

History Trees and Their Applications

Giovanni Viglietta

University of Aizu, Aizuwakamatsu, Japan viglietta@gmail.com
<https://giovanniviglietta.com>

Abstract. In the theoretical study of distributed communication networks, *history trees* are a discrete structure that naturally models the concept that anonymous agents become distinguishable upon receiving different sets of messages from neighboring agents. By conveniently organizing temporal information in a systematic manner, history trees have been instrumental in the development of optimal deterministic algorithms for networks that are both anonymous and dynamically evolving. This note provides an accessible introduction to history trees, drawing comparisons with more traditional structures found in existing literature and reviewing the latest advancements in the applications of history trees, especially within dynamic networks. Furthermore, it expands the theoretical framework of history trees in new directions, also highlighting several open problems for further investigation.

Keywords: history tree · anonymous network · dynamic network

1 Introduction

A *distributed communication network* consists of a finite number of independent computational units known as *agents*, which can send each other messages and modify their internal states based on the messages they receive. These networks generally operate in synchronous or asynchronous *communication steps*, where agents send messages through the links of a (multi)graph, which may be static or dynamic, directed or undirected, connected or disconnected, etc.¹

A network is *anonymous* if its agents are initially indistinguishable, i.e., they lack unique identifiers and can only be told apart by external input or due to the network’s layout. In such networks, it is typically assumed that all agents run the same local *deterministic* algorithm. It is important to note that allowing for randomness would enable agents to generate unique identifiers with high probability, thereby compromising the study’s focus on anonymity.

A common algorithmic approach in this setting involves assigning a “weight” to each agent, which is then distributed among neighbors at each step according to specific rules. This process is analyzed using standard stochastic methods to understand how these weights eventually converge to a common value [6,13].

¹ Although this note does not focus on *population protocols*, they can also be modeled within this framework as dynamic communication networks where finite-state agents interact in pairs consisting of an initiator and a responder having different roles.

Although this “averaging technique” enabled the development of general algorithms for anonymous networks, one of its major limitations is its disregard for the network’s structural and topological characteristics, providing minimal insight into these aspects. Furthermore, analyzing these algorithms tends to be technically cumbersome, yielding only asymptotic estimates for how quickly the network reaches convergence or stability.

Another line of research, initiated by Angluin [1], explores discrete and algebraic structures, such as the *views* of Yamashita–Kameda [17] and the *graph fibrations* and *minimum bases* of Boldi–Vigna [4]. This approach allows algorithms to fully leverage the network’s structure and usually permits a more precise analysis of running times. However, the theoretical frameworks of views and graph fibrations were developed for networks with unchanging topologies, and their successful application has been limited to such static networks.

History trees are a newer data structure that inherently includes a temporal dimension and was specifically designed for networks with dynamic topologies. The introduction of history trees has recently led to the development of optimal linear-time algorithms for anonymous dynamic networks [9,11] and state-of-the-art general algorithms for congested anonymous dynamic networks [10].

In Section 2, we outline the basic architecture of history trees, drawing comparisons with Yamashita–Kameda’s views and Boldi–Vigna’s minimum bases. Additionally, we showcase a straightforward linear-time algorithm for dynamic networks that illustrates the practicality of history trees. In Section 3, we discuss more advanced applications of history trees in challenging network situations, while also pointing out a number of areas that remain open for investigation.

2 Basic Structure and Algorithms

In this section we focus on anonymous networks operating in synchronous steps, modeled as undirected dynamic multigraphs with no port awareness. Other network models will be discussed in Section 3.

Throughout this section, the reader may find it beneficial to examine the software available at <https://github.com/viglietta/Dynamic-Networks>. The repository includes a simulator that supports the creation and visualization of history trees of user-defined dynamic networks and the testing of fundamental algorithms.

2.1 History Trees and Construction of a View

Consider a dynamic network of six anonymous agents whose first four communications steps are as in Figure 1. Before the first step (i.e., at time $t = 0$), each agent only knows its own input, which in our example is represented by a color: either cyan or yellow. In the first communication step, both cyan agents (which are indistinguishable) receive three messages from yellow agents (also indistinguishable). Thus, the cyan agents remain indistinguishable at time $t = 1$.

However, two yellow agents receive two messages from cyan agents, but the other two receive only one message. Thus, at time $t = 1$, the two yellow agents

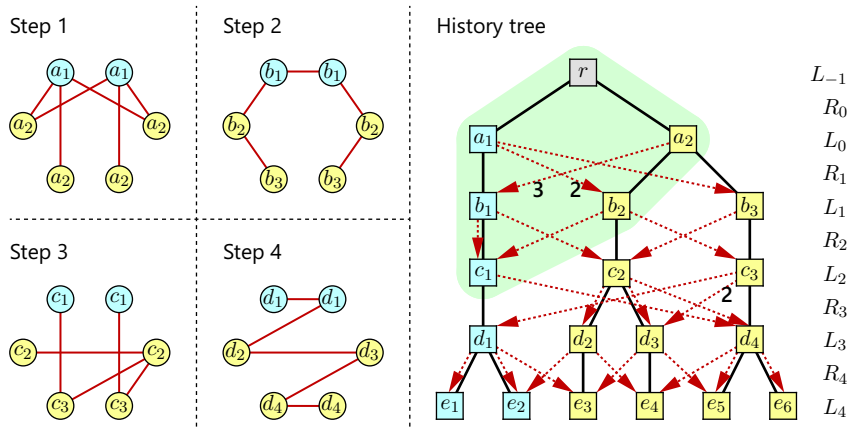


Fig. 1. The first communication steps of a dynamic network and the corresponding levels of its history tree. Initial inputs are represented by agents' colors (cyan and yellow). Labels on agents and nodes have been added for the reader's convenience and indicate classes of indistinguishable agents. The portion of the history tree with a green background is the view of the two agents labeled c_1 after two communication steps.

labeled b_2 are still indistinguishable from each other, but are distinguishable from the ones labeled b_3 , because they have a different “history”.

In general, since the network is synchronous, at step $t > 0$ (i.e., between time $t - 1$ and time t) all agents simultaneously send messages and simultaneously receive them (as unordered multisets). Two messages are assumed to be identical if and only if they are sent by indistinguishable agents (possibly by the same agent). By definition, at time $t = -1$ no two agents are distinguishable; at time $t = 0$ they are distinguishable if and only if they have different inputs. At time $t > 0$, two agents are distinguishable if and only if they are distinguishable at time $t - 1$ or receive different multisets of messages at step t .

