

# Universal Finite-State and Self-Stabilizing Computation in Anonymous Dynamic Networks\*

Giuseppe A. Di Luna ✉

DIAG, Sapienza University of Rome, Italy

Giovanni Viglietta ✉

Department of Computer Science and Engineering, University of Aizu, Japan

---

## Abstract

A communication network is said to be *anonymous* if its agents are indistinguishable from each other; it is *dynamic* if its communication links may appear or disappear unpredictably over time. Assuming that an anonymous dynamic network is always connected and each of its  $n$  agents is initially given an input, it takes  $2n$  communication rounds for the agents to compute an arbitrary (frequency-based) function of such inputs (Di Luna–Viglietta, DISC 2023).

It is known that, without making additional assumptions on the network and without knowing the number of agents  $n$ , it is impossible to compute most functions and explicitly *terminate*. In fact, current state-of-the-art algorithms only achieve *stabilization*, i.e., allow each agent to return an output after every communication round; outputs can be changed, and are guaranteed to be all correct after  $2n$  rounds. Such algorithms rely on the incremental construction of a data structure called *history tree*, which is augmented at every round. Thus, they end up consuming an unlimited amount memory, and are also prone to errors in case of memory loss or corruption.

In this paper, we provide a general *self-stabilizing* algorithm for anonymous dynamic networks that stabilizes in  $\max\{4n - 2h, 2h\}$  rounds (where  $h$  measures the amount of corrupted data initially present in the memory of each agent), as well as a general *finite-state* algorithm that stabilizes in  $3n^2$  rounds. Our work improves upon previously known methods that only apply to static networks (Boldi–Vigna, Dist. Comp. 2002). In addition, we develop new fundamental techniques and operations involving history trees, which are of independent interest.

**2012 ACM Subject Classification** Theory of computation → Distributed algorithms; Computing methodologies → Distributed algorithms

**Keywords and phrases** anonymous dynamic network, history tree, self-stabilization, finite-state stabilization

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2024.10

**Related Version** *Full Version*: <https://arxiv.org/abs/2409.00688> [7]

**Funding** *Giuseppe A. Di Luna*: Project Ateneo SAPIENZA RM1221816C1760B and project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union.

*Giovanni Viglietta*: JSPS KAKENHI grant 23K10985 and the University of Aizu Competitive Research Funding for FY 2024.

## 1 Introduction

**Dynamic networks.** The rise of technologies such as wireless sensor networks, software-defined networks, and ad-hoc networks of smart devices has led to the prevalence of network topologies characterized by frequent and continuous dynamism. In the algorithmic study

---

\* Both authors contributed equally to this research.



© Giuseppe A. Di Luna and Giovanni Viglietta;  
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles of Distributed Systems (OPODIS 2024).

Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 10; pp. 10:1–10:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of these *dynamic networks*, the most common model of computation considers a system of agents that communicate synchronously through broadcast, where the network topology is always connected but the communication links can change unpredictably.

**Anonymous networks.** A classic assumption is that each agent has a unique ID, and efficient algorithms are known for solving many problems in this context [3, 11, 12, 13, 14, 16]. An alternative model is where agents are *anonymous*, and are therefore initially indistinguishable. The study of anonymous networks is both practically and theoretically significant, with prolific research spanning several decades [1, 4, 5, 6, 10, 17, 19].

**Computing in anonymous dynamic networks.** Recent work introduced the use of *history trees* for studying anonymous dynamic networks with a unique leader [8]. This structure can be constructed using the local views of agents; unlike similar structures, it has a temporal dimension that completely captures a network’s dynamism. The theory of history trees was later extended to networks with no leader and networks with multiple leaders [9].

The introduction of history trees enabled the design of optimal linear-time *universal algorithms* for all the aforementioned network scenarios (in a sense that will be elucidated in Section 2). However, an important difference between networks with leaders and without leaders is that in the latter it is impossible to perform *terminating* computations without additional knowledge about the network (such as the number of its agents). In leaderless networks, algorithms can only *stabilize* on the correct output without explicitly terminating.

**Self-stabilization.** In real systems, temporary faults such as memory corruption or message losses are unavoidable; algorithms that withstand such faults are called *self-stabilizing*. In *static* anonymous networks, universal self-stabilizing algorithms exist [2]. However, to the best of our knowledge, a similar result for dynamic networks has never appeared in the literature. It is therefore desirable to adapt the techniques in [8, 9] to achieve self-stabilization.<sup>1</sup>

**Finite-state computation.** We remark that the stabilizing algorithms from previous work on anonymous dynamic networks use unlimited local memory.<sup>2</sup> For instance, if history trees are used, all agents are required to continually expand this data structure in their internal memory. It would be beneficial to develop a version of the stabilizing algorithms in [8, 9] that uses finite memory, polynomially dependent on network size, which is crucial for implementation in real systems with limited memory.

## Contributions and technical challenges

In this paper, we propose two separate results that address the weaknesses we have discussed. Both results were also briefly announced in [18]. The full version of this paper, including pseudocode of all algorithms and missing proofs, is found in [7].

The first result, presented in Section 4, is a self-stabilizing version of the stabilizing algorithm proposed in [9]. While the original algorithm stabilizes in  $2n$  rounds in a network of  $n$  agents, our self-stabilizing version takes  $\max\{4n - 2h, 2h\}$  rounds (where  $h$  measures the corrupted data initially present in the memory of agents). This algorithm relies on a

<sup>1</sup> Note that terminating algorithms cannot be self-stabilizing, because an adversary could corrupt the initial memory states of all agents, causing them to erroneously terminate with an arbitrary output.

<sup>2</sup> Terminating algorithms use a bounded amount of memory, and thus are automatically finite-state.

novel operation on history trees called *chop*, which deletes old information and eventually removes corrupted data from memory. The chop operation is also used to merge history trees of different heights. We believe the dependency on  $h$  to be unavoidable in dynamic networks.

The second result, presented in Section 5, is a finite-state adaptation of the stabilizing algorithm from [9], resulting in an algorithm that uses memory polynomial in the number of agents  $n$  and stabilizes in  $3n^2$  rounds. The idea of this algorithm is to avoid updating an agent's state for a round if it receives no new relevant information. This seemingly simple strategy has to be carefully implemented to guarantee the algorithm's correctness. In fact, not updating an agent's state causes a "desynchronization" of the system that has to be dealt with properly. This leads to the formulation of a generalized theory of history trees, which is of independent interest and may have important applications to asynchronous networks.

As a bonus, Section 3 offers an algorithm that is both self-stabilizing and finite-state, and stabilizes in  $2n$  rounds, which is optimal. The downside is that it assumes  $n$  to be known.

## Related Work

A thorough discussion of related work is found in [7]. Essentially, the most relevant work on self-stabilizing algorithms for anonymous network is [2], which critically assumes the network to be static. As for finite-state computation in dynamic networks, the closest work is [15], which proposes a quantized (hence inexact) algorithm in  $O(n^3 \log n)$  rounds.

