

Naming and Counting in Anonymous Unknown Dynamic Networks^{*,**}

Othon Michail¹, Ioannis Chatzigiannakis¹, and Paul G. Spirakis^{1,2}

¹ Computer Technology Institute & Press “Diophantus” (CTI), Patras, Greece

² Department of Computer Science, University of Liverpool, UK

Email: {michailo, ichtatz, spirakis}@cti.gr

Abstract. In this work, we study the fundamental naming and counting problems (and some variations) in networks that are anonymous, unknown, and possibly dynamic. In *counting*, nodes must determine the size of the network n and in *naming* they must end up with unique identities. By *anonymous* we mean that all nodes begin from identical states apart possibly from a unique leader node and by *unknown* that nodes have no a priori knowledge of the network (apart from some minimal knowledge when necessary) including ignorance of n . Network dynamicity is modeled by the 1-interval connectivity model [KLO10], in which communication is synchronous and a (worst-case) adversary chooses the edges of every round subject to the condition that each instance is connected. We first focus on static networks with broadcast where we prove that, without a leader, counting is impossible to solve and that naming is impossible to solve even with a leader and even if nodes know n . These impossibilities carry over to dynamic networks as well. We also show that a unique leader suffices in order to solve counting in linear time. Then we focus on dynamic networks with broadcast. We conjecture that dynamicity renders nontrivial computation impossible. In view of this, we let the nodes know an upper bound on the maximum degree that will ever appear and show that in this case the nodes can obtain an upper bound on n . Finally, we replace broadcast with *one-to-each*, in which a node may send a different message to each of its neighbors. Interestingly, this natural variation is proved to be computationally equivalent to a full-knowledge model, in which unique names exist and the size of the network is known.

1 Introduction

Distributed computing systems are more and more becoming dynamic. The static and relatively stable models of computation can no longer represent the

* 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 brief announcement of this work has appeared in [MCS12a].

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 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, CFQS11]), in which each edge has an associated set of time-labels indicating availability times. Even fundamental properties of classical graphs do not carry over to their temporal counterparts. See, for example, [KKK00] for a violation of Menger's theorem, [MMCS13] for a valid reformulation of Menger's theorem and the definition of several cost optimization metrics for temporal networks, and [AKL08] for the unsuitability of the standard network diameter metric.

In this work, we adopt as our dynamic network model the *1-interval connectivity* model of Kuhn *et al.* [KLO10] building upon previous work of O'Dell and Wattenhofer [OW05]. In this model, nodes proceed in *synchronous rounds* and communicate by *interchanging messages*. Message transmission is *broadcast* in which, in every round, each node issues a single message to be delivered to all its neighbors. In this model, the network may change arbitrarily from round to round subject to the condition that in each round the network is connected. We only consider deterministic algorithms.

We focus on networks in which nodes are initially identical and, unless necessary, do not have any information about the network. In any case, nodes do not know the size n of the network. By *identical* we mean that they do not have unique identities (ids) and execute identical programs. So, this is some sort of minimal reliable distributed system, like, for example, a collection of particularly cheap and bulk-produced wireless sensor nodes. Nodes may execute the same program, because it is too costly to program them individually and their lack of ids may be due to the fact that ids require customization beyond the capabilities of mass production [AFR06]. Our only assumption is the existence of

a unique leader that introduces some symmetry breaking. To further break the symmetry introduced by broadcast message transmission and in order to solve naming in dynamic networks, we allow to the nodes to send a different message to each one of their neighbors.

2 Related Work

Distributed systems with worst-case dynamicity were first studied in [OW05]. Their 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 allowed nodes detect local neighborhood changes. They studied the *flooding* and *routing* problems in this setting and among others provided 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.

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], 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 $O(n^2)$ rounds using $O(\log n)$ bits per message. The requirement for continuous connectivity was first dropped in [MCS12b]. That work proposed a set of metrics for capturing the speed of information spreading in a dynamic network that may be disconnected at every instance and efficient algorithms were developed. Some recent works, e.g. [Hae11], present information spreading algorithms in worst-case dynamic networks based on *network coding*.

