_j, t_j)\}_{i < j \leq k}$. As above, we can transform \mathcal{J}_2 into a valid journey $\hat{\mathcal{J}}_2 = \{(e_j, \hat{t}_j)\}_{i < j \leq k}$ with $\hat{t}_j = t_b + \delta + (k-j-1)\tau$ such that $\text{departure}(\hat{\mathcal{J}}_2) \geq t_b + \delta + (k-i)\tau$ and $\text{arrival}(\hat{\mathcal{J}}_2) \leq t_b + \delta$. Since $\text{arrival}(\mathcal{J}_1) \leq t_b + (k-i)\tau$, $\mathcal{J}_1 \cup \hat{\mathcal{J}}_2$ is a valid journey from u to v that leaves after t_b and arrives before $t_b + \delta$. Therefore $(u, v) \in \mathcal{R}_\delta(t_b)$. \square

We can now prove Theorem 2.

PROOF. For $\delta = 0$, \mathcal{R}_δ is always empty and the lemma is trivially true. Hereafter $\delta \in \eta\mathbb{N}^*$. For all $k \in \mathbb{N}$, and for all $k\eta \leq t_1 < (k+1)\eta$, $\mathcal{R}_\delta(t_1) \subseteq \mathcal{R}_\delta(k\eta)$ and $\mathcal{R}_\delta(t_1) \subseteq \mathcal{R}_\delta((k+1)\eta)$ (Lemma 4). For times t_2 such that $k\eta < t_2 < (k+1)\eta$, we have $\mathcal{R}_\delta(t_1) \subseteq \mathcal{R}_\delta(t_2)$ and $\mathcal{R}_\delta(t_2) \subseteq \mathcal{R}_\delta(t_1)$ (Lemma 4). Therefore $\mathcal{R}_\delta(t_1) = \mathcal{R}_\delta(t_2)$ and \mathcal{R}_δ is η -regular. \square

A.3 Proof of Theorem 3

LEMMA 5 (JOURNEYS IN REGULAR GRAPHS). *Let \mathcal{G} be an η -regular TVG whose edge traversal time is $\tau \in \eta\mathbb{N}^*$. Let \mathcal{J} be a journey in \mathcal{G} from u to v such that $\exists(a, b) \in \mathbb{N} \times \mathbb{N}$, $\text{departure}(\mathcal{J}) \geq a\eta$ and $\text{arrival}(\mathcal{J}) < (b+1)\eta$. Then there exists another journey $\hat{\mathcal{J}}$ from u to v that verifies $\text{departure}(\hat{\mathcal{J}}) \geq a\eta$ and $\text{arrival}(\hat{\mathcal{J}}) \leq b\eta$.*

PROOF. We write $\tau = n\eta$ with $n \in \mathbb{N}^*$. If $b \leq a + n$, then $\text{arrival}(\mathcal{J}) - \text{departure}(\mathcal{J}) < \tau$ which is impossible. Therefore $b > a + n$.

$\mathcal{J} = \{(e_1, t_1), \dots, (e_k, t_k)\}$ is a journey from u to v . For all i , we note c_i the integer such that $c_i\eta \leq t_i < (c_i+1)\eta$. Here $a = c_1$ and $b = c_k + n$. Indeed, we have $\text{arrival}(\mathcal{J}) = t_k + \tau < (c_k + n + 1)\eta$. Then we define $\hat{\mathcal{J}} = \{(e_1, c_1\eta), \dots, (e_k, c_k\eta)\}$.

Firstly, since \mathcal{J} is a valid journey, for all $t_i \leq t < t_i + \tau$, $e_i \in \mathcal{G}(t)$. Because \mathcal{G} is η -regular, this implies that for all $c_i\eta \leq t < c_i + \tau$, $e_i \in \mathcal{G}(t)$, and therefore all one-hop journeys $\{(e_i, c_i\eta)\}$ in $\hat{\mathcal{J}}$ are valid.

Secondly, for all i , $t_{i+1} - t_i \geq \tau$ leads to $(c_{i+1} - c_i)\eta + \eta > \tau$, i.e., $(c_{i+1} - c_i) + 1 > n$. Since these are all integers, we get $c_{i+1} - c_i \geq n$ and all the one-hop journeys in $\hat{\mathcal{J}}$ may be taken successively and $\hat{\mathcal{J}}$ is therefore a valid journey from u to v such that $\text{departure}(\hat{\mathcal{J}}) \geq a\eta$ and $\text{arrival}(\hat{\mathcal{J}}) \leq (c_k + n)\eta = b\eta$. \square

We can now prove Theorem 3.

