Time-Varying Graphs and Dynamic Networks*

Arnaud Casteigts[1], Paola Flocchini[1], Walter Quattrociocchi[2], and Nicola Santoro[3]

[1] University of Ottawa,
{casteig,flocchin}@site.uottawa.ca
[2] University of Siena,
walter.quattrociocchi@unisi.it
[3] Carleton University, Ottawa
nicola.santoro@scs.carleton.ca

Abstract. The past decade has seen intensive research efforts on highly dynamic wireless and mobile networks (variously called *delay-tolerant*, *disruptive-tolerant*, *challenged*, *opportunistic*, etc) whose essential feature is a possible absence of end-to-end communication routes at any instant. As part of these efforts, a number of important concepts have been identified, based on new meanings of distance and connectivity. The main contribution of this paper is to review and integrate the collection of these concepts, formalisms, and related results found in the literature into a unified coherent framework, called *TVG* (for *time-varying graphs*). Besides this definitional work, we connect the various assumptions through a hierarchy of classes of TVGs defined with respect to properties with algorithmic significance in distributed computing. One of these classes coincides with the family of dynamic graphs over which *population protocols* are defined. We examine the (strict) inclusion hierarchy among the classes. The paper also provides a quick review of recent stochastic models for dynamic networks that aim to enable analytical investigation of the dynamics.

Key words: Highly dynamic networks, delay-tolerant networks, challenged networks, time-varying graphs, evolving graphs, dynamic graphs.

1 Introduction

In the past few years, intensive research efforts have been devoted to the study of highly dynamic networks, whose topologies change as a function of time, and the rate of changes is too high to be reasonably modeled in terms of network faults or failures; in these systems *changes are not anomalies but rather integral part of the nature of the system*.

They include, but are not limited to, dynamic *mobile ad hoc networks* where the network's topology changes dramatically over time due to the movement of the network's nodes; *sensor networks* where links only exist when two neighbouring sensors are awake and have power; *vehicular networks* where the topology changes continuously as vehicles move. These highly dynamic infrastructure-less networks, variously called *delay-tolerant*, *disruptive-tolerant*, *challenged*, *opportunistic*, etc. (e.g., see [7, 10, 11, 19, 27, 29, 30]), have in common that the assumption of *connectivity* does not

* An extended version of this paper can be found in *arXiv:1012.0009*

necessarily hold, at least with the usual meaning of *contemporaneous end-to-end multi-hop paths* between any pair of nodes. The network may actually be disconnected at every time instant. Still, communication routes may be available over time and space, and make broadcast and routing and other computations feasible.

An extensive amount of research has been devoted, mostly by the engineering community but also by computer scientists, to the problems of operating and computing in such highly dynamical environments. As part of these efforts, a number of important concepts have been identified, often named, sometimes formally defined. In particular, most of the basic graph concepts were extended to a new *temporal* version, e.g, path and reachability [4, 20], distance [6], diameter [11], or connected components [5]. In several cases, differently named concepts identified by different researchers are actually one and the same concept. For example, the concept of *temporal distance*, formalized in [6], is the same as *reachability time* [18], *information latency* [21], and *temporal proximity* [22]; similarly, the concept of *journey* [6] has been coined *schedule-conforming path* [4], *time-respecting path* [18, 20], and *temporal path* [11, 28]. Hence, the concepts discovered in these investigations can be viewed as parts of the same conceptual universe; and the formalisms proposed so far to express them as fragments of a larger formal description of this universe.

As the notion of graph is the natural means for representing a standard network, the notion of *time-varying* graph is the natural means to represents these highly dynamic infrastructure-less networks. All the concepts and definitions advanced so far are based on or imply such a notion, as expressed even by the choices of names; e.g., Kempe et al. [20] talk of a *temporal network* (G, λ) where λ is a *time-labeling* of the edges, that associates to every edge a date corresponding to a punctual interaction; Leskovec et al. [24] talk of *graphs over time*; Ferreira [14] views the dynamic of the system in terms of a sequence of static graphs, called an *evolving graph*; Flocchini et al. [16] and Tang et al. [28] independently employ the term *time-varying graphs*; Kostakos uses the term *temporal graph* [22]; etc.[4]

The main contribution of this paper is to integrate the existing models, concepts, and results found in the literature into a unified framework that we call *TVG* (for *time-varying graphs*). This formalism, presented in Section 2, essentially consists of a set of compact and elegant notations and the possibility to switch between graph-centric and edge-centric perspectives on the dynamics. It is extended in Section 3, where we present the most central concepts identified by the research (e.g. journeys, temporal distance, connectivity over time and further concepts built on top of them). We identify in Section 4 several classes of dynamic networks defined with respects to basic properties on TVGs. Some of these classes have been extensively studied in the literature; e.g., one of them coincides with the family of dynamic graphs over which *population protocols* [1] are defined. We examine the (strict) inclusion hierarchy among the classes. As a given class typically corresponds to necessary or sufficient conditions for basic computations, the inclusion relationship implies the transferability of feasibility results (e.g., protocols) to an included class, and impossibility results (e.g., lower bounds) to an including class. Finally, Section 5 reviews recent efforts to study dynamic networks

[4] The more natural term *dynamic graph* is not often employed because it is commonly and extensively used in the context of *faulty* networks.

from a stochastic perspective, including modeling aspects (e.g. with *edge-markovian evolving graphs*), then we conclude with some remarks and open questions.

