Causality, Influence, and Computation in Possibly Disconnected Synchronous Dynamic Networks☆,☆☆

Othon Michail[a,*], Ioannis Chatzigiannakis[a], Paul G. Spirakis[a,b]

[a]*Computer Technology Institute & Press "Diophantus" (CTI), Patras, Greece*
[b]*Department of Computer Science, University of Liverpool, UK*

Abstract

In this work, we study the *propagation of influence and computation in dynamic distributed computing systems that are possibly disconnected at every instant*. We focus on a *synchronous message passing* communication model with *broadcast* and bidirectional links. Our network dynamicity assumption is a *worst-case dynamicity* controlled by an adversary scheduler, which has received much attention recently. We replace the usual (in worst-case dynamic networks) assumption that the network is connected at every instant by minimal *temporal connectivity* conditions. Our conditions only require that *another causal influence occurs within every time-window of some given length*. Based on this basic idea we define several novel metrics for capturing the speed of information spreading in a dynamic network. We present several results that correlate these metrics. Moreover, we investigate *termination criteria* in networks in which an upper bound on any of these metrics is known. We exploit our termination criteria to provide efficient (and optimal in some cases) protocols that solve the fundamental *counting* and *all-to-all token dissemination* (or *gossip*) problems.

Keywords:
dynamic graph, mobile computing, worst-case dynamicity, adversarial schedule, temporal connectivity, termination, counting, information dissemination, optimal protocol

1. Introduction

Distributed computing systems are more and more becoming dynamic. The static and relatively stable models of computation can no longer represent the plethora of recently established and rapidly emerging information and communication technologies. In recent years, we have seen a tremendous increase in the number of new mobile computing devices. Most of these devices are equipped with some sort of communication, sensing, and mobility capabilities. Even the Internet has become mobile. The design is now focused on complex collections of heterogeneous devices that should be robust, adaptive, and self-organizing, possibly moving around and serving requests that vary with time. Delay-tolerant networks are highly-dynamic, infrastructure-less networks whose essential characteristic is a possible absence of end-to-end communication routes at any instant. Mobility may be *active*, when the devices control and plan their mobility pattern (e.g. mobile robots), or *passive*, in opportunistic-mobility networks, where mobility stems from the mobility of the carries of the devices (e.g. humans carrying cell phones) or a combination of both (e.g. the devices

☆Supported in part by the project "Foundations of Dynamic Distributed Computing Systems" (FOCUS) which is implemented under the "ARISTEIA" Action of the Operational Programme "Education and Lifelong Learning" and is co-funded by the European Union (European Social Fund) and Greek National Resources.

☆☆A preliminary version of the results in this paper has appeared in [MCS12b].

*Corresponding author (Telephone number: +30 2610 960200, Fax number: +30 2610 960490, Postal Address: Computer Technology Institute & Press "Diophantus" (CTI), N. Kazantzaki Str., Patras University Campus, Rio, P.O. Box 1382, 26504, Greece).

Email addresses: michailo@cti.gr (Othon Michail), ichatz@cti.gr (Ioannis Chatzigiannakis), spirakis@cti.gr (Paul G. Spirakis)

have partial control over the mobility pattern, like for example when GPS devices provide route instructions to their carriers). Thus, it can vary from being completely predictable to being completely unpredictable. Gossip-based communication mechanisms, e-mail exchanges, peer-to-peer networks, and many other contemporary communication networks all assume or induce some sort of highly-dynamic communication network.

The formal study of dynamic communication networks is hardly a new area of research. There is a huge amount of work in distributed computing that deals with causes of dynamicity such as failures and changes in the topology that are rather slow and usually eventually stabilize (like, for example, in self-stabilizing systems [Dol00]). However the low rate of topological changes that is usually assumed there is unsuitable for reasoning about truly dynamic networks. Even graph-theoretic techniques need to be revisited: the suitable graph model is now that of a *dynamic graph* (a.k.a. *temporal graph* or *time-varying graph*) (see e.g. [MMCS13, KKK00, CFQS12]), in which each edge has an associated set of time-labels indicating availability times. Even fundamental properties of classical graphs do not easily carry over to their temporal counterparts. For example, Kempe, Kleinberg, and Kumar [KKK00] found out that there is no analogue of Menger's theorem (see e.g. [Bol98] for a definition) for arbitrary temporal networks with one label on every edge, which additionally renders the computation of the number of node-disjoint s-t paths **NP**-complete. Very recently, the authors of [MMCS13] achieved a reformulation of Menger's theorem which is valid for all temporal networks and additionally they introduced several interesting cost minimization parameters for optimal temporal network design and gave some first results on them. Even the standard network diameter metric is no more suitable and has to be replaced by a dynamic/temporal version. In a dynamic star graph in which all leaf-nodes but one go to the center one after the other in a modular way, any message from the node that enters last the center to the node that never enters the center needs $n-1$ steps to be delivered, where n is the size (number of nodes) of the network; that is the *dynamic diameter* is $n-1$ while, one the other hand, the classical diameter is just 2 [AKL08] (see also [KO11]).

2. Related Work

Distributed systems with worst-case dynamicity were first studied in [OW05]. Their outstanding novelty was to assume a communication network that may change arbitrarily from time to time subject to the condition that each instance of the network is connected. They studied asynchronous communication and considered nodes that can detect local neighborhood changes; these changes cannot happen faster than it takes for a message to transmit. They studied *flooding* (in which one node wants to disseminate one piece of information to all nodes) and *routing* (in which the information need only reach a particular destination node t) in this setting. They described a uniform protocol for flooding that terminates in $O(Tn^2)$ rounds using $O(\log n)$ bit storage and message overhead, where T is the maximum time it takes to transmit a message. They conjectured that without identifiers (IDs) flooding is impossible to solve within the above resources. Finally, a uniform routing algorithm was provided that delivers to the destination in $O(Tn)$ rounds using $O(\log n)$ bit storage and message overhead.

Computation under worst-case dynamicity was further and extensively studied in a series of works by Kuhn *et al.* in the synchronous case. In [KLO10], the network was assumed to be T-*interval connected* meaning that any time-window of length T has a static connected spanning subgraph (persisting throughout the window). Among others, *counting* (in which nodes must determine the size of the network) and *all-to-all token dissemination* (in which n different pieces of information, called tokens, are handed out to the n nodes of the network, each node being assigned one token, and all nodes must collect all n tokens) were solved in $O(n^2/T)$ rounds using $O(\log n)$ bits per message, almost-linear-time randomized approximate counting was established for $T = 1$, and two lower bounds on token dissemination were given. Several variants of *coordinated consensus* in 1-interval connected networks were studied in [KMO11]. Two interesting findings were that in the absence of a good initial upper bound on n, eventual consensus is as hard as computing deterministic functions of the input and that *simultaneous consensus* can never be achieved in less than $n-1$ rounds in any execution. [Hae11] is a recent work that presents information spreading algorithms in worst-case dynamic networks based on *network coding*. An *open* setting (modeled as high churn) in which nodes constantly join and leave has very recently been considered in [APRU12]. For an excellent introduction to distributed computation under worst-case dynamicity see [KO11]. Two very thorough surveys on dynamic networks are [Sch02, CFQS12].

Another notable model for dynamic distributed computing systems is the *population protocol* model [AAD$^+$06]. In that model, the computational agents are passively mobile, interact in ordered pairs, and the connectivity assumption is a *strong global fairness condition* according to which all events that may always occur, occur infinitely often. These assumptions give rise to some sort of structureless interacting automata model. The usually assumed *anonymity* and *uniformity* (i.e. n is not known) of protocols only allow for commutative computations that eventually stabilize to a desired configuration. Most computability issues in this area have now been established. Constant-state nodes on a complete interaction network (and several variations) compute the *semilinear predicates* [AAER07]. Semilinearity persists up to $o(\log\log n)$ local space but not more than this [CMN$^+$11]. If constant-state nodes can additionally leave and update fixed-length pairwise marks then the computational power dramatically increases to the commutative subclass of $\mathbf{NSPACE}(n^2)$ [MCS11a]. For a very recent introductory text see [MCS11b].

3. Contribution

In this work, we study worst-case dynamic networks that are *free of any connectivity assumption about their instances*. Our dynamic network model is formally defined in Section 4.1. We only impose some *temporal connectivity* conditions on the adversary guaranteeing that *another causal influence occurs within every time-window of some given length*, meaning that, in that time, another node first hears of the state that some node u had at some time t (see Section 4.3 for a formal definition of *causal influence*). Note that our temporal connectivity conditions are minimal assumptions that allow for bounded end-to-end communication in any dynamic network including those that have disconnected instances. Based on this basic idea, we define several novel generic metrics for capturing the speed of information spreading in a dynamic network. In particular, we define the *outgoing influence time* (oit) as the maximal time until the state of a node *influences* the state of another node, the *incoming influence time* (iit) as the maximal time until the state of a node *is influenced by* the state of another node, and the *connectivity time* (ct) as the maximal time until the two parts of any cut of the network become connected. These metrics are defined in Section 5, where also several results that correlate these metrics to themselves and to standard metrics, like e.g. the dynamic diameter, are presented.

In Section 5.1, we present a simple but very fundamental dynamic graph based on alternating matchings that has oit 1 (equal to that of instantaneous connectivity networks) but at the same time is *disconnected in every instance*. In Section 6, we exhibit another dynamic graph additionally guaranteeing that edges take maximal time to reappear. That graph is based on a geometric edge-coloring method due to Soifer for coloring a complete graph of even order n with $n-1$ colors [Soi09]. Similar results have appeared before but to the best of our knowledge only in probabilistic settings [CMM$^+$08, BCF09].