History trees were introduced in [9] to study how and when anonymous agents in a (dynamic) network become distinguishable. The history tree \mathcal{H}_G of a network G is an infinite tree subdivided into *levels*, as in Figure 1, where the nodes in level L_t represent the equivalence classes of agents that are indistinguishable at time t . The number of agents in the class represented by a node v is called the *anonymity* of v and is denoted as $\mathbf{a}(v)$. In our example, $\mathbf{a}(a_1) = 2$ and $\mathbf{a}(a_2) = 4$.

The root r represents all agents in the network. In general, the *children* of a node $v \in L_{t-1}$, i.e., the nodes $v_1, v_2, \dots, v_k \in L_t$ that are connected to v by a *black edge*, represent a partition of the set of agents represented by v . Thus,

$$\mathbf{a}(v) = \sum_{i=1}^k \mathbf{a}(v_i). \quad (1)$$

On the other hand, the *red edges* in R_t represent messages sent and received at step t . That is, a directed red edge (v, u) with *multiplicity* m , where $v \in L_{t-1}$




Fig. 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 matched and merged starting from the roots, v gets a child v' , and a new red edge from u to v' is added.

and $u \in L_t$, indicates that, at step t , each agent represented by node u receives exactly m (identical) messages from agents represented by v .

Assuming that their internal memory and message sizes are unbounded, the agents in a network G can locally construct portions of the history tree called *views*. More precisely, if an agent p at time t is represented by a node $v \in L_t$, then the view $\mathcal{V}_G^t(p)$ is defined as the portion of \mathcal{H}_G spanned by all directed paths from the root r to v , using black and red edges indifferently (black edges are assumed to be directed away from the root r). The node v is called the *bottom* of the view. Figure 1 shows an example of a view within a history tree.

The distributed construction of views can be achieved via an iterative process, assuming that all agents send their current view to all their neighbors at every step. Upon receiving a multiset M of messages, an agent p integrates its current view $\mathcal{V}_G^t(p)$ with all views in M . This is done by a straightforward match-and-merge algorithm, starting from the root and proceeding level by level, as exemplified in Figure 2. The result is the smallest tree containing $\mathcal{V}_G^t(p)$ and all views in M as induced subtrees. To obtain $\mathcal{V}_G^{t+1}(p)$, the agent p creates a child v' for the bottom node v of $\mathcal{V}_G^t(p)$, and connects the bottom nodes of all views in M to v' by red edges with the appropriate multiplicities. The node v' is now the bottom of $\mathcal{V}_G^{t+1}(p)$ and represents p (and possibly other agents) at time $t + 1$.

It is a simple observation that the view $\mathcal{V}_G^t(p)$ contains all the information that p can possibly extract from the network G after t communication steps [9].

2.2 Related Structures for Static Networks

We will now focus on *static* anonymous networks, which are well understood thanks to the works of Yamashita–Kameda, Boldi–Vigna, and other authors.

If G is a static network of n agents, once the classes of indistinguishable agents remain unchanged for a single communication step, they must remain unchanged forever. Thus, the number of nodes in the levels of \mathcal{H}_G strictly increases at

every level until it stabilizes, say, at level L_s . Therefore, if $|L_s| = \hat{n}$, we have $0 \leq s < \hat{n} \leq n$, because $|L_0| \geq 1$ and each node represents at least one agent.

We come to the conclusion that in any network of size n , if two agents are indistinguishable at time $t = n - 1$, they will remain indistinguishable thereafter. This observation was first made by Norris (though formulated differently) in [14], where she also raised the question of whether the same applies to agents in two distinct networks of the same size n .² We restate Norris' question below.

Open Problem 1. Let G and G' be two disjoint static networks of n agents each. If an agent in G and an agent in G' have isomorphic views at time $t = n - 1$, do they have isomorphic views at all times?

Let us consider the subgraph of \mathcal{H}_G induced by the nodes in levels L_s and L_{s+1} , as highlighted in Figure 3 with a green background. Contracting the black edges in this subgraph, we obtain a directed multigraph named \hat{G} .

This is an important structure that has been given many names and equivalent definitions in the literature. Essentially, \hat{G} could be defined as the smallest directed multigraph that gives rise to a history tree isomorphic to \mathcal{H}_G . Yamashita and Kameda call \hat{G} the *quotient graph* of G and define it in graph-theoretical terms using their own notion of *view* of a static network [16,17], whereas Boldi and Vigna call \hat{G} the *minimum base* of G and give it a topological definition in terms of *graph fibrations* [2,3,4,5].³

Notably, since the results concerning static networks from these authors are presented with reference to \hat{G} , they can be directly rephrased in the language of history trees. The advantage of using history trees lies in their ability to readily offer timing information on when agents become distinguishable. This feature makes history trees particularly suitable for dynamic networks, where \hat{G} is not well defined and the lack of topological regularity prevents the straightforward deduction of temporal data.

Before delving into applications of history trees to dynamic networks in Section 2.3, let us draw a closer comparison between views of history trees and views in the sense of Yamashita–Kameda [17]. While it is important not to confuse these two concepts, they bear a deep relationship illustrated in Figure 3.

For Yamashita and Kameda, the view $\mathcal{T}_G(p)$ of an agent p in a static network G is an infinite tree rooted at p , where each node of depth k represents a distinct

² By treating the disjoint union of the two networks as a single network of size $2n$, it becomes evident that agents that are indistinguishable at time $t = 2n - 1$ are indistinguishable forever. This point was already noted by Conway in [8, Chapter 1, Theorem 7], albeit phrased in terms of Moore machines. Although $t = 2n - 1$ might not be an optimal bound if both networks have size n (hence Norris' question), Figure 4 implies that $t = 2n$ is optimal if they have sizes n and $n + 1$, respectively. Indeed, if the two cyan agents are colored yellow, they become distinguishable after $2n$ steps, matching the upper bound given by the combined number of agents, $2n + 1$.

³ More specifically, G is fibered onto the fibration-prime graph \hat{G} , and the fibers are precisely the classes of agents represented by the nodes of L_s in the history tree \mathcal{H}_G .