The question concerning which problems can be solved by a distributed system when all processors use the same algorithm and start from the same state has a long story with its roots dating back to the seminal work of Angluin [Ang80], who investigated the problem of establishing a “center”. Further investigation led to the classification of computable functions [YK96,ASW88]. [BV99] removed the, until then, standard assumption of knowing the network size n and provided characterizations of the relations that can be computed with arbitrary knowledge. Other well-known studies on unknown networks have dealt with the problems of robot-exploration and map-drawing of an unknown graph [DP90,AH00] and on information dissemination [AGVP90]. Fraigniaud *et al.* [FPPP00] assumed a unique leader in order to break symmetry and assign short labels as fast as possible. To circumvent the further symmetry introduced by broadcast message transmission they also studied other natural message transmission models as sending only one message to a single neighbor. Recently, and independently of our work, Chalopin *et al.* [CMM12] have studied the problem of naming anonymous networks in the context of snapshot computation. Finally, Aspnes *et al.* [AFR06] studied the relative powers of reliable anonymous distributed systems with different communication mechanisms: anonymous broad-

cast, read-write registers, or read-write registers plus additional shared-memory objects.

3 Contribution

We begin, in Section 4, by formally describing our distributed models. In Section 5, we formally define the problems under consideration, that is, naming, counting and some variations of these. Our study begins, in Section 6, from static networks with broadcast. The reason for considering static networks is to arrive at some impossibility results that also carry over to dynamic networks, as a static network is a special case of a dynamic network. In particular, we prove that naming is impossible to solve under these assumptions even if a unique leader exists and even if all nodes know n . Then we prove that without a leader also counting is impossible to solve and naturally, in the sequel, we assume the existence of a unique leader. We provide an algorithm based on the *eccentricity* of the leader (greatest distance of a node from the leader) that solves counting in linear time (inspired by the findings in [FPPP00]). Then, in Section 7, we move on to dynamic networks with broadcast. We begin with a conjecture essentially stating that dynamicity renders nontrivial computations impossible even in the presence of a unique leader.³ In view of this, we allow the nodes some minimal initial knowledge, which is an upper bound on the maximum degree that any instance will ever have. This could for example be some natural constraint on the capacity of the network. We provide a protocol that exploits this information to compute an upper bound on the size of the network. However, w.r.t. naming, the strong impossibility from Section 6 still persists (after all, knowledge of n does not help in labeling the nodes). To circumvent this, in Section 8, we relax our message transmission model to *one-to-each* that allows each node to send a different message to each one of its neighbors. This is an alternative communication model that has been considered in several important works, like [Hae11], however in different contexts than ours. This further symmetry breaking, though minimal, allows us, by exploiting a leader, to uniquely label the nodes. By this, we establish that this model is equivalent to a full-knowledge model in which unique names exist and the size of the network is known. To arrive at this result, we provide four distinct naming protocols each with its own incremental value. The first presents how to assign ids in a fair context in which the leader will eventually meet every other node. The second improves on the first by allowing all nodes to assign ids in a context where no one is guaranteed to meet everybody else, but where connectivity guarantees progress. Both these are correct stabilizing solutions that do not guarantee termination. Then we provide a third protocol that builds upon the first two and manages to assign unique ids in 1-interval connected graphs while terminating in linear time. As its drawback is that messages may

³ By *nontrivial computation* we mean the ability to decide any language L on input assignments s.t. $L \neq \Sigma^*$ and $L \neq \emptyset$, where input symbols are chosen from some alphabet Σ . For example, deciding the existence of any symbol in the input is considered nontrivial.

be $\Omega(n^2)$ bit long, we refine it to a more involved fourth protocol that reduces the bits per message to $\Theta(\log n)$ by only paying a small increase in termination time.

4 Preliminaries

4.1 The models

A *dynamic network* is modeled by a *dynamic graph* $G = (V, E)$, where V is a set of n nodes (or processors) and $E : \mathbb{N} \rightarrow \mathcal{P}(E')$, where $E' = \{\{u, v\} : u, v \in V\}$, (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' . Intuitively, a dynamic graph G is an infinite sequence $G(1), G(2), \dots$ 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.

A dynamic graph/network $G = (V, E)$ is said to be *1-interval connected*, if, for all $r \in \mathbb{N}$, the static graph $G(r)$ is connected [KLO10]. Note that this allows the network to change arbitrarily from round to round always subject to the condition that it remains connected. In this work, we focus on 1-interval connected dynamic networks which also implies that we deal with connected networks in the static-network case.

Nodes in V are *anonymous* that is they do not initially have any ids and they do not know the topology or the size of the network, apart from some minimal knowledge when necessary (i.e. we say that *the network is unknown*). However, nodes have unlimited local storage. In several cases, and in order to break symmetry, we may assume a unique *leader node* (or *source*) l . If this is the case, then we assume that l starts from a unique initial state l_0 (e.g. 0) while all other nodes start from the same initial state q_0 (e.g. \perp). All nodes but the leader execute identical programs.