In Section 7, we turn our attention to terminating computations and, in particular, we investigate termination criteria in networks in which an upper bound on the ct or the oit is known. By "termination criterion" we essentially mean any locally verifiable property that can be used to determine whether a node has heard from all other nodes. Note that we do not allow to the nodes any further knowledge on the network; for instance, nodes *do not* know the dynamic diameter of the network. In particular, in Section 7.1, we study the case in which an upper bound T on the ct is known and we present an optimal termination criterion that only needs time linear in the dynamic diameter and in T. Then, in Section 7.2, we study the case in which an upper bound K on the oit is known. We first present a termination criterion that needs time $O(K \cdot n^2)$. Additionally, we establish that the optimal termination criterion for the ct case does not work in the oit case. These criteria share the fundamental property of "hearing from the past". We then develop a new technique that gives an optimal termination criterion (time linear in the dynamic diameter and in K) by "hearing from the future" (by this we essentially mean that a node is interested for its outgoing influences instead for its incoming ones). Additionally, we exploit throughout the paper our termination criteria to provide protocols that solve the fundamental *counting* and *all-to-all token dissemination* (or *gossip*) problems; in the former nodes must determine the size of the network n and in the latter each node of the network is provided with a unique piece of information, called *token*, and all nodes must collect all n tokens.

Finally, in Section 8, we conclude and discuss some interesting future research directions.

4. Preliminaries

4.1. The Dynamic Network Model

A *dynamic network* is modeled by a *dynamic graph* $G = (V, E)$, where V is a set of n nodes (or processors) and $E : \mathbb{N} \to \mathcal{P}(E')$ (wherever we use $\mathbb{N}$ we mean $\mathbb{N}_{\geq 1}$) is a function mapping a round number $r \in \mathbb{N}$ to a set $E(r)$ of bidirectional links drawn from $E' = \{\{u, v\} : u, v \in V\}$. [1] Intuitively, a dynamic graph G is an infinite sequence $G(1), G(2), \ldots$ of *instantaneous graphs*, whose edge sets are subsets of E' chosen by a *worst-case adversary.* A *static network* is just a special case of a dynamic network in which $E(i+1) = E(i)$ for all $i \in \mathbb{N}$. The set V is assumed throughout this work to be *static*, that is it remains the same throughout the execution.

We assume that nodes in V have unique identities (ids) drawn from some namespace $\mathcal{U}$ (we assume that ids are represented using $O(\log n)$ bits) and that they do not know the topology or the size of the network, apart from some minimal necessary knowledge to allow for terminating computations (usually an upper bound on the time it takes for information to make some sort of progress). Any such assumed knowledge will be clearly stated. Moreover, nodes have unlimited local storage (though they usually use a reasonable portion of it).

Communication is *synchronous message passing* [Lyn96, AW04], meaning that it is executed in discrete steps controlled by a global clock that is available to the nodes and that nodes communicate by sending and receiving messages (usually of length that is some reasonable function of n, like e.g. $\log n$). We use the terms *round*, *time*, and *step* interchangeably to refer to the discrete steps of the system. Naturally, real rounds begin to count from 1 (e.g. "first round") and we reserve time 0 to refer to the initial state of the system. We assume that the message transmission model is *anonymous broadcast*, in which, in every round r, each node u generates a single message $m_u(r)$ to be delivered to all its current neighbors in $N_u(r) = \{v : \{u, v\} \in E(r)\}$ without knowing $N_u(r)$.

In every round, the adversary first chooses the edges for the round; for this choice it can see the internal states of the nodes at the beginning of the round. At the same time and independently of the adversary's choice of edges each node generates its message for the current round. Note that a node does not have any information about the internal state of its neighbors when generating its messages (including their ids). In deterministic algorithms, nodes are only based on their current internal state to generate their messages and this implies that the adversary can infer the messages that will be generated in the current round before choosing the edges. In this work, we only consider deterministic algorithms. Each message is then delivered to the sender's neighbors, as chosen by the adversary; the nodes transition to new states, and the next round begins.

4.2. Problem Definitions

In this work, apart from studying the structural properties of possibly disconnected dynamic networks we also investigate the computability of the following two fundamental problems for distributed computing.

Counting. Nodes must determine the network size n.

All-to-all Token Dissemination (or Gossip). There is a token assignment function $I : V \to \mathcal{T}$ that assigns to each node $u \in V$ a single token $I(u)$ from some domain $\mathcal{T}$ s.t. $I(u) \neq I(v)$ for all $u \neq v$. An algorithm solves all-to-all token dissemination if for all instances (V, I), when the algorithm is executed in any dynamic graph $G = (V, E)$, all nodes eventually terminate and output $\bigcup_{u \in V} I(u)$. We assume that each token in the nodes' input is represented using $O(\log n)$ bits. The nodes know that each node starts with a unique token but they do not know n.

[1] By $\mathcal{P}(S)$ we denote the *powerset* of the set S, that is the set of all subsets of S.

4.3. Spread of Influence in Dynamic Graphs (Causal Influence)

Probably the most important notion associated with a dynamic network/graph is the *causal influence*, which formalizes the notion of one node "influencing" another through a chain of messages originating at the former node and ending at the latter (possibly going through other nodes in between). We denote by (u,t) the state of node u at time t and usually call it the *t-state of u*. The pair (u,t) is also called a *time-node*. We use $(u,r) \rightsquigarrow (v,r')$ to denote the fact that node u's state in round r influences node v's state in round r'. Formally:

Definition 1 ([Lam78])**.** *Given a dynamic graph $G = (V,E)$ we define an order $\rightarrow \subseteq (V \times \mathbb{N}_{\geq 0})^2$, where $(u,r) \rightarrow (v,r+1)$ iff $u = v$ or $\{u,v\} \in E(r+1)$. The* causal order $\rightsquigarrow \subseteq (V \times \mathbb{N}_{\geq 0})^2$ *is defined to be the reflexive and transitive closure of $\rightarrow$.*

Obviously, for a dynamic distributed system to operate as a whole there must exist some upper bound on the time needed for information to spread through the network. This is a very weak guarantee as without it global computation is in principle impossible. An abstract way to talk about information spreading is via the notion of the *dynamic diameter*. The *dynamic diameter* (also called *flooding time*, e.g., in [CMM+08, BCF09]) of a dynamic graph, is an upper bound on the time required for each node to causally influence (or, equivalently, to be causally influenced by) every other node; formally, the dynamic diameter is the minimum $D \in \mathbb{N}$ s.t. for all times $t \geq 0$ and all $u, v \in V$ it holds that $(u,t) \rightsquigarrow (v, t+D)$. A small dynamic diameter allows for fast dissemination of information. In this work, we do not allow nodes to know the dynamic diameter of the network. We only allow some minimal knowledge (that will be explained in the sequel) based on which nodes may infer bounds on the dynamic diameter.

A class of dynamic graphs with small dynamic diameter is that of *T-interval connected* graphs. T-interval connectivity was proposed in [KLO10] as an elegant way to capture a special class of dynamic networks, namely *those that are connected at every instant*. Intuitively, the parameter T represents the rate of connectivity changes. Formally, a dynamic graph $G = (V,E)$ is said to be *T-interval connected*, for $T \geq 1$, if, for all $r \in \mathbb{N}$, the static graph $G_{r,T} := (V, \bigcap_{i=r}^{r+T-1} E(r))$ is connected [KLO10]; that is, in every time-window of length T, a connected spanning subgraph is preserved. In one extreme, if $T = 1$ then the underlying connected spanning subgraph may change arbitrarily from round to round and in the other extreme if T is ∞ then a connected spanning subgraph must be preserved forever.

T-interval connected networks have the very nice feature to allow for constant propagation of information. For example, 1-interval connectivity guarantees that the state of a node causally influences the state of another uninfluenced node in every round (if one exists). To get an intuitive feeling of this fact, consider a partitioning of the set of nodes V to a subset V_1 of nodes that know the r-state of some node u and to a subset $V_2 = V \backslash V_1$ of nodes that do not know it. Connectivity asserts that there is always an edge in the cut between V_1 and V_2, consequently, if nodes that know the r-state of u broadcast it in every round, then in every round at least one node moves from V_2 to V_1.

This is formally captured by the following lemma from [KLO10].

Lemma 1 ([KLO10])**.** *For any node $u \in V$ and time $r \geq 0$, in a 1-interval connected network, we have*

1. $|\{v \in V : (u,0) \rightsquigarrow (v,r)\}| \geq \min\{r+1, n\}$,
2. $|\{v \in V : (v,0) \rightsquigarrow (u,r)\}| \geq \min\{r+1, n\}$.

Let us also define two very useful sets. For all times $0 \leq t \leq t'$, we define by $\mathrm{past}_{(u,t')}(t) := \{v \in V : (v,t) \rightsquigarrow (u,t')\}$ [KMO11] the *past set of a time-node (u,t') from time t* and by $\mathrm{future}_{(u,t)}(t') := \{v \in V : (u,t) \rightsquigarrow (v,t')\}$ the *future set of a time-node (u,t) at time t'*. In words, $\mathrm{past}_{(u,t')}(t)$ is the set of nodes whose t-state (i.e. their state at time t) has causally influenced the t'-state of u and $\mathrm{future}_{(u,t)}(t')$ is the set of nodes whose t'-state has been causally influenced by the t-state of u. If $v \in \mathrm{future}_{(u,t)}(t')$ we say that at time t' node v *has heard of/from* the t-state of node u. If it happens that $t = 0$ we say simply that v has heard of u. Note that $v \in \mathrm{past}_{(u,t')}(t)$ iff $u \in \mathrm{future}_{(v,t)}(t')$.