## 2 Model Definition and Basic Structures

**Model of computation.** A *dynamic network* is defined as an infinite sequence  $\mathcal{G} = (G_t)_{t \geq 1}$ , where  $G_t = (V, E_t)$  is an undirected multigraph,  $V = \{p_1, p_2, \dots, p_n\}$  is a system of  $n$  *anonymous agents*, and  $E_t$  is a multiset of edges representing *links* between agents.<sup>3</sup>

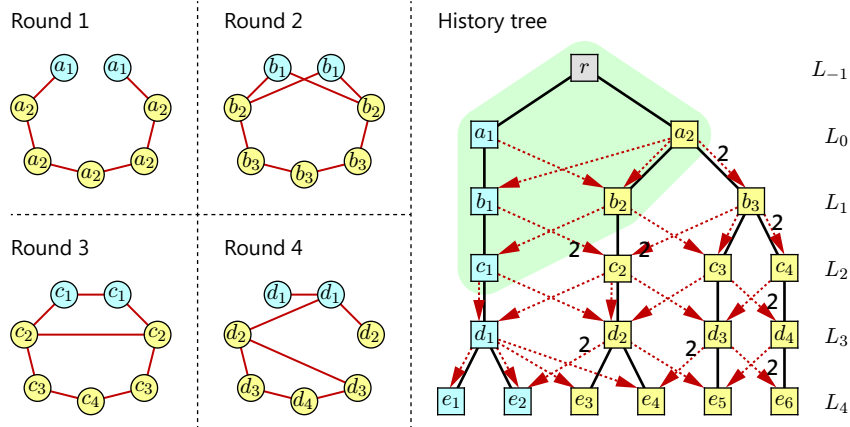
Each agent is assigned an *input* at *round* 0. Agents also have internal states, which are initially determined by their inputs (thus, agents with the same input start in the same state). At every *round*  $t \geq 1$ , each agent sends a message to its neighbors in  $G_t$  through all its incident links. Every agent creates its message as a function of its current state and sends the same message to all its current neighbors. During round  $t$ , each agent reads all messages coming from its neighbors (in no particular order). Based on the information therein, the agent updates its internal state according to a local algorithm, the same for all agents. An agent may also return an *output* at the end of a round. After that, the next round starts.

Note that our network model is *synchronous*, i.e., all agents are assumed to send and receive their messages simultaneously, and a new round starts at the same time for all agents.

**Universal computation.** If the outputs of all agents remain unchanged starting at round  $t$ , the system is said to *stabilize* in  $t$  rounds. If executing a certain algorithm causes a system to stabilize regardless of the inputs assigned to its  $n$  agents, we view such an algorithm as computing a function from an  $n$ -tuple of inputs to an  $n$ -tuple of outputs.

Not all functions can be computed by a system of anonymous agents. It is known that, without making extra assumptions on the network, the agents can essentially only compute the *Input Frequency function*, which is the function that returns the percentage of agents

<sup>3</sup> The parallel links in multigraphs can serve as a model for the multi-path propagation of radio waves. Of course, all the results in this paper also apply to networks modeled as simple graphs.



■ **Figure 1** The first communication rounds of an anonymous dynamic network of  $n = 7$  agents and the corresponding levels of its history tree. Initial inputs are represented by agents' colors (cyan or yellow). Labels on agents and nodes have been added for the reader's convenience and indicate classes of indistinguishable agents. The portion of the history tree with a green background is the view of the two agents labeled  $c_1$  after two communication rounds. Level  $L_2$  is a counting level, which yields the equations  $|c_1| = |c_2| = |c_3| = 2|c_4|$ . Knowing that  $|c_1| = 2$ , one can infer that  $n = 7$ .


that have each input value [9]. Once this fundamental function has been computed, the system can then immediately compute any *frequency-based* function, i.e., any function that depends only on input percentages, such as the weighted average of the inputs, etc.<sup>4</sup>

For this reason, any algorithm that allows a system to compute the Input Frequency function for any multiset of inputs assigned to the agents is said to be *universal*.

**History trees.** *History trees* were introduced in [8] as a tool of investigation for anonymous dynamic networks; an example is found in Figure 1. A history tree is a representation of a dynamic network given some inputs to its agents. It is an infinite graph whose nodes are partitioned into *levels*  $L_t$ , with  $t \geq -1$ ; each node in  $L_t$  represents a class of agents that are *indistinguishable* at the end of round  $t$  (with the exception of  $L_{-1}$ , which contains a single root node  $r$  representing all agents). The definition of distinguishability is inductive: At the end of round 0, two agents are distinguishable if and only if they have different inputs. At the end of round  $t \geq 1$ , two agents are distinguishable if and only if they were already distinguishable at round  $t - 1$  or if they have received different multisets of messages during round  $t$ . (We refer to “multisets” of messages, as opposed to sets, because multiple copies of identical messages may be received; thus, each message has a *multiplicity*.) The *anonymity* of a node  $v$  in a history tree is the number of agents represented by  $v$ , denoted as  $\mathbf{a}(v)$ .

Every node (other than the root  $r$ ) has a label indicating the input of the agents it represents. There are also two types of edges connecting nodes in adjacent levels. The *black edges* induce an infinite tree spanning all nodes, rooted at node  $r \in L_{-1}$ . The presence of a black edge  $(v, v')$ , with  $v \in L_t$  and  $v' \in L_{t+1}$ , indicates that the *child node*  $v'$  represents a subset of the agents represented by the *parent node*  $v$ . The *red multi-edges* represent communications between agents. The presence of a red edge  $(v, v')$  with multiplicity  $m$ , with  $v \in L_t$  and  $v' \in L_{t+1}$ , indicates that, at round  $t + 1$ , each agent represented by  $v'$  receives  $m$  (identical) messages from agents represented by  $v$ .

<sup>4</sup> If the network has a unique leader or the number of agents is known, computing the Input Frequency function also allows the system to determine the exact number of agents that were assigned each input.



■ **Figure 2** Updating the view of an agent represented by node  $v$  after it receives the view of an agent represented by  $u$  as a message. The two views are match-and-merged starting from the roots;  $v$  also gets a child  $v'$ , and a new red edge from  $u$  to  $v'$  is added, representing this interaction.

The *view*  $\mathcal{V}$  of an agent  $p$  at round  $t \geq 0$  is the subgraph of the history tree that is spanned by all the shortest paths (using black and red edges indifferently) from the root  $r$  to the node in  $L_t$  representing  $p$  (Figure 1 shows an example of a view). The node in  $\mathcal{V}$  representing  $p$  at round  $t$  is the *bottom node* of the view; this is the node of  $\mathcal{V}$  farthest from  $r$ .


**Constructing views.** It was shown in [8] that each agent can construct its view of the history tree in real time. To achieve this, each agent is required to send its current view to all its neighbors at every round, and then match-and-merge it with all the views it receives from them. An example of this operation is shown in Figure 2. At every round, the bottom node  $v$  of the view also receives a new child  $v'$ , which becomes the new bottom node. After merging an incoming view, a red edge is added from the bottom node of this view to  $v'$ . The multiplicities of such red edges match the number of isomorphic views that are being merged.