Communication is *synchronous message passing* [Lyn96,AW04], meaning that it is executed in discrete rounds controlled by a global clock that is available to the nodes and that nodes communicate by sending and receiving messages. Thus all nodes have access to the current round number via a local variable that we usually denote by r . We consider two different models of message transmission. One 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 $N_u(r) = \{v : \{u, v\} \in E(r)\}$. The other is *one-to-each* in which a different message $m_{(u,i)}(r)$, $1 \leq i \leq d_u(r)$, where $d_u(r) := |N_u(r)|$ is the degree of u in round r , may be generated for each neighbor v_i . 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. In the one-to-each message transmission model we additionally assume that the adversary also reveals to each node u a set of locally unique edge-labels $1, 2, \dots, d_u(r)$, one for each of the edges currently incident to it. Note that these labels can be reselected arbitrarily in each

round so that a node cannot infer what the internal state of a neighbor is based solely on the corresponding local edge-name. Then each node transitions to a new state based on its internal state (containing the messages received in the previous round) and generates its messages for the current round: in anonymous broadcast a single message is generated and in one-to-each a different message is generated for each neighbor of a node. Note that, in both models, a node does not have any information about the internal state of its neighbors when generating its messages. Deterministic algorithms are only based on the current internal state to generate messages. This implies that the adversary can infer the messages that will be generated in the current round before choosing the edges. Messages are then delivered to the corresponding neighbors. In one-to-each, we assume that each message m_i received by some node u is accompanied with u 's local label i of the corresponding edge, so that a node can associate a message sent through edge i with a message received from edge i . These messages will be processed by the nodes in the subsequent round so we typically begin rounds with a “receive” command referring to the messages received in the previous round. Then the next round begins.

4.2 Causal Influence

Probably the most important notion associated with a dynamic 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 use $(u, r) \rightsquigarrow (v, r')$ to denote the fact that node u 's state in round r (r -state of u) 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 .*

A very important aspect of 1-interval connectivity, that will be invoked in all our proof arguments in dynamic networks, is that it 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 $V_2 = V \setminus V_1$ 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 $r \geq 0$ we have*

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

5 Problem Definitions

k -labeling. An algorithm is said to solve the k -labeling problem if whenever it is executed on a network comprising n nodes each node u eventually terminates and outputs a *label* (or *name* or *id*) id_u so that $|\{id_u : u \in V\}| \geq k$.

Naming. The naming problem is a special case of the k -labeling problem in which it must additionally hold that $k = n$. This, in turn, implies that $id_u \neq id_v$ for all distinct $u, v \in V$ (so, unique labels are required for the nodes).

Minimal (Consecutive) Naming. It is a special case of naming in which it must additionally hold that the n nodes output the labels $\{0, 1, \dots, n - 1\}$.

Counting Upper Bound. Nodes must determine an upper bound k on the network size n .

Counting. A special case of counting upper bound in which it must hold that $k = n$.

6 Static Networks with Broadcast

We here assume that the network is described by a static graph $G = (V, E)$, where $E \subseteq \{\{u, v\} : u, v \in V\}$. Moreover, the message transmission model is broadcast, that is, in every round, each node u generates a single message to be delivered to all its neighbors. Note that any impossibility result established for static networks is also valid for dynamic networks as a static network is a special case of a dynamic network.

First of all, note that if all nodes start from the same initial state then, if we restrict ourselves to deterministic algorithms, naming is impossible to solve in general static networks, even if nodes know n . The reason is that in the worst-case they may be arranged in a ring (in which each node has precisely 2 neighbors) and it is a well-known fact [Ang80,Lyn96,AW04] that, in this case, in every round r , all nodes are in identical states.

We show now that impossibility persists even if we allow a unique leader and even if nodes have complete knowledge of the network.

Theorem 1. *Naming is impossible to solve by deterministic algorithms in general anonymous (static) networks with broadcast even in the presence of a leader and even if nodes have complete knowledge of the network.*

Proof. Consider a star graph with the leader in the center. □

An obvious generalization is that, under the same assumptions as in the statement of the theorem, it is impossible to solve k -labeling for any $k \geq 3$. In the full paper, we also provide some thoughts on a degree-based labeling.

We now turn our attention to the simpler counting problem. First we establish the necessity of assuming a unique leader.

Theorem 2. *Without a leader, counting is impossible to solve by deterministic algorithms in general anonymous networks with broadcast.*

Proof. If some algorithm counts in k rounds the n nodes of a static ring, then it fails on a ring of $k + 1$ nodes. \square

In view of Theorem 2, we assume again a unique leader in order to solve counting. Recall that the *eccentricity* of a node u is defined as the greatest geodesic distance between u and v , over all $v \in V \setminus \{u\}$, where “distance” is equivalent to “shortest path”. We first describe a protocol *Leader_Eccentricity* (inspired by the *Wake&Label* set of algorithms of [FPPP00]) that assigns to every node a label equal to its distance from the leader and then we exploit this to solve counting. We assume that all nodes have access to the current round number via a variable r .