For a distributed system to be able to perform global computation, nodes need to be able to determine for all times $0 \leq t \leq t'$ whether $\mathrm{past}_{(u,t')}(t) = V$. If nodes know n, then a node can easily determine at time t' whether $\mathrm{past}_{(u,t')}(t) = V$ by counting all different t-states that it has heard of so far (provided that every node broadcasts at every round all information it knows). If it has heard the t-states of all nodes then the equality is satisfied. If n is not known then various techniques may be applied (which is the subject of this work). By *termination criterion* we mean any locally verifiable property that can be used to determine whether $\mathrm{past}_{(u,t')}(t) = V$.

Remark 1. *Note that any protocol that allows nodes to determine whether* $\mathrm{past}_{(u,t')}(t) = V$ *can be used to solve the counting and all-to-all token dissemination problems. The reason is that if a node knows at round r that it has been causally influenced by the initial states of all other nodes, then it can solve counting by writing* $|\mathrm{past}_{(u,r)}(0)|$ *on its output and all-to-all dissemination by writing* $\mathrm{past}_{(u,r)}(0)$ *(provided that all nodes send their initial states and all nodes constantly broadcast all initial states that they have heard of so far).*

5. Our Metrics

As already stated, in this work we aim to deal with dynamic networks that are allowed to have disconnected instances. To this end, we define some novel generic metrics that are particularly suitable for capturing the speed of information propagation in such networks.

5.1. The Influence Time

Recall that the guarantee on propagation of information resulting from instantaneous connectivity ensures that any time-node (u,t) influences another node *in each step* (if an uninfluenced one exists). From this fact, we extract two novel generic influence metrics that capture the maximal time until another influence (outgoing or incoming) of a time-node occurs.

We now formalize our first influence metric.

Definition 2 (Outgoing Influence Time)**.** *We define the* outgoing influence time *(oit) as the minimum* $k \in \mathbb{N}$ *s.t. for all* $u \in V$ *and all times* $t, t' \geq 0$ *s.t.* $t' \geq t$ *it holds that*

$$|\mathrm{future}_{(u,t)}(t'+k)| \geq \min\{|\mathrm{future}_{(u,t)}(t')| + 1, n\}.$$

Intuitively, the oit is the maximal time until the t-state of a node influences the state of another node (if an uninfluenced one exists) and captures the speed of information spreading.

Our second metric is similarly defined as follows.

Definition 3 (Incoming Influence Time)**.** *We define the* incoming influence time *(iit) as the minimum* $k \in \mathbb{N}$ *s.t. for all* $u \in V$ *and all times* $t, t' \geq 0$ *s.t.* $t' \geq t$ *it holds that*

$$|\mathrm{past}_{(u,t'+k)}(t)| \geq \min\{|\mathrm{past}_{(u,t')}(t)| + 1, n\}.$$

We can now say that the oit of a T-interval connected graph is 1 and that the iit can be up to $n-2$. However, is it necessary for a dynamic graph to be T-interval connected in order to achieve unit oit? First, let us make a simple but useful observation:

Proposition 1. *If a dynamic graph* $G = (V, E)$ *has oit (or iit) 1 then every instance has at least* $\lceil n/2 \rceil$ *edges.*

Proof. $\forall u \in V$ and $\forall t \geq 1$ it must hold that $\{u, v\} \in E(t)$ for some v. In words, at any time t each node must have at least one neighbor since otherwise it influences (or is influenced by) no node during round t. A minimal way to achieve this is by a perfect matching in the even-order case and by a matching between $n-3$ nodes and a linear graph between the remaining 3 nodes in the odd-order case. □

Proposition 1 is easily generalized as: if a dynamic graph $G = (V, E)$ has oit (or iit) k then for all times t it holds that $|\bigcup_{i=t}^{t+k-1} E(i)| \geq \lceil n/2 \rceil$. The reason is that now any node must have a neighbor in any k-window of the dynamic graph (and not necessarily in every round).

Now, inspired by Proposition 1, we define a minimal dynamic graph that at the same time satisfies oit 1 and always disconnected instances:

The Alternating Matchings Dynamic Graph. Take a ring of an even number of nodes $n = 2l$, partition the edges into 2 disjoint perfect matchings A and B (each consisting of l edges) and alternate round after round between the edge sets A and B (see Figure 1).

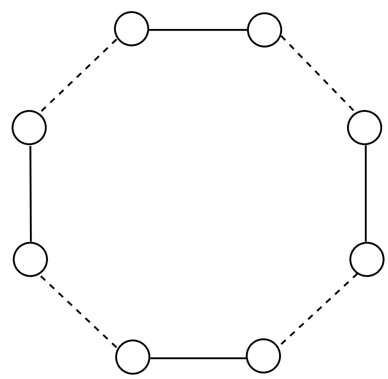

Figure 1: The Alternating Matchings dynamic graph for $n = 8$. The solid lines appear every odd round $(1, 3, 5, \ldots)$ while the dashed lines every even round $(2, 4, 6, \ldots)$.

Proposition 2. *The Alternating Matchings dynamic graph has* oit *1 and any node needs precisely* $n/2$ *rounds to influence all other nodes.*

Proof. Take any node u. In the first round, u influences its left or its right neighbor on the ring depending on which of its two adjacent edges becomes available first. Thus, including itself, it has influenced 2 nodes forming a line of length 1. In the next round, the two edges that join the endpoints of the line with the rest of the ring become available and two more nodes become influenced; the one is the neighbor on the left of the line and the other is the neighbor on the right. By induction on the number of rounds, it is not hard to see that the existing line always expands from its endpoints to the two neighboring nodes of the ring (one on the left and the other on the right). Thus, we get exactly 2 new influences per round, which gives oit 1 and $n/2$ rounds to influence all nodes. $\square$

In the alternating matchings construction any edge reappears every second step but not faster than this. We now formalize the notion of the *fastest edge reappearance* (fer) of a dynamic graph.

Definition 4. *The* fastest edge reappearance *(fer) of a dynamic graph* $G = (V, E)$ *is defined as the minimum* $p \in \mathbb{N}$ *s.t.,* $\exists e \in \{\{u, v\} : u, v \in V\}$ *and* $\exists t \in \mathbb{N}$, $e \in E(t) \cap E(t + p)$.

Clearly, the fer of the alternating matchings dynamic graph described above is 2, because no edge ever reappears in 1 step and some, at some point, (in fact, all and always) reappears in 2 steps. In Section 6, by invoking a geometric edge-coloring method, we generalize this basic construction to a more involved dynamic graph with oit 1, always disconnected instances, and fer equal to $n - 1$. [2]

We next give a proposition associating dynamic graphs with oit (or iit) upper bounded by K to dynamic graphs with connected instances.

Proposition 3. *Assume that the* oit *or the* iit *of a dynamic graph,* $G = (V, E)$*, is upper bounded by* K*. Then for all times* $t \in \mathbb{N}$ *the graph* $(V, \bigcup_{i=t}^{t+K\lfloor n/2 \rfloor - 1} E(i))$ *is connected.*

[2]It is interesting to note that in dynamic graphs with a static set of nodes (that is V does not change), if at least one change happens each time, then every instance $G(t)$ will eventually reappear after at most $\sum_{k=0}^{\binom{n}{2}} \binom{\binom{n}{2}}{k}$ steps. This counts all possible different graphs of n vertices with k edges and sums for all $k \geq 0$. Thus the fer is bounded from above by a function of n.

Proof. It suffices to show that for any partitioning (V_1, V_2) of V there is an edge in the cut labeled from $\{t, \ldots, t+K\lfloor n/2\rfloor-1\}$. W.l.o.g. let V_1 be the smaller one, thus $|V_1| \leq \lfloor n/2\rfloor$. Take any $u \in V_1$. By definition of oit, $|\text{future}_{(u,t)}(t + K\lfloor n/2\rfloor - 1)| \geq |\text{future}_{(u,t)}(t + K|V_1| - 1)| \geq |V_1| + 1$ implying that some edge in the cut has transferred u's t-state out of V_1 at some time in the interval $[t, t + K\lfloor n/2\rfloor - 1]$. The proof for the iit is similar. □

5.2. The Moi (Concurrent Progress)

Consider now the following influence metric:

Definition 5. *Define the* maximum outgoing influence *(moi) of a dynamic graph $G = (V, E)$ as the maximum k for which $\exists u \in V$ and $\exists t, t' \in \mathbb{N}$, $t' \geq t$, s.t. $|\text{future}_{(u,t)}(t' + 1)| - |\text{future}_{(u,t)}(t')| = k$.*

In words, the moi of a dynamic graph is the maximum number of nodes that are ever concurrently influenced by a time-node.

Here we show that one cannot guarantee at the same time unit oit and at most one outgoing influence per node per step. In fact, we conjecture that unit oit implies that some node disseminates in $\lfloor n/2\rfloor$ steps.

We now prove an interesting theorem stating that if one tries to guarantee unit oit then she must necessarily accept that at some steps more than one outgoing influences of the same time-node will occur leading to faster dissemination than $n - 1$ for this particular node.