2 Time-Varying Graphs

Consider a set of entities V (or *nodes*), a set of relations E between these entities (*edges*), and an alphabet L accounting for any property such a relation could have (*label*); that is, $E \subseteq V \times V \times L$. The definition of L is domain-specific, and therefore left open – a label could represent for instance the intensity of relation in a social network, a type of carrier in a transportation network, or a particular medium in a communication network. For generality, we assume L to possibly contain multi-valued elements (e.g. $<$*satellite link; bandwidth of 4 MHz; encryption available;...*$>$). The set E enables multiple relations between a given pair of entities, as long as these relations have different properties, that is, for any $e_1 = (x_1, y_1, \lambda_1) \in E, e_2 = (x_2, y_2, \lambda_2) \in E$, $(x_1 = x_2 \wedge y_1 = y_2 \wedge \lambda_1 = \lambda_2) \implies e_1 = e_2$.

Because we address dynamical systems, the relations between entities are assumed to take place over a time span $\mathcal{T} \subseteq \mathbb{T}$ called the *lifetime* of the system. The temporal domain $\mathbb{T}$ is generally assumed to be $\mathbb{N}$ for discrete-time systems or $\mathbb{R}^+$ for continuous-time systems. The dynamics of the system can be subsequently described by a time-varying graph, or TVG, $\mathcal{G} = (V, E, \mathcal{T}, \rho, \zeta)$, where

- $\rho : E \times \mathcal{T} \to \{0, 1\}$, called *presence* function, indicates whether a given edge is available at a given time.
- $\zeta : E \times \mathcal{T} \to \mathbb{T}$, 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).

One may consider variants where the presence of nodes is also conditional upon time, by adding a *node presence function* $\psi : V \times \mathcal{T} \to \{0, 1\}$. We do not do it in the general case in this paper, for conciseness of the notations, and mention instead when this could be relevant. The TVG formalism can arguably describe a multitude of different scenarios, from transportation networks to communication networks, complex systems, or social networks. Two intuitive examples are shown on Figure 1.

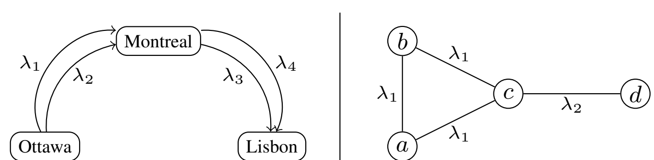

(a) Transportation network (b) Communication network

Fig. 1. Two examples of time-varying graphs, employed in very different contexts.

The meaning of what is an edge in these two examples varies drastically. In Figure 1(a), an edge from a node u to another node v represents the possibility for some agent to *move* from u to v. The edges in this example are assumed directed, and possibly multiple. The meaning of the labels λ_1 to λ_4 could be for instance "*bus*", "*car*", "*plane*", "*boat*", respectively. Except for the travel in car from Ottawa to Montreal – which could assumably be started anytime –, typical edges in this scenario are available on a *punctual* basis, *i.e.,* the presence function ρ for these edges returns 1 only at particular date(s) when the trip can be started. The latency function ζ may also vary from one edge to another, as well as for different availability dates of a same given edge (e.g. variable traffic on the road, depending on the departure time).

The second example on Figure 1(b) represents a history of connectivity between a set of moving nodes, where the possibilities of communication appear e.g. as a function of their respective distance. The two labels λ_1 and λ_2 may account here for different types of communication media, such as WiFi and Satellite, having various properties in terms of range, bandwidth, latency, or energy consumption. In this scenario, the edges are assumed to be undirected and there is no more than one edge between any two nodes. The meaning of an edge is also different here: an edge between two nodes means that any one (or both) of them can (attempt to) send a message to the other. A typical presence function for this type of edge returns 1 for some *intervals* of time, because the nodes are generally in range for a non-punctual period of time. Note that the effective delivery of a message sent at time t on an edge e could be subjected to further constraints regarding the latency function, such as the condition that $\rho(e)$ returns 1 for the whole interval $[t, t + \zeta(e, t))$.

These two examples are taken different on purpose; they illustrate the spectrum of *models* over which the TVG *formalism* can stretch. As observed, some contexts are intrisically simpler than others and call for restrictions (e.g. between any two nodes in the second example, there is at most one undirected edge). Further restrictions may be considered. For example the latency function could be decided constant over time ($\zeta : E \to \mathbb{T}$); over the edges ($\zeta : \mathcal{T} \to \mathbb{T}$); over both ($\zeta \in \mathbb{T}$), or simply ignored. In the latter case, a TVG could have its relations fully described by a graphical representation like that of Figure 2.

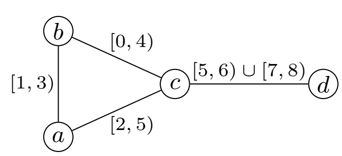

Fig. 2. A simple TVG. *The interval(s) on each edge e represents the periods of time when it is available, that is,* $\cup(t \in \mathcal{T} : \rho(e, t) = 1)$.

Note that a number of work on dynamic networks simply ignore ζ, or assume a discrete-time scenario where every time step implicitely corresponds to a constant ζ.

3 Definitions of TVG concepts

This section transposes and generalizes a number of dynamic network concepts into the framework of time-varying graphs. A majority of them emerged independently in various areas of scientific literature; some appeared more specifically; some others are original propositions.