**Computing with history trees.** As proved in [9], if the network is connected at every round, then the view of any agent at round  $2n - 2$  contains enough information to compute the Input Frequency function. This is done by finding a *counting level*, which is a level in the history tree where every node has exactly one child; level  $L_2$  in Figure 1 is an example.

For any red edge with multiplicity  $m_1 > 0$  directed from a node  $v$  in a counting level  $L_t$  to the child of another node  $u \in L_t$ , there must be a red edge with multiplicity  $m_2 > 0$  from  $u$  to the child of  $v$ . Then, we have  $m_1 \mathbf{a}(v) = m_2 \mathbf{a}(u)$ . By writing the system of such linear equations for the entire level  $L_t$  and solving it in terms of a single free variable, we can easily compute the Input Frequency function. Note that the first counting level  $L_t$  occurs when  $t \leq n - 2$ , and it takes at most  $t + n$  rounds for all the nodes in  $L_t$  and their children to appear in the views of all agents. Thus, this algorithm stabilizes in at most  $2n - 2$  rounds.

**Self-stabilization.** An algorithm computes a function in a *self-stabilizing* manner if it does so regardless of the initial states of the agents. That is, even if the states of the agents consist of incorrect history trees, a self-stabilizing algorithm is able to eventually erase all errors and stabilize on the correct output. Clearly, the above algorithm is not self-stabilizing.

**Finite-state computation.** An algorithm is *finite-state* if the states of the  $n$  agents in a system executing such an algorithm can be encoded by at most  $f(n)$  bits throughout the execution. Since the above algorithm constructs ever-growing views, it is not finite-state.



■ **Figure 3** The chop operation on a history tree view: eliminate level  $L_0$ , and then repeatedly merge nodes whose corresponding sub-views are isomorphic, also combining their outgoing red edges.

### 3 Universal Computation in Networks of Known Size

We first assume that the number of agents  $n$  in an anonymous dynamic network is known. Using this knowledge, we develop a relatively straightforward algorithm that is both self-stabilizing and finite-state and has an optimal stabilization time of  $2n - 2$  rounds.

**Chop operation.** The algorithm is based on an operation on history tree views called *chop*, illustrated in Figure 3. Given a view, eliminate the nodes in level  $L_0$  (as well as their incident black and red edges) and connect the root to all nodes in level  $L_1$  via black edges. As a result, all non-root nodes are shifted by one level: those that were in  $L_1$  are now in  $L_0$ , etc.

After that, we have to ensure that the resulting view is well-formed. Note that every node  $v$  in a view  $\mathcal{V}$  is the bottom node of a maximal view  $\mathcal{V}_v \subseteq \mathcal{V}$ , called the *sub-view* of  $\mathcal{V}$  at  $v$ . To complete the chop operation, we repeatedly merge the nodes whose sub-views are isomorphic. This is done level by level, starting from the nodes in  $L_0$ , then the nodes in  $L_1$ , etc. As a consequence, when two nodes  $u, v \in L_t$  are found whose sub-views  $\mathcal{V}_u$  and  $\mathcal{V}_v$  are isomorphic, they must be siblings. This is because the parents of  $u$  and  $v$  have isomorphic sub-views as well, and therefore they have already been merged as nodes of  $L_{t-1}$ .

Merging  $u$  and  $v$  is done as follows. First we delete one of them, say  $u$ . The children of  $u$  now become children of  $v$ . The red edges inbound to  $u$  are simply eliminated, while the red edges outbound from  $u$  are now redirected as outbound from  $v$ . Specifically, for every  $w \in L_{t+1}$ , if the red edge  $(v, w)$  originally has multiplicity  $m \geq 0$  and the red edge  $(u, w)$  has multiplicity  $m' \geq 0$ , after merging  $u$  and  $v$  the red edge  $(v, w)$  will have multiplicity  $m + m'$ .

This is justified by the fact that the agents represented by  $u$  and  $v$  have received identical multisets of messages and are therefore indistinguishable; hence the need to merge the two nodes. Thus, all messages received from agents represented by  $u$  and  $v$  are now considered identical, and the multiplicities of the corresponding red edges must be added together.

The significance of the chop operation is given by the following result, whose proof is found in [7]. Essentially, chopping a view is equivalent to “forgetting” the first communication round of the network that originated that view.

► **Lemma 1.** *Let  $\mathcal{G} = (G_1, G_2, G_3, \dots)$  be a network, and let  $\mathcal{V}$  be the view of an agent  $p$  at round  $t \geq 1$ . Then, the view  $\mathcal{V}'$  obtained by chopping  $\mathcal{V}$  is isomorphic to the view of  $p$  at round  $t - 1$  in the network  $\mathcal{G}' = (G_2, G_3, G_4, \dots)$ , where agents’ inputs are as in  $\mathcal{G}$ . ◀*

**Algorithm overview.** The pseudocode of our algorithm is found in [7]. The goal is for each agent to construct a coherent view of the history tree of the  $2n - 2$  most recent rounds, which enables the stabilizing algorithm from [9] to correctly compute the Input Frequency function based on the first visible counting level (cf. Section 2).

If the state of an agent  $p$  does not represent a well-formed view of a history tree or represents a view with more than  $2n - 2$  levels, it is simply re-initialized to a view containing a root and a single child labeled as the input of  $p$ . When  $p$  receives the views of its neighbors, it determines the view with smallest height (including its own) and trims all other views by repeatedly performing the chop operation on them, until all views have the same height. Then, the resulting views are match-and-merged into  $p$ 's view as usual. After that, if the current view of  $p$  has  $2n - 1$  levels, the oldest level is removed via a chop operation. Finally, if the resulting view contains a counting level, this is used to output the frequencies of all inputs; otherwise,  $p$  returns a default output consisting of its own input with 100% frequency.

► **Theorem 2.** *For every  $n \geq 1$ , there is a self-stabilizing, finite-state universal algorithm that operates in any connected anonymous dynamic network of known size  $n$  and stabilizes in at most  $2n - 2$  rounds.*

**Proof.** We will prove that the above algorithm satisfies the required conditions. Observe that, since the network is connected at every round, it takes at most  $n - 1$  rounds for all agents to have states representing views of the same height. In fact, trimming one's view to match the shortest incoming view is akin to broadcasting the smallest height of a view across the whole network.<sup>5</sup>

Due to Lemma 1, if the algorithm alternately performs update and chop operations on non-trivial views for  $k$  times, the result is the same as chopping all views  $k$  times first, and then doing all the update operations. Since all views are kept within  $2n - 2$  levels, it takes as many initial chops to eliminate all incorrect levels. At the same time,  $2n - 2$  update operations result in all agents constructing views correctly representing the previous  $2n - 2$  communication rounds. It follows that all agents return the correct output at any round  $t \geq 2n - 2$ .