Theorem 1. *The moi of any dynamic graph with $n \geq 3$ and unit oit is at least 2.*

Proof. For $n = 3$, just notice that unit oit implies that, at any time t, some node has necessarily 2 neighbors. We therefore focus on $n \geq 4$. For the sake of contradiction, assume that the statement is not true. Then at any time t any node u is connected to exactly one other node v (at least one neighbor is required for oit 1 - see Proposition 1 - and at most one is implied by our assumption). Unit oit implies that, at time $t + 1$, at least one of u, v must be connected to some $w \in V\backslash\{u, v\}$, let it be v. Proposition 1 requires that also u must have an edge labeled $t + 1$ incident to it. If that edge arrives at v, then v has 2 edges labeled $t + 1$. If it arrives at w, then w has 2 edges labeled $t + 1$. So it must arrive at some $z \in V\backslash\{u, v, w\}$. Note now that, in this case, the $(t-1)$-state of u first influences both w, z at time $t+1$ which is contradictory, consequently the moi must be at least 2. □

In fact, notice that the above theorem proves something stronger: Every second step at least half of the nodes influence at least 2 new nodes each. This, together with the fact that it seems to hold for some basic cases, makes us suspect that the following conjecture might be true:

Conjecture 1. *If the oit of a dynamic graph is 1 then $\forall t \in \mathbb{N}$, $\exists u \in V$ s.t. $|\text{future}_{(u,t)}(t + \lfloor n/2\rfloor)| = n$.*

That is, if the oit is 1 then, in every $\lfloor n/2\rfloor$-window, some node influences all other nodes (e.g. influencing 2 new nodes per step).

5.3. The Connectivity Time

We now propose another natural and practical metric for capturing the temporal connectivity of a possibly disconnected dynamic network that we call the *connectivity time* (ct).

Definition 6 (Connectivity Time)**.** *We define the* connectivity time *(ct) of a dynamic network $G = (V, E)$ as the minimum $k \in \mathbb{N}$ s.t. for all times $t \in \mathbb{N}$ the static graph $(V, \bigcup_{i=t}^{t+k-1} E(i))$ is connected.*

In words, the ct of a dynamic network is the maximal time of keeping the two parts of any cut of the network disconnected. That is to say, in every ct-window of the network an edge appears in every (V_1, V_2)-cut. Note that, in the extreme case in which the ct is 1, every instance of the dynamic graph is connected and we thus obtain a 1-interval connected graph. On the other hand, greater ct allows for different cuts to be connected at different times in the ct-round interval and the resulting dynamic graph can very well have disconnected instances. For an illustrating example, consider again the alternating matchings graph from

Section 5.1. Draw a line that crosses two edges belonging to matching A partitioning the ring into two parts. Clearly, these two parts communicate every second round (as they only communicate when matching A becomes available), thus the ct is 2 and every instance is disconnected. We now provide a result associating the ct of a dynamic graph with its oit.

Proposition 4. *(i) $\mathsf{oit} \leq \mathsf{ct}$ but (ii) there is a dynamic graph with oit 1 and $\mathsf{ct} = \Omega(n)$.*

Proof. (i) We show that for all $u \in V$ and all times $t, t' \in \mathbb{N}$ s.t. $t' \geq t$ it holds that $|\text{future}_{(u,t)}(t' + \mathsf{ct})| \geq \min\{|\text{future}_{(u,t)}(t')| + 1, n\}$. Assume $V \backslash \text{future}_{(u,t)}(t') \neq \emptyset$ (as the other case is trivial). In at most ct rounds at least one edge joins $\text{future}_{(u,t)}(t')$ to $V \backslash \text{future}_{(u,t)}(t')$. Thus, in at most ct rounds $\text{future}_{(u,t)}(t')$ increases by at least one.

(ii) Recall the alternating matchings on a ring dynamic graph from Section 5.1. Now take any set V of a number of nodes that is a multiple of 4 (this is just for simplicity and is not necessary) and partition it into two sets V_1, V_2 s.t. $|V_1| = |V_2| = n/2$. If each part is an alternating matchings graph for $|V_1|/2$ rounds then every u say in V_1 influences 2 new nodes in each round and similarly for V_2. Clearly we can keep V_1 disconnected from V_2 for $n/4$ rounds without violating $\mathsf{oit} = 1$. □

The following is a comparison of the ct of a dynamic graph with its dynamic diameter D.

Proposition 5. *$\mathsf{ct} \leq D \leq (n-1)\mathsf{ct}$.*

Proof. $\mathsf{ct} \leq D$ follows from the fact that in time equal to the dynamic diameter every node causally influences every other node and thus, in that time, there must have been an edge in every cut (if not, then the two partitions forming the cut could not have communicated with one another). $D \leq (n-1)\mathsf{ct}$ holds as follows. Take any node u and add it to a set S. In ct rounds u influences some node from $V \backslash S$ which is then added to S. In $(n-1)\mathsf{ct}$ rounds S must have become equal to V, thus this amount of time is sufficient for every node to influence every other node. Finally, we point out that these bounds cannot be improved in general as for each of $\mathsf{ct} = D$ and $D = (n-1)\mathsf{ct}$ there is a dynamic graph realizing it. $\mathsf{ct} = D$ is given by the dynamic graph that has no edge for $\mathsf{ct} - 1$ rounds and then becomes the complete graph while $D = (n-1)\mathsf{ct}$ is given by a line in which every edge appears at times $\mathsf{ct}, 2\mathsf{ct}, 3\mathsf{ct}, \ldots$. □

Note that the ct metric has been defined as an underapproximation of the dynamic diameter. Its main advantage is that it is much easier to compute than the dynamic diameter since it is defined on the union of the footprints and not on the dynamic adjacency itself.

6. Fast Propagation of Information Under Continuous Disconnectivity

In Section 5.1, we presented a simple example of an always-disconnected dynamic graph, namely, the alternating matchings dynamic graph, with optimal oit (i.e. unit oit). Note that the alternating matchings dynamic graph may be conceived as simple as it has small fer (equal to 2). We pose now an interesting question: Is there an always-disconnected dynamic graph with unit oit and fer as big as $n-1$? Note that this is harder to achieve as it allows of no edge to ever reappear in less than $n-1$ steps. Here, by invoking a geometric edge-coloring method, we arrive at an always-disconnected graph with unit oit and maximal fer; in particular, no edge reappears in less than $n-1$ steps.

To answer the above question, we define a very useful dynamic graph coming from the area of edge-coloring.

Definition 7. *We define the following dynamic graph S based on an edge-coloring method due to Soifer [Soi09]: $V(S) = \{u_1, u_2, \ldots, u_n\}$ where $n = 2l$, $l \geq 2$. Place u_n on the center and $u_1, \ldots, u_{n-1}$ on the vertices of a $(n-1)$-sided polygon. For each time $t \geq 1$ make available only the edges $\{u_n, u_{m_t(0)}\}$ for $m_t(j) := (t - 1 + j \bmod n-1) + 1$ and $\{u_{m_t(-i)}, u_{m_t(i)}\}$ for $i = 1, \ldots, n/2 - 1$; that is make available one edge joining the center to a polygon-vertex and all edges perpendicular to it. (e.g. see Figure 2 for $n = 8$ and $t = 1, \ldots, 7$).*

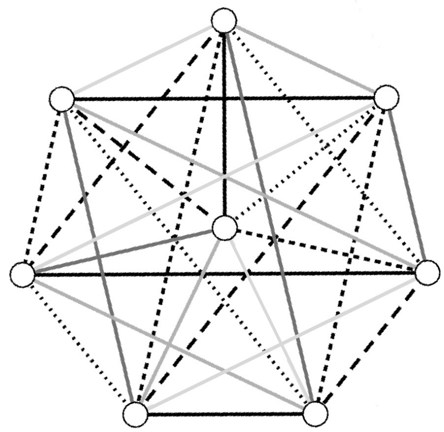

Figure 2: Soifer's dynamic graph for $n = 8$ and $t = 1, \dots, 7$. In particular, in round 1 the graph consists of the black solid edges, then in round 2 the center becomes connected via a dotted edge to the next peripheral node clockwise and all edges perpendicular to it (the remaining dotted ones) become available, and so on, always moving clockwise.

In Soifer's dynamic graph, denote by $N_u(t) := i : \{u, u_i\} \in E(t)$, that is the index of the unique neighbor of u at time t. The following lemma states that the next neighbor of a node is in almost all cases (apart from some trivial ones) the one that lies two positions clockwise from its current neighbor.

Lemma 2. *For all times $t \in \{1, 2, \dots, n-2\}$ and all u_k, $k \in \{1, 2, \dots, n-1\}$ it holds that $N_{u_k}(t+1) = n$ if $N_{u_k}(t) = (k - 3 \bmod n-1) + 1$ else $N_{u_k}(t+1) = (k + 1 \bmod n-1) + 1$ if $N_{u_k}(t) = n$ and $N_{u_k}(t+1) = (N_{u_k}(t) + 1 \bmod n-1) + 1$ otherwise.*

Proof. Since $k \notin \{n, t, t+1\}$ it easily follows that $k, N_k(t), N_k(t+1) \neq n$ thus both $N_k(t)$ and $N_k(t+1)$ are determined by $\{u_{m_t(-i)}, u_{m_t(i)}\}$ where $m_t(j) := (t-1+j \bmod n-1)+1$ and $k = m_t(-i)$. The latter implies $(t-1-i \bmod n-1)+1 = k \Rightarrow (t-1+i \bmod n-1)+1+(-2i \bmod n-1) = k \Rightarrow m_t(i) = k-(-2i \bmod n-1)$; thus, $N_k(t) = k - (-2i \bmod n-1)$. Now let us see how the i that corresponds to some node changes as t increases. When t increases by 1, we have that $(t-1+i \bmod n-1)+1 = (t+i' \bmod n-1)+1 \Rightarrow i' = i-1$, i.e. as t increases i decreases. Consequently, for $t+1$ we have $N_k(t+1) = k - [-2(i-1) \bmod n-1] = (N_{u_k}(t) + 1 \bmod n-1) + 1$. □

Theorem 2. *For all $n = 2l$, $l \geq 2$, there is a dynamic graph of order n, with* oit *equal to 1,* fer *equal to $n-1$, and in which every instance is a perfect matching.*