Protocol *Leader_Eccentricity*. The leader begins with $label \leftarrow 0$ and $max_asgnd \leftarrow 0$ and all other nodes with $label \leftarrow \perp$. In the first round, the leader broadcasts an *assign* (1) message. Upon reception of an *assign* (i) message, a node that has $label = \perp$ sets $label \leftarrow i$ and broadcasts to its neighbors an *assign* ($i + 1$) message and an *ack* (i) message. Upon reception of an *ack* (i) message, a node with $label \neq \perp$ and $label < i$ broadcasts it. Upon reception of an *ack* (i) message, the leader sets $max_asgnd \leftarrow i$ and if $r > 2 \cdot (max_asgnd + 1)$ then it broadcasts a *halt* message, outputs its label, and halts. Upon reception of a *halt* message, a node broadcasts *halt*, outputs its label, and halts.

Theorem 3. *In *Leader_Eccentricity* nodes output $\epsilon + 1$ distinct labels where ϵ is the eccentricity of the leader. In particular, every node outputs its distance from the leader.*

Proof. At time 2, nodes at distance 1 from the leader receive *assign* (1) and set their label to 1. By induction on distance, nodes at distance i get label i at round $i + 1$. In the same round, they send an *ack* that must arrive at the leader at round $2i + 1$. If not then there is no node at distance i . \square

We now use *Leader_Eccentricity* to solve counting in anonymous unknown static networks with a leader. We additionally assume that at the end of the *Leader_Eccentricity* process each node u knows the number of neighbors $up(u) = |\{\{v, u\} \in E : label(v) = label(u) - 1\}|$ it has to its upper level (it can store this during the *Leader_Eccentricity* process by counting the number of *assign* messages that arrived at it from its upper level neighbors). Moreover, we assume that all nodes know the leader’s eccentricity ϵ (just have the leader include max_asgnd in its *halt* message). Finally, let, for simplicity, the first round just after the completion of the above process be round $r = 1$. For this, we just need all nodes to end concurrently the *Leader_Eccentricity* process. This is done by having node with label i that receives or creates (this is true for the leader) a *halt* message in round r halt in round $(r + max_asgnd - i)$. Then the nodes just reset their round counters.

Protocol *Anonymous_Counting*. Nodes first execute the modified *Leader_Eccentricity*. When $\epsilon - r + 1 = label(u)$, a non-leader node u receives a

possibly empty (in case of no lower-level neighbors) set of $partial_count_i$ ($rval_i$) messages and broadcasts a $partial_count$ ($(1 + \sum_i rval_i)/up(u)$) message. When $r = \epsilon + 1$, the leader receives a set of $partial_count_i$ ($rval_i$) messages, sets $count \leftarrow 1 + \sum_i rval_i$, broadcasts a $halt$ ($count$) message, outputs $count$, and halts. When a non-leader u receives a $halt$ ($count$) message, it outputs $count$ and halts.

For a given round r denote by $rval_i(u)$ the i th message received by node u .

Theorem 4. *Anonymous Counting solves the counting problem in anonymous static networks with broadcast under the assumption of a unique leader. All nodes terminate in $O(n)$ rounds and use messages of size $O(\log n)$.*

Proof. By induction on the round number r , in the beginning of round $r \geq 2$, it holds that $\sum_{u:label(u)=\epsilon-r+1} (1 + \sum_i rval_i(u)) = |\{u : label(u) \geq \epsilon - r + 1\}|$. Clearly, in round $\epsilon + 1$ it holds that $count = 1 + \sum_i rval_i(leader) = |\{u : label(u) \geq 0\}| = n$. \square

7 Dynamic Networks with Broadcast

We now turn our attention to the more general case of 1-interval connected dynamic networks with broadcast. We begin with a conjecture stating that dynamicity renders nontrivial computation impossible (see also [OW05] for a similar conjecture in a quite different setting). Then we naturally strengthen the model to allow some computation.

Conjecture 1. It is impossible to compute (even with a leader) the predicate $N_a \geq 1$, that is “exists an a in the input”, in general anonymous unknown dynamic networks with broadcast.

In view of Theorem 1, which establishes that we cannot name the nodes of a static, and thus also of a dynamic, network if broadcast communication is assumed, and of the above conjecture, implying that in dynamic networks we cannot count even with a leader ⁴, we start strengthening our initial model.