3.1 The underlying graph G

Given a TVG $\mathcal{G} = (V, E, \mathcal{T}, \rho, \zeta)$, the graph $G = (V, E)$ is called *underlying* graph of $\mathcal{G}$. This static graph should be seen as a sort of *footprint* of $\mathcal{G}$, which flattens the time dimension and indicates only the pairs of nodes that have relations at some time in $\mathcal{T}$. It is a central concept that is used recurrently in the following.

In most studies and applications, G is assumed to be connected; in general, this is not necessarily the case. Let us stress that the connectivity of $G = (V, E)$ does not imply that $\mathcal{G}$ is connected at a given time instant; in fact, $\mathcal{G}$ could be disconnected at all times. The lack of relationship, with regards to connectivity, between $\mathcal{G}$ and its footprint G is even stronger: the fact that $G = (V, E)$ is connected does not even imply that $\mathcal{G}$ is "connected over time", as discussed in more details later.

3.2 Point of views

Depending on the problem under consideration, it may be convenient to look at the evolution of the system from the point of view of a given relation (edge) or from that of the global system (entire graph). We respectively qualify these views as *edge-centric* and *graph-centric*.

Edge-centric evolution From an edge standpoint, the notion of evolution comes down to a variation of availability and latency over time. We define the *available dates* of an edge e, noted $\mathcal{I}(e)$, as the union of all dates at which the edge is available, that is, $\mathcal{I}(e) = \{t \in \mathcal{T} : \rho(e,t) = 1\}$. When $\mathcal{I}(e)$ is expressed as a multi-interval of availability $\mathcal{I}(e) = [t_1, t_2) \cup [t_3, t_4)...$, where $t_i < t_{i+1}$, the sequence of dates $t_1, t_3, ...$ is called *appearance dates* of e, noted $App(e)$, and the sequence of dates $t_2, t_4, ...$ is called *disappearance dates* of e, noted $Dis(e)$. Finally, the sequence $t_1, t_2, t_3, ...$ is called *characteristic dates* of e, noted $\mathcal{S}_{\mathcal{T}}(e)$. In the following, we use the notation $\rho_{[t,t')}(e) = 1$ to indicate that $\forall t'' \in [t, t'), \rho(e, t'') = 1$.

Graph-centric evolution The sequence $\mathcal{S}_{\mathcal{T}}(\mathcal{G}) = sort(\cup\{\mathcal{S}_{\mathcal{T}}(e) : e \in E\})$, called *characteristic dates* of $\mathcal{G}$, corresponds to the sequence of dates when topological events (appearance/disappearance of an edge) occur in the system. Each topological event can be viewed as the transformation from one static graph to another. Hence, the evolution of the system can be described as a sequence of static graphs. More precisely, from a global viewpoint, the evolution of $\mathcal{G}$ is described as the sequence of graphs $\mathcal{S}_{\mathcal{G}} = G_1, G_2..$ where G_i corresponds to the static *snapshot* of $\mathcal{G}$ at time $t_i \in \mathcal{S}_{\mathcal{T}}(\mathcal{G})$; i.e., $e \in E_{G_i} \iff \rho_{[t_i,t_{i+1})}(e) = 1$. Note that, by definition, $G_i \neq G_{i+1}$.

In the case where the time is discrete, another possible global representation of evolution of $\mathcal{G}$ is by the sequence $\mathcal{S}_{\mathcal{G}} = G_1, G_2, \ldots$, where G_i corresponds to the static *snapshot* of $\mathcal{G}$ at time $t = i$. Note that, in this case, it is possible that $G_i = G_{i+1}$.

Observe that in both continuous and discrete cases, the underlying graph G (defined in Section 3.1) corresponds to the union of all G_i in $\mathcal{S}_{\mathcal{G}}$.

The idea of representing a dynamic graph as a sequence of static graphs, mentioned in conclusion of [17], was brought to life in [14] as a combinatorial model called *evolving graphs*. An evolving graph usually refers to either one of the two structures $(G, \mathcal{S}_\mathcal{G}, \mathcal{S}_\mathcal{T})$ or $(G, \mathcal{S}_\mathcal{G}, \mathbb{N})$, the latter used only when discrete-time is considered. Their initial version also included a latency function, which makes them a valid – graph-centric – representation of TVGs.

3.3 Subgraphs of a time-varying graph

Subgraphs of a TVG $\mathcal{G}$ can be defined in a classical manner, by restricting the set of vertices or edges of $\mathcal{G}$. More interesting is the possibility to define a *temporal subgraph* by restricting the lifetime $\mathcal{T}$ of $\mathcal{G}$, leading to the graph $\mathcal{G}' = (V, E', \mathcal{T}', \rho', \zeta')$ such that

- $\mathcal{T}' \subseteq \mathcal{T}$
- $E' = \{e \in E : \exists t \in \mathcal{T}' : \rho(e,t) = 1 \wedge t + \zeta(e,t) \in \mathcal{T}'\}$
- $\rho' : E' \times \mathcal{T}' \rightarrow \{0,1\}$ where $\rho'(e,t) = \rho(e,t)$
- $\zeta' : E' \times \mathcal{T}' \rightarrow \mathbb{T}$ where $\zeta'(e,t) = \zeta(e,t)$

In practice, we allow the notation $\mathcal{G}' = \mathcal{G}_{[t_a,t_b)}$ to denote the temporal subgraph of $\mathcal{G}$ restricted to $\mathcal{T}' = \mathcal{T} \cap [t_a, t_b)$, which includes the possible notations $\mathcal{G}_{[t_a,+\infty)}$ or $\mathcal{G}_{(-\infty,t_b)}$ regardless of whether $\mathcal{T}$ is open, semi-closed, or closed.