The algorithm is finite-state because every (non-initial) state always consists of  $O(n)$  levels of a view. If this view is correct, its size is  $O(n^3 \log n)$  bits. If the view is not correct and exceeds  $O(n^3 \log n)$  bits in size, the algorithm can chop it until it is sufficiently small. ◀

The stabilization time of this algorithm is asymptotically optimal, since a lower bound of  $2n - O(1)$  rounds was proved in [9], even for algorithms that are not self-stabilizing.

## 4 Universal Self-Stabilizing Computation

We will now propose a self-stabilizing universal algorithm for anonymous dynamic networks that operates without any knowledge of the number of agents  $n$ .

**Garbage coefficient.** Recall that a self-stabilizing algorithm should tolerate any amount of incorrect data initially present in the agents' local memory. If such initial data does not encode the view of a history tree, it is definitely incorrect and can be discarded immediately. However, even if an initial state does encode a well-formed view, the information therein may be incorrect and possibly deceiving. We define the *garbage coefficient* of an agent as the height of the view encoded by its initial state. If its initial state does not encode a well-formed view, its garbage coefficient is defined to be 0.

<sup>5</sup> It is well known that the dynamic diameter of a connected network is  $d \leq n - 1$ ; see [8, 12].

**Algorithm overview.** We will now describe our self-stabilizing algorithm. Its pseudocode is found in [7]. The structure of this algorithm is similar to the one in Section 3; in particular, the same chop operation is used to gradually eliminate all incorrect levels from every view of the history tree. The main difficulty is that, without knowledge of  $n$ , it is impossible to determine whether a view has the desired number of  $2n - 2$  levels.

To overcome this difficulty, our strategy is to chop each agent's view once every two rounds. As a result, the views grow by one level every two rounds, and eventually acquire the desired amount of correct levels.

**Algorithm details.** To control the chopping of views, the state of each agent is augmented with a binary flag that is toggled between 0 and 1 at every round. Every time an agent's flag becomes 0, its view is chopped. Accordingly, if the current state of an agent does not encode a well-formed view augmented with a binary flag, the state is discarded and reset to a view of only two nodes (a root with a single child containing the agent's input) and a flag set to 0.

An agent's binary flag is also attached to all messages it sends, alongside its view of the history tree. Upon receiving messages from its neighbors, the agent reduces all views (including its own) to the same height via chop operations, in the same way as the algorithm in Section 3 does. However, this is now extended to all flags, as well. Specifically, we define a function  $\mathbf{eval}(\mathcal{V}, b)$ , which takes a view  $\mathcal{V}$  and a flag  $b \in \{0, 1\}$ , and returns  $2h_{\mathcal{V}} + b$ , where  $h_{\mathcal{V}}$  is the height of  $\mathcal{V}$ . Let the *minimum pair*  $(\mathcal{V}^*, b^*)$  be the pair that minimizes  $\mathbf{eval}$  among the ones received from its neighbors, as well as the agent's own pair. Then, all views received by the agent, as well as the agent's own view, are chopped until they have the same height as  $\mathcal{V}^*$ . After that, the incoming views are match-and-merged into the agent's view, as usual. In addition, the agent's own flag is set to  $b^*$ . The flag is then toggled, and if it becomes 0 the agent's view is chopped once. This protocol ensures proper synchronization among agents.

Finally, if the agent's resulting view contains a counting level, this level is used to compute and output the frequencies of all inputs (as explained in Sections 2 and 3); otherwise, the agent returns a default output consisting of its own input with 100% frequency.

► **Theorem 3.** *There is a self-stabilizing universal algorithm that operates in any connected anonymous dynamic network of unknown size  $n$  and stabilizes in at most  $\max\{4n - 2h, 2h\}$  rounds, where  $h$  is the minimum garbage coefficient across all agents.*

**Proof.** The correctness of this algorithm is straightforward, given the property of the chop operation established in Lemma 1 and the observations in the proof of Theorem 2. It remains to show that the stabilization time is  $\max\{4n - 2h, 2h\}$ .

Recall that the garbage coefficient of an agent  $p$  is defined as the number of levels of the view initially encoded by the state of  $p$ . Every time a chop operation is performed on the view of  $p$ , the number of these levels is decreased by one unit. We define the *leftover garbage* of  $p$  at round  $t \geq 0$  as the garbage coefficient of  $p$  minus the number of times the view of  $p$  has been chopped up to round  $t$  (or 0 if this number is negative). Thus, the leftover garbage measures how many initial levels are still present in the view of an agent. We define  $h_t$  as the minimum leftover garbage at round  $t \geq 0$  across all agents. Clearly,  $h_0 = h$ .

As already remarked in the proof of Theorem 2, it takes at most  $n - 1$  rounds for all agents' views to have the same height and all their flags to be the same (essentially, the algorithm performs a broadcast of the smallest value returned by  $\mathbf{eval}$ ). In fact, every time two agents communicate, their leftover garbage is equalized. Thus, within  $n - 1$  rounds, all agents have the same leftover garbage. Also, every two rounds the leftover garbage of any agent with smallest  $\mathbf{eval}$  is reduced by one unit (unless it is already 0). In other words,  $h_t$  is decremented every two rounds until it reaches 0. In particular,  $h_{2h} = 0$ .



Let us assume that  $h \leq n$ ; we will prove that the stabilization time in this case is  $\max\{4n - 2h, 2h\} = 4n - 2h$ . Since  $h_{2h} = 0$ , after  $\max\{2h, n\} \leq 4n - 2h$  rounds all agents have no leftover garbage. We conclude that, after  $2(2n - h) = 4n - 2h$  rounds, every view has  $h + (2n - h) = 2n$  levels, which are all correct. From this point onward, each agent's view will consistently represent at least  $2n$  previous communication rounds of the network, and hence every agent will return the correct output.

Let us now assume that  $h > n$ ; we will prove that the stabilization time in this case is  $\max\{4n - 2h, 2h\} = 2h$ . By the previous argument, after  $\max\{2h, n\} = 2h$  rounds, all agents have no leftover garbage. At this time, there are  $2h > 2n$  correct levels in every agent's view, and so their outputs are all correct, and will be correct in all subsequent rounds. ◀

Observe that our algorithm's stabilization time exhibits a linear dependence on the minimum garbage coefficient  $h$ . We conjecture that this cannot be avoided in dynamic networks, and so this stabilization time is asymptotically optimal. Nonetheless, it is remarkable that the best stabilization time of our algorithm is achieved when  $h = n$  (as opposed to  $h = 0$ ).

## 5 Universal Finite-State Computation

Our final contribution is a finite-state universal algorithm for anonymous dynamic networks that operates without any knowledge of the number of agents  $n$ . This is achieved by preventing agents from updating their states under certain conditions, which leads to different agents potentially having views of different heights. Several new definitions are required.