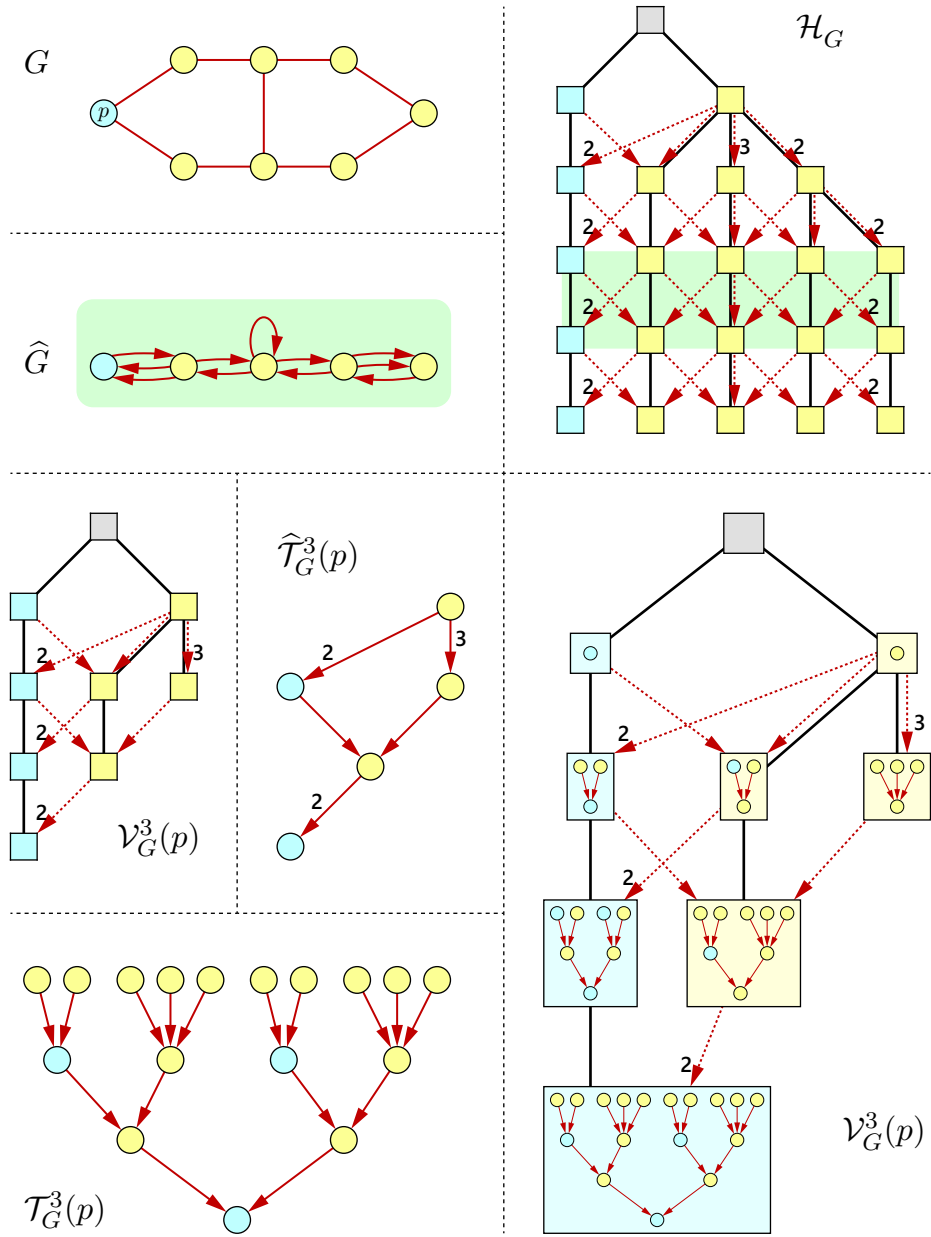


Fig. 3. A static network G with a distinguished agent p and its history tree \mathcal{H}_G (truncated at level L_4). At stabilization, the directed graph of the red edges between non-branching levels (highlighted in green in \mathcal{H}_G) is isomorphic to the Boldi–Vigna minimum base \hat{G} . The view $\mathcal{V}_G^3(p)$ after step 3 is related to the Yamashita–Kameda view $\mathcal{T}_G^3(p)$ truncated at depth 3, as well as to its Tani folded view $\hat{\mathcal{T}}_G^3(p)$, as illustrated. The precise correspondence between these structures is described in Section 2.2.

walk of length k in G terminating in p . As it turns out, $\mathcal{T}_G(p)$ truncated at depth k , denoted as $\mathcal{T}_G^k(p)$, contains precisely the information encoded in $\mathcal{V}_G^k(p)$.

Indeed, starting from $\mathcal{V}_G^k(p)$, we can construct $\mathcal{T}_G^k(p)$ as follows. Let us consider the subgraph spanned by the directed paths in $\mathcal{V}_G^k(p)$ terminating in the bottom node using only red edges. Such a subgraph is a directed acyclic graph isomorphic to the so-called *folded view* $\widehat{\mathcal{T}}_G^k(p)$. This is a structure devised by Tani in [15] to compress $\mathcal{T}_G^k(p)$ from exponential size to polynomial size by conflating equivalent nodes. Finally, $\widehat{\mathcal{T}}_G^k(p)$ can easily be unraveled to reconstruct $\mathcal{T}_G^k(p)$.

Conversely, starting from $\mathcal{T}_G^k(p)$, one can construct an inventory of “fragments” by listing all subtrees hanging from different nodes and truncating them at all possible depths. The isomorphism classes of these fragments, when ordered by their height, correspond to the levels of $\mathcal{V}_G^k(p)$, with the “empty fragment” being the root. Then, the parent of a fragment f in $\mathcal{V}_G^k(p)$ is determined by deleting all the leaves of f . Similarly, the fragments of $\mathcal{V}_G^k(p)$ that are connected to f via red edges (and the multiplicities of such red edges) match the subtrees that dangle from the root’s children within f . Thus, we can construct $\mathcal{V}_G^k(p)$ from $\mathcal{T}_G^k(p)$.

2.3 Basic Applications to Dynamic Networks

We will now describe a general algorithmic technique that can be used to solve a wide range of fundamental problems in networks operating in synchronous steps, modeled as *dynamic* undirected multigraphs that are connected at every step. Albeit being extremely straightforward, this technique achieves optimal running times and matches in efficiency the best algorithms for static networks.

As argued in Section 2.1, the agents in a network have a distributed algorithm for constructing their view of the history tree and update it at every communication step. Since an agent’s view encodes all the information that the agent can infer from the network, in principle we can reduce any problem about anonymous networks to a problem about views of history trees.