PROOF. We write $\tau = n\eta$ with $n \in \mathbb{N}^*$. At time $t = a\eta$, let (u, v) be an arc in $\mathcal{R}_{(d+k)\eta} \otimes \mathcal{R}_{(m-k)\eta}$ with $0 \leq k < n$. By setting $\epsilon = k/n$, $(u, v) \in \mathcal{R}_{(d+m)\eta}(a\eta)$ with $0 \leq \epsilon < 1$ (Theorem 1).

Conversely, let us consider an $(u, v) \in \mathcal{R}_{(d+m)\eta}(a\eta)$. Since $m\eta \geq \tau$, according to Theorem 1, there exists $0 \leq \epsilon < 1$ such that one of the following holds (we note k the integer such that $k\eta \leq \epsilon n\eta < (k+1)\eta$).

- $(u, v) \in \mathcal{R}_{d\eta + \epsilon\tau}(a\eta)$. Therefore $(u, v) \in \mathcal{R}_{(d+k)\eta}(a\eta)$ (Lemma 5).
- $(u, v) \in \mathcal{R}_{m\eta - \epsilon\tau}((a+d)\eta + \epsilon\tau)$. With the inclusion rule from Proposition 1, (u, v) belongs to $\mathcal{R}_{(m-k)\eta}((a+d)\eta + \epsilon n\eta)$. Furthermore, since $\mathcal{R}_{(m-k)\eta}$ is η -regular (Theorem 2), $(u, v) \in \mathcal{R}_{(m-k)\eta}((a+d+k)\eta)$.
- $\exists w \in V, (u, w) \in \mathcal{R}_{d\eta + \epsilon\tau}(a\eta)$ and $(w, v) \in \mathcal{R}_{m\eta - \epsilon\tau}((a+d)\eta + \epsilon\tau)$. Following the same reasoning as for the first two conditions, we show that $\exists w \in V, (u, w) \in \mathcal{R}_{(d+k)\eta}(a\eta)$ and $(w, v) \in \mathcal{R}_{(m-k)\eta}((a+d+k)\eta)$.

In all three cases, there exists $0 \leq k < n$ such that $(u, v) \in \mathcal{R}_{(d+k)\eta} \otimes \mathcal{R}_{(m-k)\eta}$. \square

A.4 Proof of Proposition 3

PROOF. For $d < 2n$, by definition $\mathcal{L}_{d\eta} = \mathcal{R}_{d\eta}$. By induction, let us suppose the proposition true until a certain $l \geq 2n - 1$, and show that it also holds for $l + 1$. Let (u, v) be an arc in $\mathcal{L}_{(l+1)\eta}(t)$. We write a the integer such that $a\eta \leq t < (a+1)\eta$. If $t = a\eta$, then $\mathcal{L}_{(l+1)\eta}(t) = \mathcal{R}_{(l+1)\eta}(t)$ (Definition 8 and Theorem 3) and $(u, v) \in \mathcal{R}_{(l+1)\eta}(t)$. Hereafter we assume that $t > a\eta$. Since $l + 1 \geq 2n$, there exists $(d, m) \in N$ such that $n \leq d \leq l$, $n \leq m \leq l$, and $l + 1 = d + m$. In this case, by definition, there exists $0 \leq k < n$ such that one of the following holds.

- $(u, v) \in \mathcal{L}_{(d+k)\eta}((a+1)\eta)$. By setting $\epsilon = \frac{(a+1)\eta - t + k\eta}{n\eta}$, one can verify that $\exists 0 \leq \epsilon < 1$, such that $(u, v) \in \mathcal{R}_{d\eta + \epsilon\tau}(t)$.
- $(u, v) \in \mathcal{L}_{(m-k)\eta}((a+d+k)\eta)$. By setting $\epsilon = \frac{t - a\eta + k\eta}{n\eta}$, one can verify that $\exists 0 \leq \epsilon < 1$, such that $(u, v) \in \mathcal{R}_{m\eta - \epsilon\tau}(t + d\eta + \epsilon\tau)$.
- $\exists w \in V, (u, w) \in \mathcal{L}_{d\eta}(t)$ and $(w, v) \in \mathcal{L}_{m\eta}(t + d\eta)$. By induction, $(u, w) \in \mathcal{R}_{d\eta}(t)$ and $(w, v) \in \mathcal{R}_{m\eta}(t + d\eta)$.

With Theorem 1, in all three cases, $(u, v) \in \mathcal{R}_{(l+1)\eta}$. By induction, this shows that $\forall d \in \mathbb{N}, d, \mathcal{L}_{d\eta} \subseteq \mathcal{R}_{d\eta}$. \square