**Generalized views.** Before we state our algorithm, we first have to define a generalized notion of view. This is necessary because we will have to deal with situations where views of different heights must be match-and-merged into each other. The result of such an operation is illustrated in Figure 4. As usual, when an agent receives some views from its neighbors, it attaches a child  $v'$  to the bottom node of its own view, and then match-and-merges the incoming views, connecting their bottom nodes to  $v'$  via red edges. If the incoming views have arbitrary heights, the resulting view may not have only red edges connecting a level to the next, but also red edges spanning any number of levels, going downward or upward.


Formally, a generalized view can be defined recursively as the result of the above update operation performed on arbitrary generalized views, where the base case is the trivial view consisting of the single root node. Accordingly, we define the bottom node of a generalized view as its unique sink, i.e., the unique node with no outgoing red or black edges (where black edges are understood as being directed away from the root). By definition, the bottom node of a view represents the agents that currently possess this view.

The notion of *sub-view* of a generalized view  $\mathcal{V}$  at a node  $v$  is as in Section 3: it is the maximal view included in  $\mathcal{V}$  whose bottom node is  $v$ . As usual, a node  $v$  in  $\mathcal{V}$  represents all agents that had a view isomorphic to the sub-view of  $\mathcal{V}$  at  $v$  at the end of some round.

From now on, when no confusion may arise, the term “generalized” will often be omitted.

**Collective trees.** Since our algorithm allows agents to skip updating their states, it follows that they are no longer constructing views of the history tree of the network as defined in Section 2. Nonetheless, it is useful to define a structure that may act as a global reference incorporating all the (generalized) views that the agents are constructing.

The *collective tree* at round  $t$ , denoted as  $\mathcal{H}_t$ , is the structure obtained by match-and-merging together all the views constructed by the agents in a network up to round  $t$ . Of course,  $\mathcal{H}_t$  depends on the local algorithm being executed by the agents, which affects the way



■ **Figure 4** The views  $\mathcal{A}$  and  $\mathcal{B}$  of two communicating agents. Since their heights are different, their updated versions  $\mathcal{A}'$  and  $\mathcal{B}'$  are generalized views, where red edges may go upward or skip multiple levels. The highlighted nodes are the bottom nodes.

they construct their views. For example, under the standard stabilizing algorithm outlined in Section 2, the collective tree  $\mathcal{H}_t$  is simply the history tree of the network truncated at level  $L_t$ . Under the finite-state algorithm that we are going to describe in this section,  $\mathcal{H}_t$  is constructed from generalized views, and may have little in common with the history tree.

**Exposed pairs.** Another feature of our finite-state algorithm is the following: Even if an agent updates its state in a given round, it may still choose to selectively discard some incoming messages from certain neighbors. However, there is a caveat: The algorithm guarantees that, if two agents send each other messages for a round, either both agents discard each other's messages or neither does.

Therefore, in any collective tree, whenever a node  $v$  has an outgoing red edge of multiplicity  $m_1 > 0$  to a child of a node  $u$ , then symmetrically  $u$  has an outgoing red edge of multiplicity  $m_2 > 0$  to a child of  $v$ . Moreover, there are a child of  $v$  and a child of  $u$  that appear in the collective tree at the same round, i.e., the round in which an agent represented by  $v$  first exchanges messages with one represented by  $u$ .

With the above notation,  $v$  and  $u$  are said to be an *exposed pair* if they both have a unique child. In this case, we have the following equation, as already noted in Section 2:

$$m_1 \mathbf{a}(v) = m_2 \mathbf{a}(u) \quad (1)$$


The *anonymity*  $\mathbf{a}(x)$  is defined, as usual, as the number of agents represented by the node  $x$ .

► **Observation 4.** *If  $v$  and  $u$  are an exposed pair in  $\mathcal{H}_t$ , then there is a round  $t' \leq t$  such that the unique children of  $v$  and  $u$  are created simultaneously in  $\mathcal{H}_{t'}$ .*

**Counting cuts.** We will now define a generalized notion of counting level. Recall from Section 2 that a counting level is a level of the history tree where all nodes have a unique child. A counting level's outgoing red edges describe all interactions that occur in the corresponding round, and this information can be used to compute the Input Frequency function.

We define a *cut* in a collective tree  $\mathcal{H}_t$  as a set of nodes  $C$  such that, for every leaf  $f$  in  $\mathcal{H}_t$ , the unique black path from the root of  $\mathcal{H}_t$  to  $f$  contains exactly one node of  $C$ .

Let  $G$  be the undirected graph on the nodes of  $\mathcal{H}_t$  whose edges are the exposed pairs. A cut  $C$  in  $\mathcal{H}_t$  is said to be a *counting cut* if  $C$  induces a connected subgraph of  $G$ . Figure 5 (left) shows two counting cuts consisting of the nodes labeled  $A$  and  $B$ , respectively. Clearly, in the history tree of a connected network, the notions of counting level and counting cut coincide.



■ **Figure 5** Left: Two counting cuts, where  $A$  dominates  $B$  (irrelevant red edges have been omitted). Right: Impossible configurations in a collective tree, implying that counting cuts are totally ordered.

In a collective tree, a cut  $A$  *dominates* a cut  $B$  if every node of  $B$  has an ancestor in  $A$ . This is the case in Figure 5 (left), where the counting cut  $A$  dominates the counting cut  $B$ .

► **Lemma 5.** *In any collective tree, dominance is a strict total order on counting cuts.*

**Proof.** The asymmetric and transitive properties of the dominance relation are obvious from the definition. We only have to prove that any two distinct counting cuts  $A$  and  $B$  in the same collective tree are comparable, i.e., one dominates the other.

Assume the opposite for a contradiction. Then, without loss of generality, there is a node  $v \in A$  with a descendant  $w \in B$ . Moreover, there are two nodes  $u \in B$  and  $z \in A$  such that either  $u$  is an ancestor of  $z$  or  $u = z$ ; both cases are sketched in Figure 5 (right).

As pointed out in Observation 4, the two children of an exposed pair make their appearance in the collective tree at the same round. Since a counting cut consists of a connected set of exposed pairs, all the children of its nodes appear in the collective tree at the same round.

Thus, the children of  $v$  and  $z$  are both created at round  $t$ , and the children of  $u$  and  $w$  are both created at round  $t'$ . Since the child of  $v$  is unique, it is an ancestor of the child of  $w$ , and therefore  $t < t'$ . Similarly, the child of  $u$  is either an ancestor of the child of  $z$  or they coincide (if  $u = z$ ). As a consequence,  $t' \leq t$ , which yields  $t < t$ , a contradiction. ◀

Due to Lemma 5, if a collective tree has at least one counting cut, then there is a well-defined *dominant* one, whose nodes are closest to the root.

The definitions of counting cut and dominance extend verbatim to generalized views. However, if an agent has a counting cut  $C$  in its view  $\mathcal{V}$ , the nodes of  $C$  as embedded in the collective tree  $\mathcal{H}_t$  are not necessarily a counting cut. For instance, although every node of  $C$  has a unique child in  $\mathcal{V}$ , it may have multiple children in  $\mathcal{H}_t$ , as some of them may be missing from  $\mathcal{V}$ . Thus, Equation (1), which we established for exposed pairs in a collective tree, may not hold in a view. Similarly, Observation 4 and Lemma 5 may not hold in a view.