Let us assume that there is a unique leader l that knows an upper bound d on maximum degree ever to appear in the dynamic network, that is $d \geq \max_{u \in V, r \in \mathbb{N}} \{d_u(r)\}$. We keep the broadcast message transmission.

Note first that impossibility of naming persists. However, we show that obtaining an upper bound on the size of the network now becomes possible, though exponential in the worst case.

Protocol Degree Counting. The leader stores in d the maximum degree that will ever appear and begins with $label \leftarrow 0$, $count \leftarrow 1$, $latest_event \leftarrow 0$, $max_label \leftarrow 0$, and $r \leftarrow 0$ while all other nodes begin with $label \leftarrow \perp$, $count \leftarrow 0$, and $r \leftarrow 0$. In the beginning of each round each node increments by one its round

⁴ This is implied because, if we could count, we could have a node wait at most $n - 1$ rounds to hear of an a (provided that all nodes that have heard of an a forward it) and if no, reject.

counter r . The leader in each round r broadcasts *assign* (r). Upon reception of an *assign* (r_label) message, a node with $label = \perp$ sets $label \leftarrow r_label$ and from now in each round r broadcasts *assign* (r) and *my_label* ($label$). A node with $label = \perp$ that did not receive an *assign* message sends an *unassigned* (r) message. All nodes continuously broadcast the maximum *my_label* and *unassigned* messages that they have received so far. Upon reception of an *unassigned* (i) message, the leader, if $i > latest_event$, it sets $count \leftarrow 1$ and, for $k = 1, \dots, i$, $count \leftarrow count + d \cdot count$, $max_label \leftarrow i$, and $latest_event \leftarrow r$ and upon reception of a *my_label* (j) message, if $j > max_label$, it sets $count \leftarrow 1$ and, for $k = 1, \dots, j$, $count \leftarrow count + d \cdot count$, $latest_event \leftarrow r$, and $max_label \leftarrow j$ (if receives both i, j it does it for $\max\{i, j\}$). When it holds that $r > count + latest_event - 1$ (eventually occurs) then the leader broadcasts a *halt* ($count$) message for $count$ rounds and then outputs $count$ and halts. Each node that receives a *halt* (r_count) message, sets $count \leftarrow r_count$, broadcasts a *halt* ($count$) message for $count$ rounds and then outputs $count$ and halts.

Theorem 5. *Degree-Counting solves the counting upper bound problem in anonymous dynamic networks with broadcast under the assumption of a unique leader. The obtained upper bound is $O(d^n)$ (in the worst case).*

Proof. In the first round, the leader assigns the label 1 to its neighbors and obtains an *unassigned* (1) message from each one of them. So, it sets $count \leftarrow (d + 1)$ (in fact, note that in the first step it can simply set $count \leftarrow d_u(1) + 1$, but this is minor), $latest_event \leftarrow 1$, and $max_label \leftarrow 1$. Now, if there are further nodes, at most by round $count + latest_event - 1$ it must have received an *unassigned* (i) message with $i > latest_event$ or a *my_label* (j) with $j > max_label$. Note that the reception of an *unassigned* (i) message implies that at least $i + 1$ distinct labels have been assigned because as long as there are unlabeled nodes one new label is assigned in each round to at least one node (this is implied by Lemma 1 and the fact that all nodes with labels constantly assign new labels). Initially, one node (the leader) assigned to at most d nodes label 1. Then the $d + 1$ labeled nodes assigned to at most $(d + 1)d$ unlabeled nodes the label 2, totalling $(d + 1) + (d + 1)d$, and so on.

In the worst-case, each label in $\{0, 1, \dots, n - 1\}$ is assigned to precisely one node (e.g., consider a static line with the leader in the one endpoint). In this case the nodes count $O(d^n)$. \square

We point out that if nodes have access to more drastic initial knowledge such as an upper bound e on the *maximum expansion*, defined as $\max_{u,r,r'} \{|\text{future}_{u,r}(r' + 1)| - |\text{future}_{u,r}(r')|\}$ (maximum number of concurrent new influences), where $\text{future}_{(u,r)}(r') := \{v \in V : (u, r) \rightsquigarrow (v, r')\}$, for $r \leq r'$, then essentially the same protocol as above provides an $O(n \cdot e)$ upper bound.

8 Dynamic Networks with One-to-Each

The result of Theorem 1, in the light of (a) the conjecture of Section 7, and (b) the assumption of a broadcast message transmission model, clearly indicates

that nontrivial computations in anonymous unknown dynamic networks are impossible even under the assumption of a unique leader. We now relax these assumptions so that we can state a correct naming protocol. We start by relaxing the assumption of a broadcast message transmission medium by offering to nodes access to a *one-to-each message transmission mechanism*. We also assume a unique leader as without it, even under one-to-each, impossibility persists.