Proof. The dynamic graph is the one of Definition 7. It is straightforward to observe that every instance is a perfect matching. We prove now that the oit of this dynamic graph is 1. We focus on the set $\text{future}_{(u_n,0)}(t)$, that is the outgoing influence of the initial state of the node at the center. Note that symmetry guarantees that the same holds for all time-nodes (it can be verified that any node can be moved to the center without altering the graph). u_n at time 1 meets u_1 and thus $\text{future}_{(u_n,0)}(1) = \{u_1\}$. Then, at time 2, u_n meets u_2 and, by Lemma 2, u_1 meets u_3 via the edge than is perpendicular to $\{u_n, u_2\}$, thus $\text{future}_{(u_n,0)}(2) = \{u_1, u_2, u_3\}$. We show that for all times t it holds that $\text{future}_{(u_n,0)}(t) = \{u_1, \dots, u_{2t-1}\}$. The base case is true since $\text{future}_{(u_n,0)}(1) = \{u_1\}$. It is not hard to see that, for $t \geq 2$, $N_{u_2}(t) = 2t-2$, $N_{u_1}(t) = 2t-1$, and for all $u_i \in \text{future}_{(u_n,0)}(t) \backslash \{u_1, u_2\}$, $1 \leq N_{u_i}(t) \leq 2t-2$. Now consider time $t+1$. Lemma 2 guarantees now that for all $u_i \in \text{future}_{(u_n,0)}(t)$ we have that $N_{u_i}(t+1) = N_{u_i}(t) + 2$. Thus, the only new influences at step $t+1$ are by u_1 and u_2 implying that $\text{future}_{(u_n,0)}(t+1) = \{u_1, \dots, u_{2(t+1)-1}\}$. Consequently, the oit is 1.

The fer is $n-1$ because the edges leaving the center appear one after the other in a clockwise fashion, thus taking $n-1$ steps to any such edge to reappear, and, by construction, any other edge appears only when its unique perpendicular that is incident to the center appears (thus, again every $n-1$ steps). □

Note that Theorem 2 is optimal w.r.t. fer as it is impossible to achieve at the same time unit oit and fer strictly greater than $n-1$. To see this, notice that if no edge is allowed to reappear in less than n steps then any node must have no neighbors once every n steps.

7. Termination and Computation

We now turn our attention to termination criteria that we exploit to solve the fundamental counting and all-to-all token dissemination problems. First observe that if nodes know an upper bound H on the iit then there is a straightforward optimal termination criterion taking time $D+H$, where D is the dynamic diameter. In every round, all nodes forward all ids that they have heard of so far. If a node does not hear of a new id for H rounds then it must have already heard from all nodes. We begin this section by assuming that nodes know an upper bound on the ct and show how this initial knowledge can be exploited for optimal termination. Then, we allow the nodes to know an upper bound on the oit. In this case, things turn out to be much harder. We give a termination criterion which, though being far from the dynamic diameter, is optimal if a node terminates based on its past set. We then develop a novel technique that gives an optimal termination criterion based on the future set of a node. Keep in mind that nodes have no *a priori* knowledge of the size of the network.

7.1. Nodes Know an Upper Bound on the ct: An Optimal Termination Criterion

We here assume that all nodes know some upper bound T on the ct. We will give an optimal condition that allows a node to determine whether it has heard from all nodes in the graph. This condition results in an algorithm for counting and all-to-all token dissemination which is optimal, requiring $D+T$ rounds in any dynamic network with dynamic diameter D. The core idea is to have each node keep track of its past sets from time 0 and from time T and terminate as soon as these two sets become equal. This technique is inspired by [KMO11], where a comparison between the past sets from time 0 and time 1 was used to obtain an optimal termination criterion in 1-interval connected networks.

Theorem 3 (Repeated Past)**.** *Node u knows at time t that* $\mathrm{past}_{(u,t)}(0)=V$ *iff* $\mathrm{past}_{(u,t)}(0)=\mathrm{past}_{(u,t)}(T)$.

Proof. If $\mathrm{past}_{(u,t)}(0)=\mathrm{past}_{(u,t)}(T)$ then we have that $\mathrm{past}_{(u,t)}(T)=V$. The reason is that $|\mathrm{past}_{(u,t)}(0)|\geq\min\{|\mathrm{past}_{(u,t)}(T)|+1,n\}$. To see this, assume that $V\backslash\mathrm{past}_{(u,t)}(T)\neq\emptyset$. At most by round T there is some edge joining some $w\in V\backslash\mathrm{past}_{(u,t)}(T)$ to some $v\in\mathrm{past}_{(u,t)}(T)$. Thus, $(w,0)\rightsquigarrow(v,T)\rightsquigarrow(u,t)\Rightarrow w\in\mathrm{past}_{(u,t)}(0)$. In words, all nodes in $\mathrm{past}_{(u,t)}(T)$ belong to $\mathrm{past}_{(u,t)}(0)$ and at least one node not in $\mathrm{past}_{(u,t)}(T)$ (if one exists) must belong to $\mathrm{past}_{(u,t)}(0)$ (see also Figure 3).

For the other direction, assume that there exists $v\in\mathrm{past}_{(u,t)}(0)\backslash\mathrm{past}_{(u,t)}(T)$. This does not imply that $\mathrm{past}_{(u,t)}(0)\neq V$ but it does imply that even if $\mathrm{past}_{(u,t)}(0)=V$ node u cannot know it has heard from everyone. Note that u heard from v at some time $T'<T$ but has not heard from v since then. It can be the case that arbitrarily many nodes were connected to no node until time $T-1$ and from time T onwards were connected only to node v (v in some sense conceals these nodes from u). As u has not heard from the T-state of v it can be the case that it has not heard at all from arbitrarily many nodes, thus it cannot decide on the count. □

We now give a time-optimal algorithm for counting and all-to-all token dissemination that is based on Theorem 3.

Protocol *A*. All nodes constantly forward all 0-states and T-states of nodes that they have heard of so far (in this protocol, these are just the ids of the nodes accompanied with 0 and T timestamps, respectively) and a node halts as soon as $\mathrm{past}_{(u,t)}(0)=\mathrm{past}_{(u,t)}(T)$ and outputs $|\mathrm{past}_{(u,t)}(0)|$ for counting or $\mathrm{past}_{(u,t)}(0)$ for all-to-all dissemination.

For the time-complexity notice that any state of a node needs D rounds to causally influence all nodes, where D is the dynamic diameter. Clearly, by time $D+T$, u must have heard of the 0-state and T-state of all nodes, and at that time $\mathrm{past}_{(u,t)}(0)=\mathrm{past}_{(u,t)}(T)$ is satisfied. It follows that all nodes terminate in at

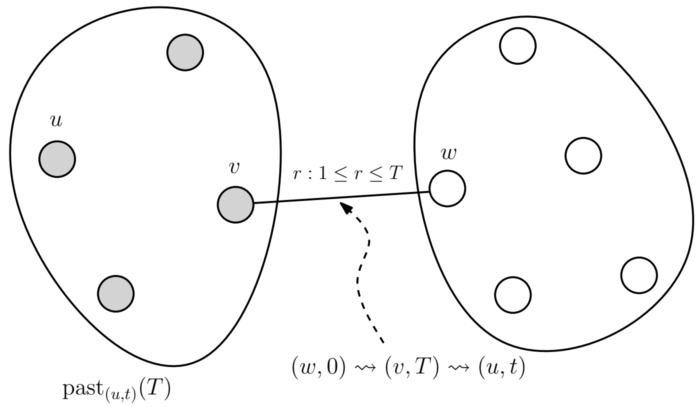

Figure 3: A partitioning of V into two sets. The left set is $\text{past}_{(u,t)}(T)$, i.e. the set of nodes whose T-state has influenced u by time t. All nodes in $\text{past}_{(u,t)}(T)$ also belong to $\text{past}_{(u,t)}(0)$. Looking back in time at the interval $[1, T]$, there should be an edge from some v in the left set to some w in the right set. This implies that v has heard from w by time T and as u has heard from the T-state of v it has also heard from the initial state of w. This implies that $\text{past}_{(u,t)}(0)$ is a strict superset of $\text{past}_{(u,t)}(T)$ as long as the right set is not empty.

most $D + T$ rounds. Optimality follows from the fact that this protocol terminates as long as $\text{past}_{(u,t)}(0) = \text{past}_{(u,t)}(T)$ which by the "only if" part of the statement of Theorem 3 is a necessary condition for correctness (any protocol terminating before this may terminate without having heard from all nodes).

7.2. Known Upper Bound on the oit*: Another Optimal Termination Criterion*

Now we assume that all nodes know some upper bound K on the oit.

7.2.1. Inefficiency of Hearing the Past

We begin by proving that if a node u has at some point heard of l nodes, then u hears of another node in $O(Kl^2)$ rounds (if an unknown one exists).