3.4 Journeys

A sequence of couples $\mathcal{J} = (e_1,t_1), (e_2,t_2)\ldots, (e_k,t_k)$, such that $e_1, e_2, ..., e_k$ is a walk in G is a *journey* in $\mathcal{G}$ if and only if $\rho(e_i,t_i) = 1$ and $t_{i+1} \geq t_i + \zeta(e_i,t_i)$ for all $i < k$. Additional constraints may be required in specific domains of application, such as the condition $\rho_{[t_i,t_i+\zeta(e_i,t_i))}(e_i) = 1$ in communication networks (the edge remains available until the message is delivered).

We denote by $departure(\mathcal{J})$, and $arrival(\mathcal{J})$, the starting date t_1 and the last date $t_k + \zeta(e_k,t_k)$ of a journey $\mathcal{J}$, respectively. Journeys can be thought of as *paths over time* from a source to a destination and therefore have both a *topological* length and a *temporal* length. The *topological length* of $\mathcal{J}$ is the number $|\mathcal{J}| = k$ of couples in $\mathcal{J}$ (i.e., the number of *hops*); its *temporal length* is its end-to-end duration: $arrival(\mathcal{J}) - departure(\mathcal{J})$.

Let us denote by $\mathcal{J}^*_\mathcal{G}$ the set of all possible journeys in a time-varying graph $\mathcal{G}$, and by $\mathcal{J}^*_{(u,v)} \subseteq \mathcal{J}^*_\mathcal{G}$ those journeys starting at node u and ending at node v. If a journey exists from a node u to a node v, that is, if $\mathcal{J}^*_{(u,v)} \neq \emptyset$, then we say that u can *reach* v, and allow the simplified notation $u \rightsquigarrow v$. Clearly, the existence of journey is not symmetrical: $u \rightsquigarrow v \nLeftrightarrow v \rightsquigarrow u$; this holds regardless of whether the edges are directed or not, because the time dimension creates its own level of direction. Given a node u, the set $\{v \in V : u \rightsquigarrow v\}$ is called the *horizon* of u.

3.5 Distance

As observed, the length of a journey can be measured both in terms of hops or time. This gives rise to two distinct definitions of *distance* in a time-varying graph $\mathcal{G}$:

- The *topological distance* from a node u to a node v at time t, noted $d_{u,t}(v)$, is defined as $Min\{|\mathcal{J}| : \mathcal{J} \in \mathcal{J}^*_{(u,v)} \wedge departure(\mathcal{J}) \geq t\}$. For a given date t, a journey whose departure is $t' \geq t$ and topological length is equal to $d_{u,t}(v)$ is qualified as *shortest* ;
- The *temporal distance* from u to v at time t, noted $\hat{d}_{u,t}(v)$ is defined as $Min\{arrival(\mathcal{J}) : \mathcal{J} \in \mathcal{J}^*_{(u,v)} \wedge departure(\mathcal{J}) \geq t\} - t$. Given a date t, a journey whose departure is $t' \geq t$ and arrival is $t + \hat{d}_{u,t}(v)$ is qualified as *foremost*. Finally, for any given date t, a journey whose departure is $\geq t$ and temporal length is $Min\{\hat{d}_{u,t'}(v) : t' \in \mathcal{T} \cap [t, +\infty)\}$ is qualified as *fastest*.

The problem of computing shortest, fastest, and foremost journeys in delay-tolerant networks was introduced in [6], and an algorithm for each of the three metrics was provided for the *centralized* version of the problem (with complete knowledge of $\mathcal{G}$).

A concept closely related to that of temporal distance is that of *temporal view*, introduced in [21] in the context of social network analysis. The temporal view (simply called view in [21]; we add the "temporal" adjective to avoid confusion with the concept of *view* in distributed computing) that a node v has of another node u at time t, denoted $\phi_{v,t}(u)$, is defined as the latest (i.e., largest) $t' \leq t$ at which a message received by time t at v could have been emitted at u; that is, in our formalism,

$$\phi_{v,t}(u) = \text{Max}\{departure(\mathcal{J}) : \mathcal{J} \in \mathcal{J}^*_{(u,v)} \wedge arrival(\mathcal{J}) \leq t\}.$$

The question of knowing whether all the nodes of a network could know their temporal views in real time was recently answered (affirmatively) in [10].

3.6 Other temporal concepts

The number of definitions built on top of temporal concepts could grow endlessly, and our aim is certainly not to enumerate all of them. Yet, here is a short list of additional concepts that we believe are general enough to be worthwhile mentioning.

The concept of eccentricity can be separated into a *topological eccentricity* and a *temporal eccentricity*, following the same mechanism as for the concept of distance. The temporal eccentricity of a node u at time t, $\hat{\varepsilon}_t(u)$, is defined as $max\{\hat{d}_{u,t}(v) : v \in V\}$, that is, the duration of the "longest" foremost journey from u to any other node. The concept of diameter can similarly be separated into those of *topological diameter* and *temporal diameter*, the latter being defined at time t as $max\{\hat{\varepsilon}_t(u) : u \in V\}$. These temporal versions of eccentricity and diameter were proposed in [6]The temporal diameter was further studied from a stochastic point of view by Chaintreau *et al.* in [11].