1st Version - Protocol *Fair*

We now present protocol *Fair* in which the unique leader assigns distinct labels to each node of the network. The labels assigned are tuples (r, h, i) , where r is the round during which the label was assigned, h is the label of the leader node and i is a unique number assigned by the leader. The labels can be uniquely ordered first by r , then by h and finally by i (in ascending order).

Each node maintains the following local variables: *clock*, for counting the rounds of execution of the protocol (due to synchrony), *label*, for storing the label assigned by the leader, *state*, for storing the local state that can be set to $\{anonymous, named, leader\}$, and *counter*, for storing the number of labels generated. All nodes are initialized to $clock \leftarrow 0$, $id \leftarrow (0, \perp, \perp)$, $state \leftarrow anonymous$, and $counter \leftarrow 0$ except from the leader that is initialized to $clock \leftarrow 0$, $id \leftarrow (0, 1, 1)$, $state \leftarrow leader$, and $counter \leftarrow 1$.

Each turn the leader u consults the one-to-each transmission mechanism and identifies a set of locally unique edge-labels $1, 2, \dots, d(u)$, one for each of the edges incident to it.⁵ The leader iterates the edge-label set and transmits to each neighboring node a different message m_i , $1 \leq i \leq d(u)$ that contains the unique label $(clock, label, counter + i)$. When the transmission is complete, it increases the variable *counter* by $d(u)$. All the other nodes of the network do not transmit any messages (or transmit a null message).

All nodes under $state = anonymous$, upon receiving a (non-null) message set the local *label* to the contents of the message and change *state* to *named*. All the other nodes of the network simply ignore all the messages received.

Recall that a naming assignment is correct if *all nodes* are assigned *different* labels. It is clear that *Fair* is a non-terminating correct protocol, given the following *fairness assumption*: the leader node at some point has become directly connected with each other node of the network (i.e., eventually meets all nodes).

Lemma 2. *With one-to-each transmission, under the fairness assumption, and in the presence of a unique leader, protocol Fair eventually computes a unique assignment for all the nodes in any anonymous unknown dynamic network.*

2nd Version - Protocol *Delegate*

We now proceed by presenting a stronger protocol *Delegate* (based on *Fair*) that is correct even without the fairness assumption. To achieve correctness

⁵ Recall that these edge-labels can be reselected arbitrarily in each round (even if the neighbors remain the same) by the adversary so that a node cannot infer the internal state of a neighbor, based solely on the corresponding local edge-name.

the leader node delegates the role of assignment of labels to all the nodes that it encounters. Thus, without loss of generality, even if the leader does not encounter all other nodes of the network, due to the *connectivity property*, all nodes will eventually hear from the leader. Therefore, all nodes will either receive a unique label from the leader or from another labeled node. The uniqueness among the labels generated is guaranteed since each label can be traced back to the node that issued it using the h parameter.

In *Delegate* the nodes maintain the same variables as in *Fair*. Each turn the leader performs the same actions as in *Fair*. Also similarly to *Fair*, each node that is in $state = anonymous$ does not transmit any message (or transmits a null message if message transmission is compulsory). Each node u that is in $state = named$ performs similar actions as the leader node and transmits to each edge-label i a message containing the unique label $(clock_u, label_u, counter_u + i)$ and then increases the variable $counter_u$ by $d(u)$. All nodes under $state = anonymous$, upon receiving one or more (non-null) messages that contain a label, select the message that contains the lowest label (i.e., the one with the lowest h parameter) and set the local $label$ to the contents of the message and change $state$ to $named$.

Lemma 3. *With one-to-each transmission, and in the presence of a unique leader, protocol Delegate correctly computes a unique assignment for all the nodes in any anonymous unknown dynamic network.*

3rd Version - Protocol *Dynamic_Naming* (terminating)

The protocols *Fair* and *Delegate* compute a correct naming assignment (based on different assumptions) but do not terminate. Essentially the nodes continue to transmit labels forever. We now present protocol *Dynamic_Naming* (based on *Delegate, Fair*) that manages to terminate.

Dynamic_Naming is an $O(n)$ -time protocol that assigns unique ids to the nodes and informs them of n . As usual, there is a unique leader l with id 0 while all other nodes have id \perp .