Of course, this is true assuming that all agents have enough internal memory and can send large-enough messages to store their views. After t steps, the size of a view is $O(tn^2 \log M)$ bits, where n is the total number of agents and M is the maximum number of messages sent by any agent in a single step. Indeed, such a view has t levels, each of which contains at most n nodes and n^2 incoming red edges of multiplicity at most M . In particular, if the view construction algorithm runs for a polynomial number of steps, it requires internal memory and messages of polynomial size. For the time being, we will assume that this is not an issue.

The basic technique relies on the following observation. If two nodes u and v in the same level L_i of the history tree are *non-branching*, i.e., they have a unique child u' and v' respectively, and there are red edges (u, v') with multiplicity $m_{u,v'} > 0$ and (v, u') with multiplicity $m_{v,u'} > 0$, then we can count the number of messages exchanged by the agents represented by u and by v as

$$m_{v,u'} \mathbf{a}(u) = m_{u,v'} \mathbf{a}(v) \tag{2}$$

because communication links are bidirectional. Thus, since we know the multiplicities of the red edges, we can infer the ratio of the anonymities involved.

Now, if L_i is a level where all nodes are non-branching, and since the network is connected at every step by assumption, we can repeatedly apply Equation (2) to compute the ratios of the anonymities of all nodes in L_i . For example, since level L_1 in Figure 1 is non-branching, we can easily deduce that $\mathbf{a}(b_1) = \mathbf{a}(b_2) = \mathbf{a}(b_3)$. Then, using Equation (1), we can extend the computation to all previous levels, and obtain for instance that $2\mathbf{a}(a_1) = \mathbf{a}(a_2)$. The reader may apply the same method to the history tree in Figure 3 to deduce that the ratio between the number of yellow agents and the number of cyan agents is 7.

Once we have this information, we can solve problems such as *Average Consensus*: if the agents are given input numbers, we can compute the fraction of agents that have each input and use it as a weight to compute the mean input.

If we are given additional information, such as the total number of agents or the number of agents that have a certain input, we can also compute how many agents have each input. For instance, in Figure 1, if we know that there are two cyan agents, we deduce that there are four yellow ones, because $2\mathbf{a}(a_1) = \mathbf{a}(a_2)$.

In particular, if we know that the network contains a unique distinguished agent, typically referred to as a *leader*, then we can solve the *Counting problem*, which asks to compute the total number of agents, n . Thus, in Figure 3, if we know that p is the unique leader, we can deduce from the history tree that $n = 8$.

Let us discuss the efficiency of this method. Since $|L_{-1}| = 1$ and the number of nodes per level is at most n , it follows that it is sufficient to inspect the history tree up to level L_{n-1} to find a non-branching level and the relative red edges.

A subtler issue is how long it takes before the views of all agents acquire all nodes (and hence all edges) up to L_{n-1} , so that all agents can carry out correct computations locally. In Figure 1, for instance, since level L_0 is non-branching in the highlighted view (because node b_3 is not in the view), at time $t = 2$ the cyan agents may be deceived and incorrectly deduce that $3\mathbf{a}(a_1) = 2\mathbf{a}(a_2)$.

Recall that the network is assumed to be connected at all steps. So, it takes fewer than n steps for information to travel from any agent to any other agent; in other words, the *dynamic diameter* d is at most $n - 1$. Indeed, if $k < n$ agents have some information at time t , then at least $k + 1$ agents have it at time $t + 1$.

Hence, it takes at most d steps for a node in the history tree to appear in the views of all agents. In particular, at time $n + d - 1 \leq 2n - 2$, all agents have the entirety of L_{n-1} in their views, and thus have enough information to do correct computations. For example, in Figure 3, knowing all levels of \mathcal{H}_G up to L_3 is enough, but these levels entirely appear in the view of p only at step 7. In fact, the diameter of the network G is $d = 4$.

This technique yields an algorithm for Average Consensus that *stabilizes* in $2n - 2$ steps. That is, if all agents attempt to compute the average of the input values after every step, they may guess it incorrectly for some time, but will be all correct after $2n - 2$ steps at the latest. The same upper bound holds for the Counting problem, assuming the presence of a unique leader in the network.

A downside of this method is that it provides no certificate of correctness, and therefore the algorithm is supposed to run indefinitely. However, if n is known to the agents, they can simply count $2n - 2$ steps, do all computations on the




Fig. 4. A static network with two connected components of sizes n and $n + 1$. The history tree shows that the two agents in cyan become distinguishable after $2n - 1$ communication steps. Therefore, in static networks of n agents with a unique leader, no Counting algorithm can terminate in $2n - 2$ steps or stabilize in $2n - 4$ steps.

first $n - 1$ levels, and then terminate returning the correct output. Similarly, if an agent knows the dynamic diameter d , it can terminate at time t if its view contains a non-branching level L_i with $i < t - d$, because all levels up to L_{i+1} are guaranteed to be entirely in the view. This occurs by the time $t = n + d - 1$.

Finally, let us discuss lower bounds. Figure 4 shows two networks of sizes n and $n + 1$ whose leaders have the same view up to step $2n - 2$. Since agents with equal views must return equal outputs (assuming they execute the same algorithm), it easily follows that there is no Counting algorithm that terminates in $2n - 2$ steps or stabilizes in $2n - 4$ steps in both networks. For if there were such an algorithm, then both leaders would return the same output, and at least one of them would be incorrect. Note that this almost matches the running time of the stabilizing Counting algorithm given above, which is $2n - 2$ steps. This leaves only a small gap: namely, whether $2n - 3$ steps are actually sufficient.

Open Problem 2. Can a Counting algorithm stabilize in $2n - 3$ steps in all connected undirected dynamic networks with a unique leader?

As for Average Consensus, if the two cyan agents in Figure 4 are colored yellow, then they have the same view up to step $2n - 1$. This implies that no Average Consensus algorithm can stabilize in $2n - 3$ steps in all networks, and therefore the one described above, which stabilizes in $2n - 2$ steps, is optimal.

3 Variations and Extensions

3.1 Leader Election

Another application of the technique in Section 2.3 is *Leader Election*, where all agents have to agree on a unique representative to be identified as the “leader”. Of course, this problem has a solution only if the history tree contains a node of anonymity 1; as it turns out, this condition is also sufficient.

A simple Leader Election algorithm, at step t , computes the ratios of all anonymities in the first non-branching level occurring after $L_{\lfloor t/2 \rfloor}$ and deterministically picks a node of smallest anonymity as representing the leader. This strategy eventually elects a unique leader if and only if it is possible to do so, but no certificate of correctness is provided, and the algorithm has to run indefinitely.