Clementi *et al.* introduced in [13] a concept of *dynamic expansion* – the dynamic counterpart of the concept of *node expansion* in static graphs – which accounts for the maximal speed of information propagation. Given a subset of nodes $V' \subseteq V$, and two dates $t_1, t_2 \in \mathcal{T}$, the dynamic expansion of V' from time t_1 to time t_2 is the size of the set $\{v \in V \setminus V' : \exists \mathcal{J}_{(u,v)} \in \mathcal{J}^*_{\mathcal{G}[t_1,t_2)} : u \in V'\}$, that is roughly speaking, the "collective" horizon of V' in $\mathcal{G}_{[t_1,t_2)}$.

The concept of journey was dissociated in [10] into *direct* and *indirect* journeys. A journey $\mathcal{J} = \{(e_1, t_1), (e_2, t_2) \ldots, (e_k, t_k)\}$ is said *direct* iff $\forall i, 1 \leq i < k$, $\rho(e_{i+1}, t_i + \zeta(e_i, t_i)) = 1$, that is, every *next* edge in $\mathcal{J}$ is directly available; it is said indirect otherwise. The knowledge of whether a journey is direct or indirect was directly exploited by the distributed algorithm in [10] to compute temporal views between nodes. Such parameter can also play a role in the context of delay-tolerant routing, indicating whether a store-carry-forward mechanism is required (for indirect journeys).

4 TVG Classes

This section discusses the impact of temporal properties on the feasibility and complexity of distributed problems, unifying existing works from the literature. In particular, we identify a hierarchy of *classes* of TVGs based on properties that are formulated using the concepts presented in the previous section. These class-defining properties, organized in an ascending order of assumptions, are important in that they imply necessary conditions and impossibility results for distributed computations. Let us start with the simplest class.

Class 1 $\exists u \in V : \forall v \in V : u \leadsto v$.

That is, at least one node can reach all the others. This condition is necessary, for example, for broadcast to be feasible from at least one node.

Class 2 $\exists u \in V : \forall v \in V : v \leadsto u$.

That is, at least one node can be reached by all the others. This condition is necessary to be able to compute a function whose input is spread over all the nodes, with at least one node capable of generating the output. Any algorithm for which a terminal state must be causally related to all the nodes initial states also falls in this category, such as the election of a leader in an anonymous network or the counting of the number of nodes by at least one node.

Class 3 (Connectivity over time) $\forall u, v \in V, u \leadsto v$.

That is, every node can reach all the others; in other words, the TVG is connected over time. By the same discussions as for Class 1 and Class 2, this condition is necessary to be able to broadcast from any node, to compute a function whose output is known by all the nodes, or to ensure that every node has a chance to be elected. These three basic classes were used e.g. in [8] to investigate how relations between TVGs properties and feasibility of algorithms could be *canonically* proven.

Class 4 (Round connectivity)

$$\forall u, v \in V, \exists \mathcal{J}_1 \in \mathcal{J}^*_{(u,v)}, \exists \mathcal{J}_2 \in \mathcal{J}^*_{(v,u)} : arrival(\mathcal{J}_1) \leq departure(\mathcal{J}_2).$$

That is, every node can reach all the others and be reached back *afterwards*. Such a condition may be required e.g. for adding explicit termination to broadcast, election, or counting algorithms.

The classes defined so far are in general relevant to the case when the lifetime is *finite* and a limited number of topological events is considered. When the lifetime is *infinite*, connectivity over time is generally assumed on a regular basis, and more elaborated assumptions can be considered.

Class 5 (Recurrent connectivity) $\forall u, v \in V, \forall t \in \mathcal{T}, \exists \mathcal{J} \in \mathcal{J}^*_{(u,v)} : departure(\mathcal{J}) > t.$

That is, at any point t in time, the temporal subgraph $\mathcal{G}_{[t,+\infty)}$ remains connected over time. This class is implicitely considered in most works on delay-tolerant networks. It indeed represents those DTNs where routing can always be achieved over time. It has been explicitly referred to as *eventually transportable dynamic networks* in [27].

As discussed in Section 3.1, the fact that the underlying graph $G = (V, E)$ is connected does not imply that $\mathcal{G}$ is connected over time – the ordering of topological events matters. Such a condition is however *necessary* to allow connectivity over time and thus to perform any type of global computation. Therefore, the following three classes assume that the underlying graph G is connected.

Class 6 (Recurrence of edges) $\forall e \in E, \forall t \in \mathcal{T}, \exists t' > t : \rho(e, t') = 1$ *and* G *is connected.*

That is, if an edge appears once, it appears infinitely often. Since the underlying graph G is connected, we have Class 6 $\subseteq$ Class 5. Indeed, if all the edges of a connected graph appear infinitely often, then there must exist, by transitivity, a journey between any pairs of nodes infinitely often.

In a context where connectivity is recurrently achieved, it becomes interesting to look at problems where more specific properties of the journeys are involved, e.g. the possibility to broadcast a piece of information in a shortest, foremost, or fastest manner (see Section 3.5 for definitions). Interestingly, these three declinations of the same problem have different requirements in terms of TVG properties. It is for example possible to broadcast in a foremost fashion in Class 6, whereas shortest and fastest broadcasts are not possible [9].