Theorem 4. *In any given dynamic graph with* oit *upper bounded by K, take a node u and a time t and denote $|\text{past}_{(u,t)}(0)|$ by l. It holds that $|\{v : (v, 0) \rightsquigarrow (u, t + Kl(l + 1)/2)\}| \geq \min\{l + 1, n\}$.*

Proof. Consider a node u and a time t and define $A_u(t) := \text{past}_{(u,t)}(0)$ (we only prove it for the initial states of nodes but easily generalizes to any time), $I_u(t') := \{v \in A_u(t) : A_v(t') \backslash A_u(t) \neq \emptyset\}$, $t' \geq t$, that is $I_u(t')$ contains all nodes in $A_u(t)$ whose t'-states have been influence by nodes not in $A_u(t)$ (these nodes know new info for u), $B_u(t') := A_u(t) \backslash I_u(t')$, that is all nodes in $A_u(t)$ that do not know new info, and $l := |A_u(t)|$. The only interesting case is for $V \backslash A_u(t) \neq \emptyset$. Since the oit is at most K we have that at most by round $t + Kl$, (u, t) influences some node in $V \backslash B_u(t)$ say via some $u_2 \in B_u(t)$. By that time, u_2 leaves B_u. Next consider $(u, t + Kl + 1)$. In $K(l - 1)$ steps it must influence some node in $V \backslash B_u$ since now u_2 is not in B_u. Thus, at most by round $t + Kl + K(l - 1)$ another node, say e.g. u_3, leaves B_u. In general, it holds that $B_u(t' + K|B_u(t')|) \leq \max\{|B_u(t')| - 1, 0\}$. It is not hard to see that at most by round $j = t + K(\sum_{1 \leq i \leq l} i)$, B_u becomes empty, which by definition implies that u has been influenced by the initial state of a new node. In summary, u is influenced by another initial state in at most $K(\sum_{1 \leq i \leq l} i) = kl(l + 1)/2$ steps. □

The good thing about the upper bound of Theorem 4 is that it associates the time for a new incoming influence to arrive at a node only with an upper bound on the oit, which is known, and the number of existing incoming influences which is also known, and thus the bound is locally computable at any time. So, there is a straightforward translation of this bound to a termination criterion and, consequently, to an algorithm for counting and all-to-all dissemination based on it.

Protocol *B*. All nodes constantly broadcast all ids that they have heard of so far. Each node u keeps a set $A_u(r)$ containing the ids it knows at round r and a termination bound $H_u(r)$ initially equal to K. If, at round r, u hears of new nodes, it inserts them in $A_u(r)$ and sets $H_u(r) \leftarrow r + Kl(l+1)/2$, where $l = |A_u(r)|$. If it ever holds that $r > H_u(r)$, u halts and outputs $|A_u(r)|$ for counting or $A_u(r)$ for all-to-all dissemination.

In the worst case, u needs $O(Kn)$ rounds to hear from all nodes and then another $Kn(n+1)/2 = O(Kn^2)$ rounds to realize that it has heard from all. So, the time complexity is $O(Kn^2)$.

Note that the upper bound of Theorem 4 is loose. The reason is that if a dynamic graph has oit upper bounded by K then in $O(Kn)$ rounds all nodes have causally influenced all other nodes and clearly the iit can be at most $O(Kn)$. We now show that there is indeed a dynamic graph that achieves this worst possible gap between the iit and the oit.

Theorem 5. *There is a dynamic graph with oit k but iit $k(n-3)$.*

Proof. Consider the dynamic graph $G = (V, E)$ s.t. $V = \{u_1, u_2, \ldots, u_n\}$ and u_i, for $i \in \{1, n-1\}$, is connected to u_{i+1} via edges labeled jk for $j \in \mathbb{N}_{\geq 1}$, u_i, for $i \in \{2, 3, \ldots, n-2\}$, is connected to u_{i+1} via edges labeled jk for $j \in \mathbb{N}_{\geq 2}$. and u_2 is connected to u_i, for $i \in \{3, \ldots, n-1\}$ via edges labeled k. In words, at step k, u_1 is only connected to u_2, u_2 is connected to all nodes except from u_n and u_n is connected to u_{n-1}. Then every multiple of k there is a single linear graph starting from u_1 and ending at u_n. At step k, u_2 is influenced by the initial states of nodes $\{u_3, \ldots, u_{n-1}\}$. Then at step $2k$ it forwards these influences to u_1. Since there are no further shortcuts, u_n's state needs $k(n-1)$ steps to arrive at u_1, thus there is an incoming-influence-gap of $k(n-2)$ steps at u_1. To see that oit is indeed k we argue as follows. Node u_1 cannot use the shortcuts, thus by using just the linear graph it influences a new node every k steps. u_2 influences all nodes apart from u_n at time k and then at time $2k$ it also influences u_n. All other nodes do a shortcut to u_2 at time k and then for all multiples of k their influences propagate to both directions from two sources, themselves and u_2, influencing 1 to 4 new nodes every k steps. □

Next we show that the $Kl(l+1)/2$ ($l := |\text{past}_{(u,t)}(0)|$) upper bound (of Theorem 4), on the time for another incoming influence to arrive, is optimal in the following sense: a node cannot obtain a better upper bound based solely on K and l. We establish this by showing that it is possible that a new incoming influence needs $\Theta(Kl^2)$ rounds to arrive, which excludes the possibility of a $o(Kl^2)$-bound to be correct as a protocol based on it may have nodes terminate without having heard of arbitrarily many other nodes. This, additionally, constitutes a tight example for the bound of Theorem 4.

Theorem 6. *For all n, l, K s.t. $n = \Omega(Kl^2)$, there is a dynamic graph with oit upper bounded by K and a round r such that, a node that has heard of l nodes by round r does not hear of another node for $\Theta(Kl^2)$ rounds.*

Proof. Consider the set $\text{past}_{(u,t)}(0)$ and denote its cardinality by l. Take any dynamic graph on $\text{past}_{(u,t)}(0)$, disconnected from the rest of the nodes, that satisfies oit $\leq K$ and that all nodes in $\text{past}_{(u,t)}(0)$ need $\Theta(Kl)$ rounds to causally influence all other nodes in $\text{past}_{(u,t)}(0)$; this could, for example, be the alternating matchings graph from Section 5.1 with one matching appearing in rounds that are odd multiples of K and the other in even. In $\Theta(Kl)$ rounds, say in round j, some *intermediary* node $v \in \text{past}_{(u,t)}(0)$ must get the outgoing influences of nodes in $\text{past}_{(u,t)}(0)$ outside $\text{past}_{(u,t)}(0)$ so that they continue to influence new nodes. Assume that in round $j-1$ the adversary directly connects all nodes in $\text{past}_{(u,t)}(0)\backslash\{v\}$ to v. In this way, at time j, v forwards outside $\text{past}_{(u,t)}(0)$ the $(j-2)$-states (and all previous ones) of all nodes in $\text{past}_{(u,t)}(0)$. Provided that $V\backslash\text{past}_{(u,t)}(0)$ is sufficiently big (see below) the adversary can now keep $S = \text{past}_{(u,t)}(0)\backslash\{v\}$ disconnected from the rest of the nodes for another $\Theta(Kl)$ rounds (in fact, one round less this time) without violating oit $\leq K$ as the new influences of the $(j-2)$-states of nodes in S may keep occurring outside S. The same process repeats by a new *intermediary* $v_2 \in S$ playing the role of v this time. Each time the process repeats, in $\Theta(|S|)$ rounds the intermediary gets all outgoing influences outside S and is then removed from S. It is straightforward to observe that a new incoming influence needs $\Theta(Kl^2)$ rounds to arrive at u in such a dynamic network. Moreover, note that $V\backslash\text{past}_{(u,t)}(0)$ should also satisfy oit $\leq K$ but this is easy to achieve by e.g. another alternating matchings dynamic graph on $V\backslash\text{past}_{(u,t)}(0)$ this time. Also, $n - l = |V\backslash\text{past}_{(u,t)}(0)|$ should satisfy $n - l = \Omega(Kl^2) \Rightarrow n = \Omega(Kl^2)$ so that the time needed for a $w \in V\backslash\text{past}_{(u,t)}(0)$ (in an alternating matchings dynamic graph on $V\backslash\text{past}_{(u,t)}(0)$) to influence all nodes in

$V\backslash\text{past}_{(u,t)}(0)$ and start influencing nodes in $\text{past}_{(u,t)}(0)$ is asymptotically greater than the time needed for S to extinct. To appreciate this, observe that if $V\backslash\text{past}_{(u,t)}(0)$ was too small then the outgoing influences of some $w \in V\backslash\text{past}_{(u,t)}(0)$ that occur every K rounds would reach u before the $\Theta(Kl^2)$ bound was achieved. Finally, we note that whenever the number of nodes in $V\backslash S$ becomes odd we keep the previous alternating matchings dynamic graph and the new node becomes connected every K rounds to an arbitrary node (the same in every round). When $|V\backslash S|$ becomes even again we return to a standard alternating matchings dynamic graph. $\square$

We now show that even the criterion of Theorem 3, that is optimal if an upper bound on the ct is known, does not work in dynamic graphs with known an upper bound K on the oit. In particular, we show that for all times $t' < K(n/4)$ there is a dynamic graph with oit upper bounded by K, a node u, and a time $t \in \mathbb{N}$ s.t. $\text{past}_{(u,t)}(0) = \text{past}_{(u,t)}(t')$ while $\text{past}_{(u,t)}(0) \neq V$. In words, for any such t' it can be the case that while u has not been yet causally influenced by all initial states its past set from time 0 may become equal to its past set from time t', which violates the termination criterion of Theorem 3.