However, if n is known, then anonymities in non-branching levels can be computed exactly with a known delay of $n - 1$ steps. In this case, any agent that becomes distinguishable from all others at time t can be identified with certainty by time $t + 2n - 2$, allowing all agents to terminate.

3.2 Terminating Counting

In Section 2.3 we gave a Counting algorithm for connected undirected dynamic networks with a unique leader that stabilizes in $2n - 2$ steps. At the cost of n additional communication steps, we can implement a correctness certificate and make all agents return n and terminate. The algorithm’s details are somewhat intricate and can be found in [9]; here we will merely sketch the main ideas.

If we know the anonymities of a node u and of all its children, then u is called a *guesser*. If there is a red edge directed from u to a node v of unknown anonymity, we can write an equation similar to Equation (2) to count messages exchanged by the corresponding agents. Solving this equation yields a *guess* $\mathbf{g}(v)$ on $\mathbf{a}(v)$ in terms of known anonymities and multiplicities of red edges. As it turns out, $\mathbf{g}(v) \geq \mathbf{a}(v)$ always holds; moreover, if v has no siblings, $\mathbf{g}(v) = \mathbf{a}(v)$.

Clearly, the nodes representing the leader are always guessers, because their anonymity is 1. Using these nodes, we can make initial guesses on at least one node per level. The question is how we can determine which guesses are correct without knowing the whole history tree, but only a view of it.

Let us define the *weight* $\mathbf{w}(v)$ of a node v as the number of guesses that have been made on nodes in the subtree hanging from v (including on v itself). Then, v is said to be *heavy* if $\mathbf{w}(v) \geq \mathbf{g}(v)$. The key observation is that, if guesses are *well spread*, i.e., no two sibling nodes are assigned a guess, then the deepest heavy node necessarily has a correct guess.

Hence, in Figure 5, the nodes in (a)–(d) with a blue number have a correct guess, but the one in (e) may not. For (a), this is obvious: since $\mathbf{g}(v) = 1$ and $\mathbf{g}(v) \geq \mathbf{a}(v)$, we must conclude that $\mathbf{a}(v) = 1$ and the guess is correct. As for (b), if the lower node u has a correct guess, then the upper node v has a correct guess too, because $2 = \mathbf{g}(v) \geq \mathbf{a}(v) \geq \mathbf{a}(u) = \mathbf{g}(u) = 2$, hence $\mathbf{g}(v) = \mathbf{a}(v)$. Otherwise, the guess on u is incorrect, and so u must have a sibling (perhaps not in the




Fig. 5. The numbers in the nodes are guesses made according to the algorithm in Section 3.2. Blue numbers indicate heavy nodes; their guesses are necessarily correct except for the one in (e), where guesses are not well spread.

view). Hence, $2 = \mathbf{g}(v) \geq \mathbf{a}(v) \geq 2$, and the guess on v is correct. The reader may verify that a similar line of reasoning applies to (c) and (d) but not to (e).

It is not difficult to see that, if there are at least n well-spread nodes with guesses, then some of them must be heavy, and thus the deepest of them is necessarily correct. In this way, we can steadily determine the anonymities of new nodes. Eventually, some of these nodes become guessers themselves and can be used to make new guesses, and so on. Since the network is connected at all steps, it can be shown that only $2n - 2$ levels of the history tree are required for this algorithm to eventually identify nodes with correct guesses on all branches.

Adding up the anonymities of nodes on all branches produces an estimate n' on the actual size of the network, n . In general we have $n' \leq n$, because the current view may not include all branches of the history tree. To confirm this estimate, it is sufficient to wait an additional n' steps, which is the longest time it may take for at least one missing branch to appear in the view, if one exists.

Overall, this Counting algorithm terminates in $3n - 2$ steps in the worst case, leaving a small gap with the lower bound of $2n - 2$ steps given in Section 2.3.

Open Problem 3. Can a Counting algorithm terminate in $2n + O(1)$ steps in all connected undirected dynamic networks with a unique leader?

The previous Counting algorithm was generalized in [11] to networks where the leader may not be unique, but there is a known number $\ell \geq 1$ of indistinguishable leaders, or *supervisors*. Note that the history tree may contain several branches corresponding to leaders; the challenge posed by this scenario is that the individual anonymities of such branches are unknown (although their sum is ℓ).

The state-of-the-art solution given in [11] involves subdividing the history tree into ℓ intervals, each of at most $(\ell + 1)n$ levels. The total running time is roughly $(\ell^2 + \ell + 1)n$ steps, making this algorithm impractical unless ℓ is small.

In fact, it is a major open question whether a larger ℓ may actually be helpful.

Open Problem 4. In connected undirected dynamic networks with a known number ℓ of leaders, how does the optimal running time of terminating Counting algorithms scale with ℓ ?




Fig. 6. Two directed static networks with a unique leader and no outdegree awareness that have a different number of agents but the same history tree.

3.3 Directed Networks

Generalizing the scenario of Section 2, we may consider *directed networks*, i.e., networks with unidirectional links. It is important to recognize that, in these networks, even solving basic problems such as Average Consensus or Counting with a unique leader requires some additional assumptions, as Figure 6 demonstrates.

A reasonable approach is to assume that, in each step, agents are aware of their outdegree either before (*early outdegree awareness*) or after (*late outdegree awareness*) they transmit their messages. The key distinction is that in the early awareness model, an agent can incorporate its current outdegree into the messages it sends. Here we will focus on the late outdegree awareness model, which is less advantageous especially in dynamic networks.

The basic history tree architecture can be extended by attaching outdegrees to black edges, as in Figure 7. We can then identify a non-branching level L_i in the history tree, as in Section 2.3, and use outdegrees and multiplicities of red edges to write linear equations generalizing Equation (2), as shown in Figure 7.

The resulting homogeneous system is represented by an *irreducible* matrix A , assuming the network is strongly connected at every step. Note that $A = \lambda I - P$, where $\lambda > 0$ and $P \geq 0$ is also an irreducible matrix. In our example,

$$A = \begin{bmatrix} 3 & -1 & 0 & 0 \\ 0 & 2 & -1 & -2 \\ -2 & 0 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix} - \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 2 \\ 2 & 0 & 2 & 0 \\ 0 & 0 & 1 & 2 \end{bmatrix} = 3I - P.$$

We have $Ax = 0$, where $x > 0$ is the vector of anonymities of the nodes in L_i . Hence, $Px = \lambda x$. By the Perron-Frobenius theorem, λ is a *simple* eigenvalue of P , and thus 0 is a simple eigenvalue of A . In other words, the *nullity* of A is 1, which means we can solve the linear system in terms of a single free variable.