Shortest broadcast becomes however possible if the recurrence of edges is bounded in time, and the bound known to the nodes, a property characterizing the next class:

Class 7 (Time-bounded recurrence of edges) $\forall e \in E, \forall t \in \mathcal{T}, \exists t' \in [t, t+\Delta), \rho(e, t') = 1$*, for some* $\Delta \in \mathbb{T}$ *and* G *is connected.*

Some implications of this class include a temporal diameter that is bounded by $\Delta Diam(G)$, as well as the possibility for the nodes to wait a period of Δ to discover all their neighbors (if Δ is known). The feasibility of shortest broadcast follows naturally by using a Δ-rounded breadth-first strategy that minimizes the topological length of journeys.

A particular important type of bounded recurrency is the periodic case:

Class 8 (Periodicity of edges) $\forall e \in E, \forall t \in \mathcal{T}, \forall k \in \mathbb{N}, \rho(e, t) = \rho(e, t + kp)$*, for some* $p \in \mathbb{T}$ *and* G *is connected.*

The periodicity assumption holds in practice in many cases, including networks whose entities are mobile with periodic movements (satellites, guards tour, subways, or buses). The periodic assumption within a delay-tolerant network has been considered, among others, in the contexts of network exploration [15, 16] and routing [25]. Periodicity enables also the construction of foremost broadcast trees that can be re-used (modulo p in time) for subsequent broadcasts [10] (whereas the more general classes of recurrence requires the use of a different tree for every foremost broadcast).

More generally, the point in exploiting TVG properties is to rely on invariants that are generated by the dynamics (*e.g.* recurrent existence of journeys, periodic optimality of a broadcast tree, *etc.*). In some works, particular assumptions on the network dynamics are made to obtain invariants of a more classic nature. Below are some examples of classes, formulated using the graph-centric point of view of (discrete-time) evolving graphs, *i.e.*, where $\mathcal{G} = (G, \mathcal{S}_\mathcal{G}, \mathbb{N})$.

Class 9 (Constant connectivity) $\forall G_i \in \mathcal{S}_\mathcal{G}, G_i$ *is connected.*

Here, the dynamics of the network is not constrained as long as it remains connected in every time step. Such a class was used for example in [26] to enable progression hypotheses on the broadcast problem. Indeed, if the network is always connected, then at every time step there must exist an edge between an informed node and a non-informed node, which allows to upper-bound the broadcast time by $n = |V|$ time steps (worst case scenario).

Class 10 (T-interval connectivity) $\forall i \in \mathbb{N}, T \in \mathbb{N}, \exists G' \subseteq G : V_{G'} = V_G, G'$ *is connected, and* $\forall j \in [i, i+T-1), G' \subseteq G_j$.

This class is a particular case of constant connectivity in which a same spanning connected subgraph of the underlying graph G is available for any period of T consecutive time steps. It was introduced in [23] to study problems such as counting, token dissemination, and computation of functions whose input is spread over all the nodes (considering an adversary-based edge schedule). The authors shown that the computation of these problems could be sped up of a factor T compared to the 1-interval connected graphs, that is, graphs of Class 9.

Other classes of TVGs can be found in [27], based on intermediate properties between constant connectivity and connectivity over time. They include Class 11 and Class 12 below.

Class 11 (Eventual connectivity) $\forall i \in \mathbb{N}, \exists j \in \mathbb{N} : j \geq i$*,* G_j *is connected.*

In other words, there is always a future time step in which the network is instantly connected.

Class 12 (Eventual routability) $\forall u, v \in V, \forall i \in \mathbb{N}, \exists j \in \mathbb{N} : j \geq i$ *and a path from* u *to* v *exists in* G_j.

That is, for any two nodes, there is always a future time step in which a instant path exists between them. The difference with Class 11 is that the paths can appear at different time for different pairs of nodes. Classes 11 and 12 were used in [27] to represent networks where routing protocols for (connected) mobile ad hoc networks eventually work if they tolerate transient topological faults.

For all the classes discussed so far, the referenced investigations studied the impact that various TVG properties have on problems or algorithms. A reverse approach was considered by Angluin *et al.* in the field of *population protocols* [1]. Instead of studying the impact of various assumptions on given problems, they assumed a given assumption – that any pair of node interacts infinitely often –, and characterized all problems that can be solved in that context. This class is generally referred to as that of *complete graphs of interaction*.

Class 13 (Complete graph of interaction) *The underlying graph* $G=(V,E)$ *is complete, and* $\forall e \in E, \forall t \in \mathcal{T}, \exists t' > t : \rho(e,t')=1$.

From a time-varying graph perspective, this class is the specific subset of Class 6, in which the underlying graph G is complete. Various types of schedulers have been considered in the area of population protocols that add further fairness constraints on Class 13 (e.g. weak fairness, strong fairness, bounded, or k-bounded schedulers). Each of these could further be seen as a distinct subclass of Class 13.

An interesting aspect of unifying these properties within the same formalism is the possibility to see how they relate to one another, and to compare the associated solutions or algorithms. An insight for example can be gained by looking at the short classification shown in Figure 3, where basic relations of inclusion between the above classes are reported. These inclusion are *strict*: for each relation, the parent class contains some time-varying graphs that are not in the child class.