**Algorithm overview.** We are finally ready to provide our finite-state algorithm. Its pseudocode is in [7]. Each agent sends its view to all its neighbors and updates its own view by match-and-merging incoming views as usual, with one exception: If the view of an agent contains a dominant counting cut (i.e., a counting cut that dominates all others in the view), then all incoming views that contain the same dominant counting cut are discarded. Moreover, if all incoming views are discarded, the agent skips updating its own view altogether for that round (hence it does not even add a child to the bottom node).

The rationale is that an agent with a dominant counting cut “believes” that this cut is dominant for all agents and is sufficient to compute the Input Frequency function. Therefore, this agent deems unnecessary to further update its view, unless its belief is proven incorrect.

**Algorithm details.** To make our algorithm more precise, we have to define the concept of *isomorphism* between counting cuts in different views. A counting cut  $C$  in a view  $\mathcal{V}$  is *isomorphic* to a counting cut  $C'$  in a view  $\mathcal{V}'$  if there is a bijection  $f: C \rightarrow C'$  such that, for every  $v \in C$ , the sub-view of the unique child of  $v$  in  $\mathcal{V}$  is isomorphic to the sub-view of the unique child of  $f(v)$  in  $\mathcal{V}'$ . This definition extends verbatim to collective trees.

According to the algorithm, if two agents have dominant counting cuts in their respective views, and these cuts are isomorphic, both agents discard each other’s messages. In all other cases (i.e., if either of them does not have a counting cut, or does not have a dominant counting cut, or both have dominant counting cuts which are not isomorphic), the agents use each other’s messages to update their respective views. Hence Observation 4, which relies on the fact that interactions between pairs of agents occur symmetrically and simultaneously, is indeed correct, and so are its consequences, such as Lemma 5.

If, due to this rule, an agent ends up discarding all messages incoming from its neighbors, it does not update its view at all, i.e., it does not even add a child to the bottom node of its view (cf. Figure 2). There is only one trivial exception: If an agent’s view consists of a single node  $r$ , then this is necessarily round 0; the agent does not receive any messages in this round, but it still updates its view by adding a child to  $r$  and labeling it as its input.

At the end of every round, if an agent’s view has a dominant counting cut, this cut is used to compute the Input Frequency function by repeated application of Equation (1); the result is then returned as output by the agent. Otherwise, the agent simply returns the default output: its own input with a frequency of 100%.

**Total agreement.** There are several challenges to proving the correctness of this relatively simple algorithm. Even if an agent has a dominant counting cut in its view, using it to compute the Input Frequency function may lead to incorrect results, because Equation (1) might not hold in that view. Additionally, the rule that permits agents to discard messages could potentially result in a situation where all agents have an incorrect dominant counting cut, yet none of them updates its view, preventing any progress from being made.

The next lemma addresses these difficulties (its proof is found in [7]). We say that two agents *agree* on a counting cut  $C$  at round  $t$  if both of their views at round  $t$  contain isomorphic copies of  $C$ . There is *total agreement* on  $C$  if all agents in the system agree on  $C$ .

► **Lemma 6.** *If there is total agreement on a counting cut  $C$  at round  $t$ , then  $C$  is also a counting cut in the collective tree  $\mathcal{H}_t$ . Moreover, if  $C$  is the dominant counting cut in the views of all agents at round  $t$ , then all agents return the correct output at round  $t$ .*

Lemma 6 provides a sufficient condition for the algorithm to be correct. Thus, we only have to prove that eventually total agreement on a dominant counting cut is achieved.

**Collective tree dynamics.** To begin with, the collective tree does not have any counting cuts, except in the trivial case where all agents have the same input (then the root of the collective tree by itself constitutes a counting cut). We also remark that the existence of a counting cut in the collective tree does not necessarily imply the presence of a counting cut in any individual agent’s view. Next, we will study the way counting cuts are formed in the collective tree and how they eventually become counting cuts in the agents’ views.

A *branch* in a tree is any path from the root to a leaf. Thus, the number of branches in a tree is precisely the number of its leaves. We say that a *branching* occurs in the collective tree during round  $t$  if  $\mathcal{H}_t$  has more branches than  $\mathcal{H}_{t-1}$ . The collective tree *loses* a counting cut  $C$  at round  $t$  if  $C$  is a counting cut in  $\mathcal{H}_{t-1}$  but not in  $\mathcal{H}_t$ . Similarly, the collective tree *acquires* a counting cut  $C$  at round  $t$  if  $C$  is a counting cut in  $\mathcal{H}_t$  but not in  $\mathcal{H}_{t-1}$ .

The above definitions straightforwardly extend to views.

► **Lemma 7.** *If the collective tree (respectively, an agent's view) loses a counting cut at round  $t$ , then a branching occurs in the collective tree (respectively, in the same agent's view) at round  $t$ .*

**Proof.** We will only prove our claim for collective trees, as the proof for views is identical. Observe that all the red edges in  $\mathcal{H}_{t-1}$  also appear in  $\mathcal{H}_t$ . Thus, the only way a counting cut  $C$  may be lost at round  $t$  is if a node of  $C$  acquires an additional child in  $\mathcal{H}_t$  or if  $C$  is no longer a cut in  $\mathcal{H}_t$ , which therefore has a branch that does not intersect  $C$ . Clearly, both scenarios imply that a branching occurs at round  $t$  in the collective tree. ◀

The proof of the next lemma is found in [7].

► **Lemma 8.** *If at the beginning of round  $t$  no agent's view has a counting cut and no branching occurs in the collective tree during round  $t$ , then the collective tree at round  $t$  acquires exactly one counting cut.*

**Agent interactions.** Suppose that two agents  $p$  and  $p'$  send each other messages at round  $t$ , which are not discarded. Let  $C$  be a counting cut in the view  $\mathcal{V}$  of  $p$  at the end of round  $t-1$ , and assume that the view  $\mathcal{V}'$  of  $p'$  at the end of round  $t-1$  does not contain  $C$ . At the end of round  $t$ , there are four possibilities:

- The view of  $p$  loses  $C$  and the view of  $p'$  acquires  $C$ . In this case, a branching occurs in the view of  $p$ , due to Lemma 7.
- The view of  $p$  loses  $C$  and the view of  $p'$  does not acquire  $C$ . Also in this case, a branching occurs in the view of  $p$ , due to Lemma 7.
- The view of  $p$  does not lose  $C$  and the view of  $p'$  acquires  $C$ .
- The view of  $p$  does not lose  $C$  and the view of  $p'$  does not acquire  $C$ .

We will prove that, in the latter case, a branching occurs in the view of  $p'$ .