Theorem 7. *For all n, K and all times $t' < K(n/4)$ there is a dynamic graph with oit upper bounded by K, a node u, and a time $t > t'$ s.t. $\text{past}_{(u,t)}(0) = \text{past}_{(u,t)}(t')$ while $\text{past}_{(u,t)}(0) \neq V$.*

Proof. For simplicity assume that n is a multiple of 4. As in Proposition 4 (ii), by an alternating matchings dynamic graph, we can keep two parts V_1, V_2 of the network, of size $n/2$ each, disconnected up to time $K(n/4)$. Let $u \in V_1$. At any time t, s.t. $t' < t \leq K(n/4)$, the adversary directly connects $u \in V_1$ to all $w \in V_1$. Clearly, at that time, u learns the t'-states (and thus also the 0-states) of all nodes in V_1 and, due to the disconnectivity of V_1 and V_2 up to time $K(n/4)$, u hears (and has heard up to then) of no node from V_2. It follows that $\text{past}_{(u,t)}(0) = \text{past}_{(u,t)}(t')$ and $|\text{past}_{(u,t)}(0)| = n/2 \Rightarrow \text{past}_{(u,t)}(0) \neq V$ as required. $\square$

7.2.2. Hearing the Future

In contrast to the previous negative results, we now present an optimal protocol for counting and all-to-all dissemination in dynamic networks with known an upper bound K on the oit, that is based on the following termination criterion. By definition of oit, if $\text{future}_{(u,0)}(t) = \text{future}_{(u,0)}(t+K)$ then $\text{future}_{(u,0)}(t) = V$. The reason is that if there exist uninfluenced nodes, then at least one such node must be influenced in at most K rounds, otherwise no such node exists and $(u, 0)$ must have already influenced all nodes (see also Figure 4). So, a fundamental goal is to allow a node to know its future set. Note that this criterion has a very basic difference from all termination criteria that have so far been applied to worst-case dynamic networks: instead of keeping track of its past set(s) and waiting for new incoming influences a node now directly keeps track of its future set and is informed by other nodes of its progress. We assume, for simplicity, a unique leader ℓ in the initial configuration of the system (this is not a necessary assumption and we will soon show how it can be dropped).

Protocol *Hear_from_known.* We denote by r the current round. Each node u keeps a list $Infl_u$ in which it keeps track of all nodes that first heard of $(\ell, 0)$ (the initial state of the leader) by u (u was between those nodes that first delivered $(\ell, 0)$ to nodes in $Infl_u$), a set A_u in which it keeps track of the $Infl_v$ sets that it is aware of initially set to $(u, Infl_u, 1)$, and a variable *timestamp* initially set to 1. Each node u broadcast in every round (u, A_u) and if it has heard of $(\ell, 0)$ also broadcasts $(\ell, 0)$. Upon reception of an id w that is not accompanied with $(\ell, 0)$, a node u that has already heard of $(\ell, 0)$ adds (w, r) to $Infl_u$ to recall that at round r it notified w of $(\ell, 0)$ (note that it is possible that other nodes also notify w of $(\ell, 0)$ at the same time without u being aware of them; all these nodes will write (w, r) in their lists). If it ever holds at a node u that $r > \max_{(v \neq u, r') \in Infl_u}\{r'\} + K$ then u adds (u, r) in $Infl_u$ (replacing any existing $(u, t) \in Infl_u$) to denote the fact that r is the maximum known time until which u has performed no further propagations of $(\ell, 0)$. If at some round r a node u modifies its $Infl_u$ set, it sets $timestamp \leftarrow r$. In every round, a node u updates A_u by storing in it the most recent $(v, Infl_v, timestamp)$ triple of each node v that it has heard of so far (its own $(u, Infl_u, timestamp)$ inclusive), where the "most recent" triple of a node v is the one with the greatest *timestamp* between those whose first component is v. Moreover, u clears multiple (w, r) records from the $Infl_v$ lists of A_u. In particular, it keeps (w, r) only in the $Infl_v$ list of the node v with the smallest id between

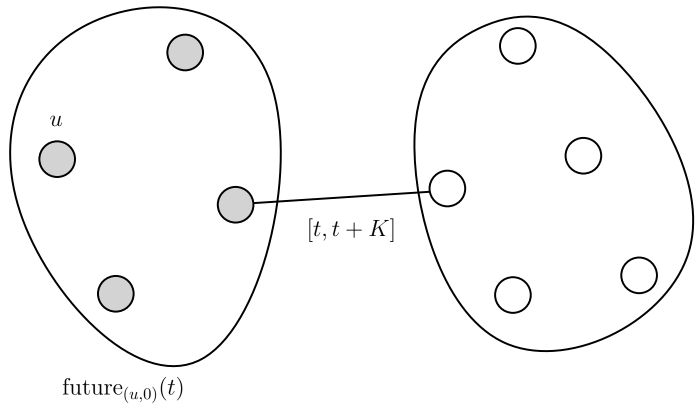

Figure 4: If there are still nodes that have not heard from u, then if K is an upper bound on the oit, in at most K rounds another node will hear from u (by definition of the oit).

those that share (w, r). Similarly, the leader collects all $(v, Infl_v, timestamp)$ triples in its own A_ℓ set. Let $tmax$ denote the maximum timestamp appearing in A_l, that is the maximum time for which the leader knows that some node was influenced by $(\ell, 0)$ at that time. Moreover denote by I the set of nodes that the leader knows to have been influenced by $(\ell, 0)$. Note that I can be extracted from A_ℓ by $I = \{v \in V : \exists u \in V, \exists timestamp, r \in \mathbb{N}$ s.t. $(u, Infl_u, timestamp) \in A_\ell$ and $(v, r) \in Infl_u\}$. If at some round r it holds at the leader that for all $u \in I$ there is a $(u, Infl_u, timestamp) \in A_\ell$ s.t. $timestamp \geq tmax + K$ and $\max_{(w \neq u, r') \in Infl_u}\{r'\} \leq tmax$ then the leader outputs $|I|$ or I depending on whether counting or all-to-all dissemination needs to be solved and halts (it can also easily notify the other nodes to do the same in $K \cdot |I|$ rounds by a simple flooding mechanism and then halt).

The above protocol can be easily made to work without the assumption of a unique leader. The idea is to have all nodes begin as leaders and make all nodes prefer the leader with the smallest id that they have heard of so far. In particular, we can have each node keep an $Infl_{(u,v)}$ only for the smallest v that it has heard of so far. Clearly, in $O(D)$ rounds all nodes will have converged to the node with the smallest id in the network.

Theorem 8. *Protocol* Hear_from_known *solves counting and all-to-all dissemination in $O(D + K)$ rounds by using messages of size $O(n(\log K + \log n))$, in any dynamic network with dynamic diameter D, and with* oit *upper bounded by some K known to the nodes.*

Proof. In time equal to the dynamic diameter D, all nodes must have heard of ℓ. Then in another $D + K$ rounds all nodes must have reported to the leader all the direct outgoing influences that they performed up to time D (nodes that first heard of ℓ by that time) together with the fact that they managed to perform no new influences in the interval $[D, D + K]$. Thus by time $2D + K = O(D + K)$, the leader knows all influences that were ever performed, so no node is missing from its I set, and also knows that all these nodes for K consecutive rounds performed no further influence, thus outputs $|I| = n$ (for counting) or $I = V$ (for all-to-all dissemination) and halts.

Can these termination conditions be satisfied while $|I| < n$, which would result in a wrong decision? Thus, for the sake of contradiction, assume that $tmax$ is the time of the latest influence that the leader is aware of, that $|I| < n$, and that all termination conditions are satisfied. The argument is that if the termination conditions are satisfied then (i) $I = \text{future}_{(\ell,0)}(tmax)$, that is the leader knows precisely those nodes that have been influenced by its initial state up to time $tmax$. Clearly, $I \subseteq \text{future}_{(\ell,0)}(tmax)$ as every node in I has been influenced at most by time $tmax$. We now show that additionally $\text{future}_{(\ell,0)}(tmax) \subseteq I$. If $\text{future}_{(\ell,0)}(tmax) \backslash I \neq \emptyset$, then there must exist some $u \in I$ that has influenced a $v \in \text{future}_{(\ell,0)}(tmax) \backslash I$ at most by time $tmax$ (this follows by observing that $\ell \in I$ and that all influence paths originate from ℓ). But now observe that when the termination conditions are satisfied, for each $u \in I$ the leader knows a $timestamp_u \geq tmax + K$, thus the leader knows all influences that u has performed up to time $tmax$ and it should be aware of the fact that $v \in \text{future}_{(\ell,0)}(tmax)$, i.e. it should hold that $v \in I$, which contradicts the fact that $v \in \text{future}_{(\ell,0)}(tmax)\backslash I$. (ii) The leader knows that in the interval $[tmax, tmax + K]$ no node in $I = \text{future}_{(\ell,0)}(tmax)$ performed a new influence. These result in a contradiction as $|\text{future}_{(\ell,0)}(tmax)| = |I| < n$ and a new influence should have occurred in the interval $[tmax, tmax + K]$ (by the fact that the oit is upper bounded by K).

Optimality follows from the fact that *a node u can know at time t that* $\text{past}_{(u,t)}(0) = V$ *only if* $\text{past}_{(u,t)}(K) = V$. This means that u must have also heard of the K-states of all nodes, which requires $\Theta(K + D)$ rounds in the worst case. If $\text{past}_{(u,t)}(K) \neq V$, then it can be the case that there is some $v \in V\backslash\text{past}_{(u,t)}(K)$ s.t. u has heard v's 0-state but not its K-state. Such a node could be a neighbor of u at round 1 that then moved far away. Again, similarly to Theorem 3, we can have arbitrarily many nodes to have no neighbor until time K (e.g. in the extreme case were oit is equal to K) and then from time K onwards are only connected to node v. As u has not heard from the K-state of v it also cannot have heard of the 0-state of arbitrarily many nodes. □

An interesting improvement is to limit the size of the messages to $O(\log n)$ bits probably by paying some increase in time to termination. We almost achieve this by showing that an improvement of the size of the messages to $O(\log D + \log n)$ bits is possible (note that $O(\log D) = O(\log Kn)$) if we have the leader initiate *individual conversations* with the nodes that it already knows to have been influenced by its initial state. We have already successfully applied a similar technique in [MCS12a], however in a different context. The protocol, that we call *Talk_to_known*, solves counting and all-to-all dissemination in $O(Dn^2 + K)$ rounds by using messages of size $O(\log D + \log n)$, in any dynamic network with dynamic diameter D, and with oit upper bounded by some K known to the nodes.