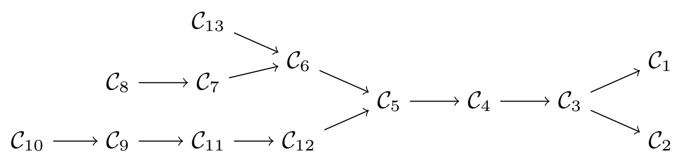

Fig. 3. Relations of inclusion between classes (*from specific to general*).

Clearly, one should try to solve a problem in the most general context possible. The right-most classes are so general that they offer little properties to be exploited by an algorithm, but some intermediate classes, such as Class 5, appear quite central in the hierarchy. This class indeed contains all the classes where significant work was done. A problem solved in this class would therefore apply to virtually all the contexts considered heretofor in the literature.

Such a classification may also be used to categorize problems themselves. As mentioned above, shortest broadcast is not generally achievable in Class 6, whereas foremost broadcast is. Similarly, it was shown in [9] that fastest broadcast is not feasible in Class 7, whereas shortest broadcast can be achieved with some knowledge. Since Class 7 $\subset$ Class 6, we have

$$foremostBcast \prec shortestBcast \prec fastestBcast$$

where $\prec$ is the partial order on these problems' topological requirements.

5 Non-Deterministic TVGs

Non-determinism in time-varying graphs can be introduced at several different levels. The most direct one is clearly that provided by *probabilistic time-varying graphs*, where the presence function $\rho : E \times \mathcal{T} \to [0,1]$ indicates the *probability* that a given edge is available at a given time. In a context of mobility, the probability distribution of ρ

is intrinsically related to the random *mobility pattern* defining the network. Popular example of random mobility models are the Random Waypoint and Random Direction models, where waypoints of consecutive movements are chosen uniformly at random.

Definitions of *random* TVG differ depending on whether the time is discrete or continuous. A *(discrete-time) random TVG* is one whose lifetime is an interval of $\mathbb{N}$ and whose *sequence of characteristic graphs* $\mathcal{S}_\mathcal{G} = G_1, G_2, ..$ is such that every G_i is a Erdös and Rényi random graph; that is, $\forall e \in V^2, \mathbb{P}[e \in E_{G_i}] = p$ for some p; this definition is introduced by Chaintreau *et al.* [11].

One particularity of discrete-time random TVGs is that the G_is are independent with respect to each other. While this definition allows purely random graphs, it does not capture some properties of real world networks, such as the fact that an edge may be more likely to be present in G_{i+1} if it is already present in G_i. This question is addressed by Clementi *et al.* [12] by introducing *Edge-Markovian Evolving Graphs*. These are discrete-time evolving graphs in which the presence of every edge follows an individual Markovian process. More precisely, the sequence of characteristic graph $\mathcal{S}_\mathcal{G} = G_1, G_2, ..$ is such that

$$\begin{cases} \mathbb{P}[e \in E_{G_{i+1}} | e \notin E_{G_i}] = p \\ \mathbb{P}[e \notin E_{G_{i+1}} | e \in E_{G_i}] = q \end{cases}$$

for some p and q called *birth rate* and *death rate*, respectively. The probability that a given edge remains absent or present from G_i to G_{i+1} is obtained by complement of p and q. The very idea of considering a *Markovian* Evolving Graph seems to have appeared in [2], in which the authors consider a particular case that is substantially equivalent to the discrete-time random TVG from [11]. Edge-Markovian EGs were used in [12], along with the concept of dynamic expansion (see Section 3.6) to address analytically some fundamental questions such as *does dynamics necessarily slow down a broadcast?* Or *can random node mobility be exploited to speed-up information spreading?* Baumann *et al.* extended this work in [3] by establishing tight bounds on the propagation time for any birth and death rates.

A *continuous-time random TVG* is one in which the appearance of every edge obeys a Poisson process, that is, $\forall e \in V^2, \forall t_i \in App(e), \mathbb{P}[t_{i+1} - t_i < d] = \lambda e^{\lambda d}$ for some λ; this definition is introduced by Chaintreau *et al.* in [11].[5]

Random time-varying graphs, both discrete- and continuous-time, were used in [11] to characterize phase transitions between no-connectivity and connectivity over time as a function of the number of nodes, a given time-window duration, and constraints on both the topological and temporal lengths of journeys.

6 Research Problems and Directions

The first most obvious research task is that of exploring the universe of dynamic networks using the formal tools provided by the TVG formalism. The long-term goal is that of providing a comprehensive map of this universe, identifying both the commonality and the natural differences between the various types of dynamical systems modeled

[5] It is interesting to note that the authors rely on a graph-centric point of view in discrete time and on an edge-centric point of view in continuous time. This trend seems to be general.

by TVG. Additionally, several, more specific research areas can be identified including the ones described below.

Distributed TVG algorithms design and analysis. The design and analysis of distributed algorithms and protocols for time-varying graphs is an open research area. In fact very few problems have been attacked so far: routing and broadcasting in delay-tolerant networks; broadcasting and exploration in opportunistic-mobility networks; new self-stabilization techniques; detection of emergence and resilience of communities, and viral marketing in social networks.

Design and optimization of TVG. If the interactions in a network can be planned – decided by a designer –, then a number of new interesting optimization problems arise with the design of time-varying graph. They may concern for example the minimization of the temporal diameter or the balancing of nodes eccentricities. Is a given setting optimal? How to prove it? What if the underlying graph can also be modified? *etc.* A whole field is opening that promises exciting research avenues.