Thus, we can find the ratios between the anonymities of all nodes in L_i , or compute them exactly if there is a known number $\ell \geq 1$ of leaders in the network.

This technique yields stabilizing Average Consensus and Counting algorithms for strongly connected directed dynamic networks with late outdegree awareness, generalizing the ones in Section 2.3. The stabilization time of these algorithms is again $2n - 2$ steps, which is optimal.




Fig. 7. A communication step of a directed network with late outdegree awareness and a non-branching level of its history tree. Blue numbers indicate outdegrees.

We now present a *terminating* Counting algorithm for directed networks with late outdegree awareness and a known number $\ell \geq 1$ of leaders. Every agent waits until its view contains a long-enough interval \mathcal{I} of non-branching levels. The goal is to find an upper bound U_i on the anonymity of a node v_i in each branch B_i within \mathcal{I} . To start, we take $U_1 = \ell$, where B_1 is any leader branch.

The algorithm repeatedly uses branches with a known upper bound to determine new upper bounds. Namely, let B_i have a known upper bound $U_i \geq \mathbf{a}(v_i)$ and consider all red edges (v, v') , where v is a descendant of v_i within \mathcal{I} . Each of these yields an *estimate* δU_i on $\mathbf{a}(v')$, where δ is the outdegree corresponding to the child of v in B_i . Note that, if v has a unique child in the history tree, at most δU_i messages are sent from agents represented by v , and so $\delta U_i \geq \mathbf{a}(v')$. Nonetheless, even if v is not branching in a view, it may still branch in the history tree, and in this case δU_i may not be a correct upper bound. However, this undesirable event may occur at most $U_i - 1$ times. Thus, as soon as a branch B_j receives estimates from B_i on U_i nodes, we take the maximum estimate as a correct upper bound U_j on the anonymity of the *deepest* such node v_j in B_j .

When we finally have an upper bound for all branches in \mathcal{I} , we wait $\sum_i U_i$ additional steps to confirm that there are no branches missing from the view. Then we run the previous stabilizing Counting algorithm and output the result.

Assuming that the network is simple and strongly connected at all steps, this algorithm terminates in $2^{O(n \log n)}$ steps, which is likely far from optimal.

Open Problem 5. Can a Counting algorithm terminate in a *polynomial* number of steps in all strongly connected directed dynamic simple networks with (early or late) outdegree awareness and a unique leader?

3.4 Disconnected Networks

For networks that are not necessarily connected at all steps, we define a *communication round* as a minimal sequence of consecutive steps whose communication multigraphs have a (strongly) connected *sum* (constructed by adding together their adjacency matrices). A network is τ -*union-connected* if every block of τ consecutive steps contains a communication round. It was observed in [11] that τ and the dynamic diameter d are related by the tight inequalities $\tau \leq d \leq \tau(n-1)$.




Fig. 8. Two agents in a semi-synchronous network share their views \mathcal{A} and \mathcal{B} , which are updated as \mathcal{A}' and \mathcal{B}' . Note that red edges may go upward or skip multiple levels. The highlighted nodes are the “bottom nodes” of their respective views, which can be identified as the only sinks (assuming black edges are directed away from the root).

Clearly, any non-trivial *terminating* computation is impossible if the agents know nothing about τ . On the other hand, with knowledge of τ , all of the previous algorithms can be straightforwardly adapted to τ -union-connected networks. It is sufficient for each agent to accumulate incoming messages at every step, updating its view only once every τ steps. This adaptation slows down all running times by a factor of τ . However, this is worst-case optimal, since a τ -union-connected network may be devoid of links except at steps that are multiples of τ .

As for *stabilizing* computations, they can be done even with no knowledge of τ , again with a worst-case optimal slowdown by a factor of τ , as detailed in [11].

3.5 Semi-Synchronous Networks

In a network operating *semi-synchronously*, any agent may unpredictably be *inactive* at any step. An inactive agent does not communicate with other agents and does not even update its state. This model is called “asynchronous” by Boldi-Vigna [3,5] and generalizes the *asynchronous starts* of Charron-Bost–Moran [7], but we will give the term “asynchronous” a more extensive meaning in Section 3.6.

The concept of a round and the parameter τ are defined as in Section 3.4. A notable distinction from the synchronous models discussed so far is that agents in a semi-synchronous network cannot reliably count communication steps, because they do not know for how many steps they have been inactive. Hence, two views shared by two agents may not have the same height, as the ones in Figure 8.

Consequently, a red edge no longer has to connect a level to the next, but can span any number of levels. However, it is always possible to restore this property by adding *dummy nodes* to represent inactive agents, resulting in an *equalized* view (Figure 9). Thus, we can reduce any semi-synchronous network to an equivalent synchronous one with the same parameter τ and apply the methods outlined in Section 3.4 to obtain algorithms with the same running times.




Fig. 9. Equalizing the view of an agent in a semi-synchronous network by adding dummy nodes. In the resulting view, every red edge connects a level to the next.

3.6 Asynchronous Networks

A network is *asynchronous* if messages can take an arbitrary, independent, and unpredictable amount of time to reach their destinations. Such a network is necessarily directed, and therefore some form of outdegree awareness is needed, as argued in Section 3.3. A *round* is now any minimal interval of time such that the messages that are sent and received form a strongly connected multigraph.

Observe that, even in asynchronous networks, time can always be discretized, assuming that no agent can send an infinite number of messages in a finite time.

Also, doing non-trivial computations with termination in asynchronous networks with no knowledge about the duration of a round is impossible.

However, there is a simple stabilization algorithm, as follows. When an agent sends some messages, it expands its view by adding a child to the bottom node and attaches its outdegree to the corresponding black edge, as in Section 3.3. When it receives a multiset of messages, it updates its view as in Section 3.5.

Then, the agent seeks the first interval of levels constituting a round where all nodes are non-branching. The outdegrees and red edges in this interval yield a system of equations that is solved as in the stabilizing algorithm of Section 3.3.

Such a non-branching interval occurs in the history tree after at most $n - 1$ rounds and appears in the agent's view in at most $n - 1$ additional rounds. The total stabilization time is therefore $2n - 2$ rounds, which is worst-case optimal.