► **Lemma 9.** *With the above notation, if the view of  $p$  does not lose  $C$  and the view of  $p'$  does not acquire  $C$ , then a branching occurs in the view of  $p'$  at round  $t$ .*

**Proof.** Let  $\mathcal{T}$  be the tree obtained by match-and-merging  $\mathcal{V}$  and  $\mathcal{V}'$ . Since the view of  $p$  does not lose  $C$ , it follows that  $\mathcal{T}$  contains  $C$  as a counting cut. However, the view of  $p'$  fails to acquire  $C$  as a counting cut, which implies that other agents send their views to  $p'$  at round  $t$ , causing  $\mathcal{T}$  to lose  $C$ . Hence, a branching must occur in the view of  $p'$ , as argued in the proof of Lemma 7. ◀

**Undominated counting cuts.** We say that a counting cut  $C$  in an agent's view  $\mathcal{V}$  at round  $t$  is *undominated* if no counting cut in  $\mathcal{V}$  dominates  $C$ . In other words,  $C$  is a maximal element in the dominance partial order within  $\mathcal{V}$ .<sup>6</sup>

<sup>6</sup> As previously pointed out, Lemma 5 does not necessarily hold in a view, hence the term “partial order”.

► **Lemma 10.** *If an agent's view at round  $t$  has an undominated counting cut  $C$ , then at every round  $t' \geq t$  the counting cut  $C$  is undominated in every view that contains it.*

**Proof.** Let the view  $\mathcal{V}$  at round  $t$  contain the undominated counting cut  $C$ , and let a view  $\mathcal{V}'$  at round  $t'$  contain both  $C$  and a counting cut  $C'$  that dominates  $C$ . Then, the nodes of  $C'$  are ancestors of the nodes of  $C$ , and hence appear in every view that contains  $C$ . We infer that  $\mathcal{V}$  also contains  $C'$ , contradicting the assumption that  $C$  is undominated in  $\mathcal{V}$ . ◀

► **Lemma 11.** *Let  $p$  be an agent whose view at the beginning of round  $t$  contains an undominated counting cut  $C$ , and let  $p'$  be an agent whose view at the beginning of round  $t$  does not contain  $C$ . Then, if  $p$  and  $p'$  are neighbors at round  $t$ , they do not discard each other's messages.*

**Proof.** The algorithm makes  $p$  and  $p'$  discard each other's messages only if they have the same dominant counting cut at the beginning of round  $t$ . In this case, since  $C$  is undominated in the view of  $p$ , then  $C$  is dominant. However, by assumption the view of  $p'$  does not contain  $C$ , and therefore it does not have the same dominant counting cut as the view of  $p$ . ◀

**Algorithm correctness.** We conclude this section with a proof of correctness of our finite-state algorithm.

► **Theorem 12.** *There is a finite-state universal algorithm that operates in any connected anonymous dynamic network of unknown size  $n$  and stabilizes in at most  $3n^2$  rounds.*

**Proof.** Since each agent starts with a view containing a single branch, a branching in its view may occur at most  $n - 1$  times (note that the leaves represent disjoint and non-empty sets of agents). Thus, in total, the agents' views may branch at most  $n(n - 1)$  times. For a similar reason, a branching in the collective tree may occur at most  $n - 1$  times.

We classify the rounds into three types:

- (i) Rounds at the beginning of which no agent's view has a counting cut.
- (ii) Rounds at the beginning of which some agents' views have a counting cut, but there is no total agreement on any counting cut.
- (iii) Rounds at the beginning of which there is total agreement on a counting cut.

During rounds of type (i), according to our algorithm, no messages are discarded. In each of these rounds, if no branching occurs in the collective tree, a new counting cut is acquired by the collective tree, due to Lemma 8. If such a counting cut  $C$  already exists in the collective tree at the beginning of a round of type (i), then either  $C$  is lost by the collective tree (and a branching occurs, due to Lemma 7) or progress is made toward total agreement on  $C$  by the agents. Specifically, for every child  $v$  of a node in  $C$ , at every round of type (i) where  $C$  is not lost, at least one more agent's view acquires  $v$  (since the network is always connected and no messages are discarded). Thus, after  $n - 1$  (not necessarily consecutive) rounds of type (i) where  $C$  is not lost, total agreement on  $C$  is reached.

So, if no branching occurs, it takes  $n$  rounds of type (i) for a counting cut to be acquired by the collective tree and then be acquired by all agents' views. This cut may be lost during the process if a branching occurs in the collective tree, but this may happen only  $n - 1$  times. Therefore, there can be at most  $n^2$  (not necessarily consecutive) rounds of type (i) before total agreement on a counting cut is reached.

Let us consider now the rounds of type (ii). Let  $C$  be any undominated cut in an agent's view. Due to Lemma 7, in every view that loses  $C$  a branching must occur. Moreover, since the network is always connected, there is at least one agent  $p$  whose view contains  $C$  that

exchanges messages with an agent  $p'$  whose view does not contain  $C$ . Due to Lemma 11, these messages are not discarded. Hence, by to our previous analysis of agent interactions and by Lemma 9, a branching occurs in the view of  $p$  or in the view of  $p'$ , or the counting cut  $C$  is acquired by the view of  $p'$  and is not lost by the view of  $p$ . Recall that, by Lemma 10,  $C$  remains undominated in any view that acquires it, and therefore the same argument holds for the next rounds of type (ii).

Thus, at every round of type (ii), either the number of agents whose views contain  $C$  increases, or it remains stable and at least one branching occurs in the view of some agent, or it decreases by  $k > 0$ , and at the same time at least  $k$  branchings occur.  $C$  may be lost by all views, but then another undominated counting cut will replace it at the next round of type (ii). Since the total number of branchings in agents' views may be at most  $n(n-1)$ , we conclude that there can be at most  $2b_1 + n$  (not necessarily consecutive) rounds of type (ii), with  $b_1 \leq n(n-1)$ , before total agreement on a counting cut is reached. Indeed, the number of agents' views containing a same counting cut must start from 0 and reach  $n$ , and this number is decremented (and then incremented again)  $b_1 \leq n(n-1)$  times.

It follows that, eventually, total agreement on a counting cut  $C$  is reached, and the following round will be of type (iii). By Lemma 6,  $C$  is also a counting cut in the collective tree. Without loss of generality, we may assume that  $C$  is the dominant counting cut in the collective tree. For if it were not, the dominant counting cut (which exists by Lemma 5) would consist of ancestors of nodes of  $C$ , which are contained (as well as their children) in all agents' views.

It is easy to see that  $C$  is never lost by any agent's view, and therefore all subsequent rounds will be of type (iii). Indeed, in the collective view, all agents are represented by descendants of children of nodes in  $C$ . Hence, none of the new nodes that are added to the collective view may ever cause  $C$  to be lost as a counting cut. Similarly, no view may acquire such a node either, or the same node would also appear in the collective tree.

Observe that  $C$  is undominated in all views. Indeed, assume that a counting cut  $C'$  dominates  $C$  in some agent's view. Since  $C'$  consists of ancestors of nodes of  $C$ , and  $C$  appears in all views, it follows that  $C'$  is a counting cut in all views, and therefore it is a counting cut in the collective view as well (by Lemma 6), contradicting the fact that  $C$  is dominant in the collective view.