Complexity Analysis. Analyzing the complexity of a distributed algorithm in a TVG – e.g. in number of messages – is not trivial, partly because contrarily to the static cases, the complexity of an algorithm in a dynamic network has a strong dependency, not only on the usual network parameters (number of nodes, edges, etc.), but also on the number of topological events taking place during its execution. In many of the algorithms we have encountered, the majority of messages is in fact directly triggered by topological events, e.g., in reaction to the local appearance or disappearance of an edge. The number of topological events therefore represents a new complexity parameter, whose impact on various problems remains to study.

References

1. D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
2. C. Avin, M. Koucky, and Z. Lotker. How to explore a fast-changing world. In *Proc. 35th Intl. Colloquium on Automata, Languages and Programming* (ICALP), pages 121–132, 2008.
3. H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. In Proc. 28th ACM *Symp. on Principles of Distributed Computing* (PODC), pages 260–269, 2009.
4. K.A. Berman. Vulnerability of scheduled networks and a generalization of Menger's Theorem. *Networks*, 28(3):125–134, 1996.
5. S. Bhadra and A. Ferreira. Complexity of connected components in evolving graphs and the computation of multicast trees in dynamic networks. In *Proc. 2nd Intl. Conf. on Ad Hoc Networks and Wireless (AdHoc-Now)*, pages 259–270, 2003.
6. B. Bui-Xuan, A. Ferreira, and A. Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *Intl. J. of Foundations of Comp. Science*, 14(2):267–285, April 2003.
7. J. Burgess, B. Gallagher, D. Jensen, and B.N. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networks. In *Proc. 25th IEEE Conference on Computer Communications (INFOCOM)*, pages 1–11, 2006.
8. A. Casteigts, S. Chaumette, and A. Ferreira. On the assumptions about network dynamics in distributed computing. *Arxiv preprint arXiv:1102.5529*, 2011. A preliminary version appeared in *SIROCCO'09*.
9. A. Casteigts, P. Flocchini, B. Mans, and N. Santoro. Deterministic computations in time-varying graphs: Broadcasting under unstructured mobility. In *Proc. 5th IFIP Conference on Theoretical Computer Science (TCS)*, pages 111–124, 2010.

10. A. Casteigts, P. Flocchini, B. Mans, and N. Santoro. Measuring temporal lags in delay-tolerant networks. In *Proc. 25th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2011.
11. A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot. The diameter of opportunistic mobile networks. *Communications Surveys & Tutorials*, 10(3):74–88, 2008.
12. A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time in edge-markovian dynamic graphs. In *Proc. 27th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 213–222, 2008.
13. A. Clementi and F. Pasquale. *Information Spreading in Dynamic Networks: An Analytical Approach*. In: S. Nikoletseas, and J. Rolim (Eds), *Theoretical Aspects of Distributed Computing in Sensor Networks*, Springer, 2010.
14. A. Ferreira. Building a reference combinatorial model for MANETs. *IEEE Network*, 18(5):24–29, 2004.
15. P. Flocchini, M. Kellett, P. Mason, and N. Santoro. Mapping an unfriendly subway system. In *Proc. 5th Intl. Conf. on Fun with Algorithms*, pages 190–201, 2010.
16. P. Flocchini, B. Mans, and N. Santoro. Exploration of periodically varying graphs. In *Proc. 20th Intl. Symp. on Algorithms and Computation (ISAAC)*, pages 534–543, 2009.
17. F. Harary and G. Gupta. Dynamic graph models. *Mathematical and Computer Modelling*, 25(7):79–88, 1997.
18. P. Holme. Network reachability of real-world contact sequences. *Physical Review E*, 71(4):46119, 2005.
19. S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proc. Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 145–158, 2004.
20. D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proc. 32nd ACM Symp. on Theory of Computing (STOC)*, page 513, 2000.
21. G. Kossinets, J. Kleinberg, and D. Watts. The structure of information pathways in a social communication network. In *Proc. 14th Intl. Conf. on Knowledge Discovery and Data Mining (KDD)*, pages 435–443, 2008.
22. V. Kostakos. Temporal graphs. *Physica A*, 388(6):1007–1023, 2009.
23. F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proc. 42nd ACM Symp. on Theory of Computing*, pages 513–522, 2010.
24. J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Trans. on Knowledge Discovery from Data*, 1(1), 2007.
25. C. Liu and J. Wu. Scalable routing in cyclic mobile networks. *IEEE Trans. Parallel Distrib. Syst.*, 20(9):1325–1338, 2009.
26. R. O'Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proc. Workshop on Foundations of Mobile Computing (DIALM-POMC)*, pages 104–110, 2005.
27. R. Ramanathan, P. Basu, and R. Krishnan. Towards a formalism for routing in challenged networks. In *Proc. 2nd ACM Workshop on Challenged Networks (CHANTS)*, pages 3–10, 2007.
28. J. Tang, M. Musolesi, C. Mascolo, and V. Latora. Characterising temporal distance and reachability in mobile and online social networks. *ACM Computer Communication Review*, 40(1):118–124, 2010.
29. X. Zhang, J. Kurose, B.N. Levine, D. Towsley, and H. Zhang. Study of a bus-based disruption-tolerant network: mobility modeling and impact on routing. In *Proc. 13th ACM Int. Conf. on Mobile Computing and Networking*, pages 195–206, 2007.
30. Z. Zhang. Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: Overview and challenges. *IEEE Comm. Surveys & Tutorials*, 8(1):24–37, 2006.