We now describe the *Talk_to_known* protocol by assuming again for simplicity a unique leader (this, again, is not a necessary assumption).

Protocol *Talk_to_known*. As in *Hear_from_known*, nodes that have been influenced by the initial state of the leader (i.e. $(\ell, 0)$) constantly forward it and whenever a node v manages to deliver it then it stores the id of the recipient node in its local $Infl_v$ set. Nodes send in each round the time of the latest influence (i.e. the latest new influence of a node by $(\ell, 0)$), call it $tmax$, that they know to have been performed so far. Whenever the leader hears of a greater $tmax$ than the one stored in its local memory it reinitializes the process of collecting its future set. By this we mean that it waits K rounds and then starts again from the beginning, talking to the nodes that it has influenced itself, then to the nodes that were influenced by these nodes, and so on. The goal is for the leader to collect precisely the same information as in *Hear_from_known*. In particular, it sorts the nodes that it has influenced itself in ascending order of id and starts with the smallest one, call it v, by initiating a $talk(\ell, v, current\_round)$ message. All nodes forward the most recent $talk$ message (w.r.t. to their timestamp component) that they know so far. Upon receipt of a new $talk(\ell, v, timestamp)$ message (the fact that it is "new" is recognized by the timestamp), v starts sending $Infl_v$ to the leader in packets of size $O(\log n)$, for example a single entry each time, via $talk(v, \ell, current\_round, data\_packet)$ messages. When the leader receives a $talk(v, \ell, timestamp, data\_packet)$ message where $data\_packet = END\_CONV$ (for "END of CONVersation") it knows that it has successfully received the whole $Infl_v$ set and repeats the same process for the next node that it knows to have been already influenced by $(\ell, 0)$ (now also including those that it learned from v). The termination criterion is the same as in *Hear_from_known*.

Theorem 9. *Protocol* Talk_to_known *solves counting and all-to-all dissemination in $O(Dn^2 + K)$ rounds by using messages of size $O(\log D + \log n)$, in any dynamic network with dynamic diameter D, and with* oit *upper bounded by some K known to the nodes.*

Proof. Correctness follows from the correctness of the termination criterion proved in Theorem 8. For the bit complexity we notice that the timestamps and $tmax$ are of size $O(\log D)$ (which may be $O(\log Kn)$ in the worst case). The data packet and the id-components are all of size $O(\log n)$. For the time complexity, clearly, in $O(D)$ rounds the final outgoing influence of $(\ell, 0)$ will have occurred and thus the maximum $tmax$ that will ever appear is obtained by some node. In another $O(D)$ rounds, the leader hears of that $tmax$ and thus reinitializes the process of collecting its future set. In that process and in the worst case the leader must talk to $n-1$ nodes each believing that it performed $n-1$ deliveries (this is because in the worst case it can hold that any new node is concurrently influenced by all nodes that were already influenced and in the end all nodes claim that they have influenced all other nodes) thus, in total, it has to wait for $O(n^2)$ data packets each taking $O(D)$ rounds to arrive. The K in the bound is from the fact that the leader waits K rounds after reinitializing in order to allow nodes to also report whether they performed any new assignments in the $[tmax, tmax + K]$ interval. $\square$

8. Conclusions

To the best of our knowledge, this is the first study of worst-case dynamic networks that are free of any connectivity assumption about their instances. To enable a quantitative study we proposed some novel generic metrics that capture the speed of information propagation in a dynamic network. We proved that fast dissemination and computation are possible even under continuous disconnectivity. In particular, we presented optimal termination conditions and protocols based on them for the fundamental counting and all-to-all token dissemination problems.

There are many open problems and promising research directions related to this work. We would like to improve the existing lower and upper bounds for counting and information dissemination. The best lower bound known for k-token dissemination is $\Omega(nk/\log n)$ (even for centralized algorithms on networks with connected instances and messages of size $O(\log n)$) [DPR+13] while our upper bound is $O(Dn^2 + K)$ (for messages of size $O(\log D + \log n)$). Techniques from [HCAM12] or related ones may be applicable to achieve quick token dissemination. It would be also important to refine the metrics proposed in this work so that they become more informative. For example, the oit metric, in its present form, just counts the time needed for another outgoing influence to occur. It would be useful to define a metric that counts the number of new nodes that become influenced per round which would be more informative w.r.t. the speed of information spreading. An interesting variation of our metrics, which is due to a reviewer of this article, is to define them in an amortized way. For example, $\mathsf{oit} = k$ can alternatively be defined as $h = |\text{future}_{(u,t)}(t')| \geq \min\{h + \lfloor (t'-t)/k \rfloor, n\}$ and this weaker definition may be of high practical value as, this way, the class of dynamic graphs having $\mathsf{oit} = k$ (which now can also be fractional) will be much larger. An asynchronous communication model in which nodes can broadcast when there are new neighbors would be a very natural extension of the synchronous model that we studied in this work. Note that in our work (and all previous work on the subject) information dissemination is only guaranteed under continuous broadcasting. How can the number of redundant transmissions be reduced in order to improve communication efficiency? Is there a way to exploit visibility to this end? Does predictability help (i.e. some knowledge of the future)? Finally, randomization will be certainly valuable in constructing fast and symmetry-free protocols. We strongly believe that these and other known open questions and research directions will motivate the further growth of this emerging field.

Acknowledgements

We would like to thank some anonymous reviewers for carefully reading the manuscript and for providing valuable comments that have helped us to improve our work substantially.

References

[AAD+06] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, pages 235–253, March 2006.

[AAER07] D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*, 20[4]:279–304, November 2007.

[AKL08] C. Avin, M. Koucký, and Z. Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *Proceedings of the 35th international colloquium on Automata, Languages and Programming (ICALP), Part I*, pages 121–132, Berlin, Heidelberg, 2008. Springer-Verlag.

[APRU12] J. Augustine, G. Pandurangan, P. Robinson, and E. Upfal. Towards robust and efficient computation in dynamic peer-to-peer networks. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 551–569. SIAM, 2012.

[AW04] H. Attiya and J. Welch. *Distributed computing: fundamentals, simulations, and advanced topics*, volume 19. Wiley-Interscience, 2004.

[BCF09] H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. In *Proceedings of the 28th ACM symposium on Principles of distributed computing (PODC)*, pages 260–269. ACM, 2009.

[Bol98] B. Bollobás. *Modern Graph Theory*. Springer, corrected edition, July 1998.

[CFQS12] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27[5]:387–408, 2012.

[CMM+08] A. E. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time in edge-markovian dynamic graphs. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing (PODC)*, pages 213–222, New York, NY, USA, 2008. ACM.

[CMN+11] I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P. G. Spirakis. Passively mobile communicating machines that use restricted space. *Theor. Comput. Sci.*, 412[46]:6469–6483, October 2011.

[Dol00] S. Dolev. *Self-stabilization*. MIT Press, Cambridge, MA, USA, 2000.

[DPR+13] C. Dutta, G. Pandurangan, R. Rajaraman, Z. Sun, and E. Viola. On the complexity of information spreading in dynamic networks. In *Proc. of the 24th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2013.

[Hae11] B. Haeupler. Analyzing network coding gossip made easy. In *Proceedings of the 43rd annual ACM symposium on Theory of computing (STOC)*, pages 293–302. ACM, 2011.

[HCAM12] B. Haeupler, A. Cohen, C. Avin, and M. Médard. Network coded gossip with correlated data. *CoRR*, abs/1202.1801, 2012.

[KKK00] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing (STOC)*, pages 504–513, 2000.

[KLO10] F. Kuhn, N. Lynch, and R. Oshman. Distributed computation in dynamic networks. In *Proceedings of the 42nd ACM symposium on Theory of computing (STOC)*, pages 513–522, New York, NY, USA, 2010. ACM.

[KMO11] F. Kuhn, Y. Moses, and R. Oshman. Coordinated consensus in dynamic networks. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing (PODC)*, pages 1–10, 2011.

[KO11] F. Kuhn and R. Oshman. Dynamic networks: models and algorithms. *SIGACT News*, 42:82–96, March 2011. Distributed Computing Column, Editor: Idit Keidar.

[Lam78] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21[7]:558–565, 1978.

[Lyn96] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann; 1st edition, 1996.

[MCS11a] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Mediated population protocols. *Theor. Comput. Sci.*, 412[22]:2434–2450, May 2011.

[MCS11b] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. N. A. Lynch (Ed), Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool, 2011.

[MCS12a] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Brief announcement: Naming and counting in anonymous unknown dynamic networks. In *Proceedings of the 26th international conference on Distributed Computing (DISC)*, pages 437–438, Berlin, Heidelberg, 2012. Springer-Verlag.

[MCS12b] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Causality, influence, and computation in possibly disconnected synchronous dynamic networks. In *16th International Conference on Principles of Distributed Systems (OPODIS)*, pages 269–283, 2012.

[MMCS13] G. B. Mertzios, O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Temporal network optimization subject to connectivity constraints. In *40th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 7966 of *Lecture Notes in Computer Science*, pages 663–674. Springer, July 2013.

[OW05] R. O'Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proceedings of the 2005 joint workshop on Foundations of mobile computing (DIALM-POMC)*, pages 104–110, 2005.

[Sch02] C. Scheideler. Models and techniques for communication in dynamic networks. In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 27–49, 2002.

[Soi09] A. Soifer. *The Mathematical Coloring Book: Mathematics of Coloring and the Colorful Life of its Creators*. Springer; 1st edition, 2009.