3.7 Port Awareness

A widely studied scenario, especially within static networks, is when each agent has a local numbering for its incoming links (*input port awareness*) or outgoing links (*output port awareness*). It is evident that, although input port awareness enables each agent in a static network to identify all messages sent by the same neighbor, this feature has no clear meaning or effect in dynamic networks.

In contrast, output port awareness has a significant impact on both static and dynamic networks, as it not only implies outdegree awareness, but allows agents to assign a different tag to each message they send within a communication step. This is helpful in breaking network symmetry, thus facilitating certain computations. The history tree architecture for outdegree awareness of Section 3.3 can also be adapted to this model by simply attaching a port number to each red edge.

As an example, output port awareness combined with a unique leader makes the Counting problem particularly simple even in strongly connected directed dynamic networks. Indeed, all messages sent by the leader in a communication step have different tags, and so all agents that receive them become distinguishable. Generalizing, every node on a red path starting at a leader node must have an anonymity of 1. When a whole level L_i in a view consists of such nodes, we can check if the outdegree of each node in L_i matches the number of its outgoing red edges. If so, all agents have been accounted for, and we can return $n = |L_i|$.

This Counting algorithm terminates in $2n - 1$ steps, greatly improving upon the one in Section 3.2. Observe that the lower bound in Figure 4 does not hold for networks with output port awareness, which may allow for even faster solutions.

Open Problem 6. Can a Counting algorithm terminate in fewer than $2n - 2$ steps in all strongly connected directed dynamic networks with output port awareness and a unique leader?

A less obvious fact is that, in any non-branching level of the history tree of a strongly connected directed dynamic network with output port awareness, all nodes must have the same anonymity.⁴ Indeed, if A and B are two classes of indistinguishable agents represented by nodes in a non-branching level, and an agent in A receives a message tagged k from an agent in B , then so do all agents in A (or else their node would branch). Since no agent in B can send more than one message tagged k , we have $|A| \leq |B|$. Repeating this reasoning for all pairs of communicating classes, and recalling that the network is strongly connected, we conclude that all classes must have the same size.

A remarkable consequence is that, with output port awareness, a unique leader can be elected if and only if it is possible to assign unique identifiers to all agents (also known as the *Naming problem*). A terminating algorithm for both problems, if n is known, is easily obtained by combining the Leader Election algorithm in Section 3.1 (which was designed for undirected networks only) with the stabilizing technique for directed networks in Section 3.3.

3.8 Varying Inputs

Suppose that agents have inputs that may change at every step. Note that the history tree architecture already supports attaching inputs to nodes, and so the basic stabilizing algorithms previously discussed can become *streaming algorithms* that adaptively return the correct output with an amortized delay of $n - 1$ steps.

⁴ This observation, limited to static networks, is also found in [2, Theorems 3 and 10].




Fig. 10. A routine in the universal self-stabilizing protocol: delete level L_0 in the view, and then repeatedly merge nodes whose corresponding sub-views are isomorphic.

3.9 Self-Stabilization

A network algorithm is *self-stabilizing* if it returns the correct output regardless of the initial state of the agents. This implies tolerance to memory corruption, transient faults, and agents dynamically joining and leaving the network.

Let us assume the network to be synchronous, directed, and dynamic. We remark that most existing self-stabilizing protocols, such as Boldi–Vigna’s [5], are specifically designed for static networks, and cause agents to perpetually and incorrectly reset their states when executed in highly dynamic networks.

We will give a *universal* self-stabilizing protocol that constructs coherent views, allowing to convert any stabilizing algorithm into a self-stabilizing one (note that non-trivial terminating computations are impossible in this scenario).

The core idea is that an agent can deliberately “forget” old information by deleting level L_0 of its current view, connect level L_1 to the root, and merge equivalent nodes from top to bottom to restore a well-formed view, as in Figure 10.

If the number of agents n is known, they can simply update their views as usual (resetting their state if it does not encode a well-formed view) and delete old levels when their number exceeds some fixed threshold, e.g., $2n - 2$. Eventually, this protocol produces views that correctly describe the last $2n - 2$ communication steps, which are enough for the stabilizing algorithms of Sections 2.3 and 3.3.

On the other hand, if n is unknown, every agent updates its view as usual, but deletes level L_0 every *two* steps. This is controlled by a local binary flag that is toggled at every step. However, before merging its view with a neighbor’s view of different height, it deletes the top levels of the taller view to match the other. Also, if the taller view was its own, the agent does not toggle its flag for a step.

In $O(n)$ steps, all agents have views of equal height, as well as equal flags. Moreover, their views eventually describe an increasingly large number of previous communication steps, enabling any stabilizing algorithm to work correctly.

Open Problem 7. Is there a universal self-stabilizing protocol for *semi-synchronous* strongly connected directed dynamic networks?

3.10 Finite-State Stabilization

The stabilizing algorithms described so far assume that all agents constantly update their views at every step, which requires an unlimited amount of internal memory. To mitigate this, we will now give a *universal finite-state* protocol that enables the conversion of any stabilizing algorithm into one that uses a finite amount of memory (as a function of n), albeit with an extra delay.

Let us assume the network to be synchronous, undirected, connected, and dynamic. A level L_i in a view is said to be *suitable* if every node in L_{i-1} has a unique child in L_i and the red edges between these two levels are compatible with a connected network. Recall that the basic stabilizing algorithms of Section 2.3 do all computations using only the shallowest suitable level of a view.

Now, when two neighboring agents share views whose shallowest suitable levels are isomorphic, they do not merge their views. Essentially, this removes a link from the communication network for that step. Also, when an agent receives no views that it can merge, it skips updating its view altogether, i.e., it remains inactive for that step. Note that this makes the network semi-synchronous, so views must be *equalized* prior to being used by the protocol (see Section 3.5).

According to this protocol, if all agents remain inactive forever, they all share the same suitable level, which is enough to perform correct computations. Otherwise, some inactive agent acquires relevant information within $n - 1$ steps and reactivates. Moreover, since there are at most $n - 1$ branching nodes in the history tree, the agents can incorrectly guess a suitable level at most $n - 1$ times. Thus, the protocol is finite-state, but introduces an overhead of $O(n^2)$ steps.

Open Problem 8. Is there a universal finite-state stabilizing protocol for connected undirected dynamic networks with an overhead of $O(n)$ steps?

Another open problem is to design a protocol that is both self-stabilizing and finite-state. In Section 3.9 we already gave one, but it assumes n to be known.