The idea here is as follows. Similarly to *Delegate*, all nodes that have obtained an id assign ids and these ids are guaranteed to be unique. Additionally to *Delegate*, we have nodes that have obtained an id to acknowledge their id to the leader. Thus, all nodes send their ids and all nodes continuously forward the received ids so that they eventually arrive at the leader (simple flooding mechanism). So, at some round r , the leader knows a set of assigned ids $K(r)$. We describe now the termination criterion. If $|K(r)| \neq |V|$ then in at most $|K(r)|$ additional rounds the leader must hear (be causally influenced) from a node outside $K(r)$ (to see why, see Lemma 1). Such a node, either has an id that the leader first hears of, or has no id yet. In the first case, the leader updates $K(r)$ and in the second waits until it hears of a new id (which is guaranteed to appear in the future). On the other hand, if $|K(r)| = |V|$ no new info will ever arrive at the leader in the future and the leader may terminate after the $|K(r)|$ -round waiting period ellapses.

Protocol *Dynamic_Naming*. Initially, every node has three variables $count \leftarrow 0$, $acks \leftarrow \emptyset$, and $latest_unassigned \leftarrow 0$ and the leader additionally has $latest_new \leftarrow 0$, $time_bound \leftarrow 1$, and $known_ids \leftarrow \{0\}$. A node with $id \neq \perp$ for $1 \leq i \leq k$ sends $assign(id, count + i)$ message to its i th neighbor and sets $count \leftarrow count + k$. In the first round, the leader additionally sets $known_ids \leftarrow \{0, (0, 1), (0, 2), \dots, (0, k)\}$, $latest_new \leftarrow 1$, and $time_bound \leftarrow 1 + |known_ids|$. Upon receipt of l $assign$ messages (rid_j), a node with $id = \perp$ sets $id \leftarrow \min_j\{rid_j\}$ (in number of bits), $acks \leftarrow acks \cup id$, sends an $ack(acks)$ message to all its k current neighbors, for $1 \leq i \leq k$ sends $assign(id, count + i)$ message to its i th neighbor, and sets $count \leftarrow count + k$. Upon receipt of l ack messages ($acks_j$), a nonleader sets $acks \leftarrow acks \cup (\bigcup_j acks_j)$ and sends $ack(acks)$. A node with $id = \perp$ sends $unassigned(current_round)$. Upon receipt of $l \geq 0$ $unassigned$ messages (val_j), a node with $id \notin \{0, \perp\}$ sets $latest_unassigned \leftarrow \max\{latest_unassigned, \max_j\{val_j\}\}$ and sends $unassigned(latest_unassigned)$. Upon receipt of l ack messages ($acks_j$), the leader if $(\bigcup_j acks_j) \setminus known_ids \neq \emptyset$ sets $known_ids \leftarrow known_ids \cup (\bigcup_j acks_j)$, $latest_new \leftarrow current_round$ and $time_bound \leftarrow current_round + |known_ids|$ and upon receipt of l $unassigned$ messages (val_j), it sets $latest_unassigned \leftarrow \max\{latest_unassigned, \max_j\{val_j\}\}$. If, at some round r , it holds at the leader that $r > time_bound$ and $latest_unassigned < latest_new$, the leader sends a $halt(|known_ids|)$ message for $|known_ids| - 1$ rounds and then outputs id and halts. Any node that receives a $halt(n)$ message, sends $halt(n)$ for $n - 2$ rounds and then outputs id and halts.

Denote by $S(r) = \{v \in V : (l, 0) \rightsquigarrow (v, r)\}$ the set of nodes that have obtained an id at round r and by $K(r)$ those nodes in $S(r)$ whose id is known by the leader at round r , that is $K(r) = \{u \in V : \exists r' \text{ s.t. } u \in S(r') \text{ and } (u, r') \rightsquigarrow (l, r)\}$.

Theorem 6. *Dynamic_Naming solves the naming problem in anonymous unknown dynamic networks under the assumptions of one-to-each message transmission and of a unique leader. All nodes terminate in $O(n)$ rounds and use messages of size $\Theta(n^2)$.*

Proof. Unique names are guaranteed as in *Delegate*. Termination is as follows. Clearly, if $V \setminus K(r) \neq \emptyset$, either $|K(r + |K(r)|)| \geq |K(r)| + 1$ or $(u, r) \rightsquigarrow (l, r + |K(r)|)$ for some $u \in V \setminus S(r)$. The former is recognized by the leader by the arrival of a new id and the latter by the arrival of an *unassigned(timestamp)* message, where $timestamp \geq r$. On the other hand, if $K(r) = V$ then $|K(r + |K(r)|)| = |K(r)|$ and $\nexists u \in V \setminus S(r)$ s.t. $(u, r) \rightsquigarrow (l, r + |K(r)|)$ as $V \setminus S(r) = \emptyset$. Finally, note that connectivity implies that $|S(r + 1)| \geq \min\{|S(r)| + 1, n\}$ which in turn implies $O(n)$ rounds until unique ids are assigned. Then another $O(n)$ rounds are required until nodes terminate. \square

Clearly, by executing a simple $O(n)$ -time process after *Dynamic_Naming* we can easily reassign minimal (consecutive) names to the nodes. The leader just floods a list of (old_id, new_id) pairs, one for each node in the network.