Consider an agent's view  $\mathcal{V}$  where  $C$  is not dominant. Since  $C$  is undominated, there must be another counting cut  $C'$  in  $\mathcal{V}$  that intersects  $C$ . There cannot be total agreement on  $C'$ , for otherwise  $C'$  would be a counting cut in the collective tree (by Lemma 6), contradicting the fact that different counting cuts in the collective tree are disjoint (by Lemma 5).

Since  $C$  is undominated in all views, no counting cut other than  $C$  can be dominant. Therefore, none of the messages between agents whose view contains  $C'$  and agents whose view does not contain  $C'$  are discarded (because such views cannot have the same dominant counting cut). Thus, as we argued for rounds of type (ii), at every round either the number of agents whose views contain  $C'$  increases, or it remains stable and at least one branching occurs in some view, or it decreases by  $k > 0$  and at least  $k$  branchings occur. Since there can never be total agreement on  $C'$ , any such counting cut will be lost by all views within  $2b_2 + n$  rounds, where  $b_2$  is the number of branchings that occur in agents' views during these rounds. Of course,  $b_1 + b_2 \leq n(n-1)$ .

Thus, eventually  $C$  becomes the dominant counting cut in all views; by Lemma 6, at this point all agents return the correct output. Moreover, since all agents have now the same dominant counting cut, they stop updating their states and keep returning the same correct output. In total, this takes at most  $(n^2) + (2b_1 + n) + (2b_2 + n) \leq 3n^2$  rounds. Note that a history tree with  $O(n^2)$  levels can be encoded in  $O(n^4 \log n)$  bits, implying that the algorithm is finite-state. ◀

## 6 Concluding Remarks

We have proposed the first self-stabilizing universal algorithm for anonymous dynamic networks; this algorithm has a linear stabilization time (Section 4). We have also provided the first finite-state universal algorithm for anonymous dynamic networks; this algorithm has a quadratic stabilization time (Section 5).

It is natural to ask whether a universal algorithm that is both self-stabilizing and finite-state exists. For static networks, such an algorithm is found in [2]; as for dynamic networks, we gave an optimal solution in Section 3 assuming that the number of agents is known.

Another open problem is to improve the stabilization times of our algorithms. For self-stabilizing algorithms in dynamic networks, we believe that a linear dependency on the garbage coefficient is unavoidable, and therefore our algorithm of Section 4 is asymptotically optimal. As for finite-state algorithms, we believe that a linear stabilization time is achievable.

It would be interesting to investigate disconnected networks, as well as universal algorithms that are not only finite-state, but also allow agents to stop sending messages after stabilization has been achieved.

---

### References

- 1 P. Boldi and S. Vigna. An Effective Characterization of Computability in Anonymous Networks. In *Proceedings of the 15th International Conference on Distributed Computing (DISC '01)*, pages 33–47, 2001.
- 2 P. Boldi and S. Vigna. Universal Dynamic Synchronous Self-Stabilization. *Distributed Computing*, 15:137–153, 2002. doi:10.1007/S004460100062.
- 3 A. Casteigts, F. Flocchini, B. Mans, and N. Santoro. Shortest, Fastest, and Foremost Broadcast in Dynamic Networks. *International Journal of Foundations of Computer Science*, 26(4):499–522, 2015. doi:10.1142/S0129054115500288.
- 4 J. Chalopin, S. Das, and N. Santoro. Groupings and Pairings in Anonymous Networks. In *Proceedings of the 20th International Conference on Distributed Computing (DISC '06)*, pages 105–119, 2006.
- 5 J. Chalopin, E. Godard, and Y. Métivier. Local Terminations and Distributed Computability in Anonymous Networks. In *Proceedings of the 22nd International Symposium on Distributed Computing (DISC '08)*, pages 47–62, 2008.
- 6 J. Chalopin, Y. Métivier, and T. Morsellino. Enumeration and Leader Election in Partially Anonymous and Multi-hop Broadcast Networks. *Fundamenta Informaticae*, 120(1):1–27, 2012. doi:10.3233/FI-2012-747.
- 7 G. A. Di Luna and G. Viglietta. Universal Finite-State and Self-Stabilizing Computation in Anonymous Dynamic Networks. [arXiv:2409.00688](https://arxiv.org/abs/2409.00688).
- 8 G. A. Di Luna and G. Viglietta. Computing in Anonymous Dynamic Networks Is Linear. In *Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science (FOCS '22)*, pages 1122–1133, 2022.
- 9 G. A. Di Luna and G. Viglietta. Optimal computation in leaderless and multi-leader disconnected anonymous dynamic networks. In *Proceedings of the 37th International Symposium on Distributed Computing (DISC '23)*, pages 18:1–18:20, 2023.
- 10 P. Fraigniaud, A. Pelc, D. Peleg, and S. Pérennes. Assigning Labels in Unknown Anonymous Networks. In *Proceedings of the 19th ACM Symposium on Principles of Distributed Computing (PODC '00)*, pages 101–111, 2000.
- 11 F. Kuhn, T. Locher, and R. Oshman. Gradient Clock Synchronization in Dynamic Networks. *Theory of Computing Systems*, 49(4):781–816, 2011. doi:10.1007/S00224-011-9348-1.



- 12 F. Kuhn, N. Lynch, and R. Oshman. Distributed Computation in Dynamic Networks. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC '10)*, pages 513–522, 2010.
- 13 F. Kuhn, Y. Moses, and R. Oshman. Coordinated Consensus in Dynamic Networks. In *Proceedings of the 30th ACM Symposium on Principles of Distributed Computing (PODC '11)*, pages 1–10, 2011.
- 14 O. Michail and P. G. Spirakis. Elements of the Theory of Dynamic Networks. *Communications of the ACM*, 61(2):72, 2018. doi:10.1145/3156693.
- 15 A. Nedić, A. Olshevsky, A. E. Ozdaglar, and J. N. Tsitsiklis. On Distributed Averaging Algorithms and Quantization Effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009. doi:10.1109/TAC.2009.2031203.
- 16 R. O'Dell and R. Wattenhofer. Information Dissemination in Highly Dynamic Graphs. In *Proceedings of the 5th Joint Workshop on Foundations of Mobile Computing (DIALM-POMC '05)*, pages 104–110, 2005.
- 17 J. Seidel, J. Uitto, and R. Wattenhofer. Randomness vs. Time in Anonymous Networks. In *Proceedings of the 29th International Symposium on Distributed Computing (DISC '15)*, pages 263–275, 2015.
- 18 G. Viglietta. History Trees and Their Applications. In *Proceedings of the 31st International Colloquium on Structural Information and Communication Complexity (SIROCCO '24)*, pages 3–23, 2024.
- 19 M. Yamashita and T. Kameda. Computing on an Anonymous Network. In *Proceedings of the 7th ACM Symposium on Principles of Distributed Computing (PODC '88)*, pages 117–130, 1988.