Open Problem 9. Is there a universal finite-state self-stabilizing protocol for connected undirected dynamic networks of unknown size?

3.11 Memoryless Computation

An agent is *memoryless* if its state is reset at every communication step, meaning its entire memory is erased after it sends messages and before receiving messages from neighbors. The goal is to design a universal protocol that enables memoryless agents to construct coherent views of *some* history tree related to the network. This would allow basic algorithms to run correctly within the memoryless model.

Open Problem 10. Under what assumptions is there a universal memoryless protocol for anonymous networks?

3.12 Congested Networks

In a *congested network* of n agents, communication links have a logarithmic bandwidth, and therefore the size of each message is limited to $O(\log n)$ bits.⁵

It is evident that the basic technique of encoding a view as a single message is not applicable in a congested network, because the size of a view grows polynomially at every step (see Section 2.3). Furthermore, any naive approach that simply subdivides a view into smaller pieces to be transmitted over multiple steps is clearly ineffective in anonymous dynamic networks.

Correct protocols for constructing views in congested dynamic networks are found in [10]. If the dynamic diameter d is known, then $O(\log n)$ bits of information can be reliably *broadcast* in phases of d steps. This makes it possible for agents to share enough information to construct a history tree that is compatible with their network, one level at a time. Every node in this history tree has a unique label of logarithmic size, which also identifies all agents represented by that node. This protocol can be used to solve the Counting problem in $O(dn^2)$ steps.

If d is unknown but there is a unique leader, there is a more complex protocol that attempts to estimate d by implementing a *reset module* that repeatedly doubles the estimate every time a broadcasting error is detected. The leader coordinates the process, ensuring that all agents agree on the same version of the history tree. This yields a Counting algorithm that terminates in $O(n^3)$ steps.

Open Problem 11. Is there a terminating Counting algorithm for congested dynamic networks *with more than one leader*?

It was proved in [12] that solving the Counting problem in congested networks by broadcasting “tokens” requires at least $\Omega(n^2/\log n)$ steps. On the other hand, there is a solution in $O(n^2)$ steps if messages have size $O(n \log n)$ bits [10].

Open Problem 12. Can a Counting algorithm terminate in $O(n^2)$ steps in all congested dynamic networks with a unique leader?

We find it fitting to conclude this note with an open problem that is unlikely to have a solution using history trees.

Open Problem 13. Is there a terminating Counting algorithm for congested dynamic networks where agents have *logarithmic-sized memory*?

Acknowledgments. The author expresses profound gratitude to collaborators Jérémie Chalopin, Giuseppe A. Di Luna, and Federico Poloni for numerous inspiring discussions. This work was supported by the JSPS KAKENHI grant 23K10985, the University of Aizu Competitive Research Funding for FY 2023, and the University of Salerno.

⁵ Agents do not necessarily know this limit, since they generally do not know n . Nonetheless, if they attempt to send larger messages, the network’s behavior is undefined. Ideally, algorithms should be designed to automatically avoid this situation.

Disclosure of Interests. The author has no competing interests to declare that are relevant to the content of this article.

References

1. Angluin, D.: Local and Global Properties in Networks of Processors. In: Proceedings of the 12th ACM Symposium on Theory of Computing (STOC '80). pp. 82–93 (1980)
2. Boldi, P., Codenotti, B., Gemmel, P., Shammah, S., Simon, J., Vigna, S.: Symmetry Breaking in Anonymous Networks: Characterizations. In: Proceedings of the 4th Israeli Symposium on Theory of Computing and Systems (ISTCS '96). pp. 16–26 (1996)
3. Boldi, P., Vigna, S.: Computing Anonymously with Arbitrary Knowledge. In: Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing (PODC '99). pp. 181–188 (1999)
4. Boldi, P., Vigna, S.: Fibrations of Graphs. *Discrete Mathematics* **243**(1), 21–66 (2002)
5. Boldi, P., Vigna, S.: Universal Dynamic Synchronous Self-Stabilization. *Distributed Computing* **15**, 137–153 (2002)
6. Charron-Bost, B., Lambein-Monette, P.: Computing Outside the Box: Average Consensus over Dynamic Networks. In: Proceedings of the 1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND '22). pp. 10:1–10:16 (2022)
7. Charron-Bost, B., Moran, S.: The Firing Squad Problem Revisited. *Theoretical Computer Science* **793**, 100–112 (2019)
8. Conway, J.H.: *Regular Algebra and Finite Machines*. Chapman & Hall (1971)
9. Di Luna, G.A., Viglietta, G.: Computing in Anonymous Dynamic Networks Is Linear. In: Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science (FOCS '22). pp. 1122–1133 (2022)
10. Di Luna, G.A., Viglietta, G.: Brief Announcement: Efficient Computation in Congested Anonymous Dynamic Networks. In: Proceedings of the 42nd ACM Symposium on Principles of Distributed Computing (PODC '23). pp. 176–179 (2023)
11. Di Luna, G.A., Viglietta, G.: Optimal computation in leaderless and multi-leader disconnected anonymous dynamic networks. In: Proceedings of the 37th International Symposium on Distributed Computing (DISC '23). pp. 18:1–18:20 (2023)
12. Dutta, C., Pandurangan, G., Rajaraman, R., Sun, Z., Viola, E.: On the Complexity of Information Spreading in Dynamic Networks. In: Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '13). pp. 717–736 (2013)
13. Kowalski, D.R., Mosteiro, M.A.: Polynomial Counting in Anonymous Dynamic Networks with Applications to Anonymous Dynamic Algebraic Computations. *Journal of the ACM* **67**(2), 11:1–11:17 (2020)
14. Norris, N.: Universal Covers of Graphs: Isomorphism to Depth $n - 1$ Implies Isomorphism to All Depths. *Discrete Applied Mathematics* **56**(1), 61–74 (1995)
15. Tani, S.: Compression of View on Anonymous Networks—Folded View—. *IEEE Transactions on Parallel and Distributed Systems* **23**(2), 255–262 (2012)
16. Yamashita, M., Kameda, T.: Computing on an Anonymous Network. In: Proceedings of the 7th ACM Symposium on Principles of Distributed Computing (PODC '88). pp. 117–130 (1988)
17. Yamashita, M., Kameda, T.: Computing on Anonymous Networks. I. Characterizing the Solvable Cases. *IEEE Transactions on Parallel and Distributed Systems* **7**(1), 69–89 (1996)