Note that the messages sent by *Dynamic_Naming* may be of size $\Omega(n^2)$. We now refine *Dynamic_Naming* to arrive at a more involved construction that re-

duces the message size to $\Theta(\log n)$ by paying a small increase in termination time. We call this 4th version of our naming protocols *Individual Conversations*. Due to space restrictions, we only give the main idea here.

Protocol *Individual Conversations* [Main Idea]. To reduce the size of the messages (i) the assigned names are now of the form $k \cdot d + id$, where id is the id of the node, d is the number of *unique consecutive* ids that the leader knows so far, and $k \geq 1$ is a name counter (ii) Any time that the leader wants to communicate to a remote node that has a unique id it sends a message with the id of that node and a timestamp equal to the current round. The timestamp allows all nodes to prefer this message from previous ones so that the gain is twofold: the message is delivered and no node ever issues a message containing more than one id. The remote node then can reply in the same way. For the assignment formula to work, nodes that obtain ids are not allowed to further assign ids until the leader freezes all named nodes and reassigns to them unique consecutive ids. During freezing, the leader is informed of any new assignments by the named nodes and terminates if all report that no further assignments were performed.

Theorem 7. *Individual Conversations solves the (minimal) naming problem in $O(n^3)$ rounds using messages of size $\Theta(\log n)$.*

References

- [AFR06] J. Aspnes, F. E. Fich, and E. Ruppert. Relationships between broadcast and shared memory in reliable anonymous distributed systems. *Distributed Computing*, 18[3]:209–219, February 2006.
- [AGVP90] B. Awerbuch, O. Goldreich, R. Vainish, and D. Peleg. A trade-off between information and communication in broadcast protocols. *Journal of the ACM (JACM)*, 37[2]:238–256, 1990.
- [AH00] S. Albers and M. Henzinger. Exploring unknown environments. *SIAM J. Comput.*, 29[4]:1164–1188, 2000.
- [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. Springer-Verlag, 2008.
- [Ang80] D. Angluin. Local and global properties in networks of processors (extended abstract). In *Proceedings of the twelfth annual ACM symposium on Theory of computing (STOC)*, pages 82–93. ACM, 1980.
- [ASW88] H. Attiya, M. Snir, and M. K. Warmuth. Computing on an anonymous ring. *J. ACM*, 35[4]:845–875, October 1988.
- [AW04] H. Attiya and J. Welch. *Distributed computing: fundamentals, simulations, and advanced topics*, volume 19. Wiley-Interscience, 2004.
- [BV99] P. Boldi and S. Vigna. Computing anonymously with arbitrary knowledge. In *Proceedings of the 18th annual ACM symposium on Principles of distributed computing (PODC)*, pages 181–188. ACM, 1999.
- [CFQS11] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *Ad-hoc, Mobile, and Wireless Networks*, pages 346–359, 2011.

- [CMM12] J. Chalopin, Y. Métivier, and T. Morsellino. On snapshots and stable properties detection in anonymous fully distributed systems (extended abstract). In *Structural Information and Communication Complexity*, volume 7355 of *LNCS*, pages 207–218. Springer, 2012.
- [Dol00] S. Dolev. *Self-stabilization*. MIT Press, Cambridge, MA, USA, 2000.
- [DP90] X. Deng and C. Papadimitriou. Exploring an unknown graph. In *31st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 355–361. IEEE, 1990.
- [FPPP00] P. Fraigniaud, A. Pelc, D. Peleg, and S. Pérennes. Assigning labels in unknown anonymous networks (extended abstract). In *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing (PODC)*, pages 101–111. ACM, 2000.
- [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.
- [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 '00*, pages 504–513. ACM, 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. ACM, 2010.
- [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.
- [MCS12a] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Brief announcement: naming and counting in anonymous unknown dynamic networks. In *26th International Symposium on Distributed Computing (DISC)*, pages 437–438. Springer, 2012.
- [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. Also in *Journal of Parallel and Distributed Computing* (To appear).
- [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 *LNCS*, 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 '05*, pages 104–110. ACM, 2005.
- [YK96] M. Yamashita and T. Kameda. Computing on anonymous networks. i. characterizing the solvable cases. *Parallel and Distributed Systems, IEEE Transactions on*, 7[1]:69–89, 1996.