

Efficient Computation in Congested Anonymous Dynamic Networks*

Giuseppe A. Di Luna¹ ✉

DIAG, Sapienza University of Rome, Italy

Giovanni Viglietta¹ ✉

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

Abstract

An *anonymous dynamic network* is a network of indistinguishable processes whose communication links may appear or disappear unpredictably over time. Previous research has shown that deterministically computing an arbitrary function of a multiset of input values given to these processes takes only a linear number of communication rounds (Di Luna–Viglietta, FOCS 2022).

However, fast algorithms for anonymous dynamic networks rely on the construction and transmission of large data structures called *history trees*, whose size is polynomial in the number of processes. This approach is unfeasible if the network is *congested*, and only messages of logarithmic size can be sent through its links. Observe that sending a large message piece by piece over several rounds is not in itself a solution, due to the anonymity of the processes combined with the dynamic nature of the network. Moreover, it is known that certain basic tasks such as all-to-all token dissemination (by means of single-token forwarding) require $\Omega(n^2/\log n)$ rounds in congested networks (Dutta et al., SODA 2013).

In this work, we develop a series of practical and efficient techniques that make it possible to use history trees in congested anonymous dynamic networks. Among other applications, we show how to compute arbitrary functions in such networks in $O(n^3)$ communication rounds, greatly improving upon previous state-of-the-art algorithms for congested networks.

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

Keywords and phrases anonymous dynamic network, congested network, history tree

Digital Object Identifier 10.4230/LIPIcs.MFCS.2024.49

Related Version *Full Version*: <https://arxiv.org/abs/2301.07849>

Funding Giuseppe A. Di Luna was Giovanni Viglietta was

Giuseppe A. Di Luna: partially supported by Sapienza Progetto Ateneo RM1221816C1760BF.

Giovanni Viglietta: partially supported by the JSPS KAKENHI grant 23K10985.

1 Introduction

Dynamic networks. In recent years, distributed computing has seen a remarkable increase in research on the algorithmic aspects of networks that constantly change their topology [7, 30, 32]. The study of these *dynamic networks* is motivated by technologies such as wireless sensors networks, software-defined networks, and networks of smart devices. Typically, the distributed system consists of n processes that communicate with each other in synchronous rounds. At each round, the network topology is rearranged arbitrarily, and communication links appear or disappear unpredictably.

* A preliminary version of this paper appeared as a Brief Announcement in the Proceedings of PODC 2023 [16].

¹ Both authors contributed equally to this research.



Anonymity and leadership. There are efficient algorithms for various tasks that work under the assumption that processes have *unique IDs* [6, 27, 28, 29, 32, 34]. However, unique IDs may not be available due to operational limitations [34] or to protect user privacy; for instance, assigning temporary random IDs to users of COVID-19 tracking apps was not sufficient to eliminate privacy concerns [40]. Systems where processes are indistinguishable are called *anonymous*.

It is known that many fundamental problems for anonymous networks cannot be solved without additional “symmetry-breaking” assumptions: A notable example is the *Counting problem*, i.e., determining the total number of processes n . The most typical symmetry-breaking choice is assuming the presence of a single distinguished process in the system, called *leader* [1, 2, 3, 4, 14, 19, 21, 23, 31, 39, 42]. A leader process may represent a base station in a sensor network, a super-node in a P2P network, etc.

Disconnected networks. Another common assumption is that the network is connected at every round [28, 34]. However, this assumption appears somewhat far-fetched when one considers the highly dynamic nature of some real-world networks, such as P2P networks of smart devices moving unpredictably. A weaker and more reasonable assumption is that the union of all the network’s links across any T consecutive rounds induces a connected (multi)graph on the processes [25, 36]. Such a network is said to be *T -union-connected*, and $T \geq 1$ is its *dynamic disconnectivity* [17].

Congested networks. Almost all previous research on anonymous dynamic networks pertains to models that impose no limit on the size of messages exchanged by processes [11, 15, 21, 22, 23, 24, 26, 34, 36]. Unfortunately, in most mobile networks, sending small-size messages is not only desirable but also a necessity; for example, in sensor networks, short communication times significantly increase battery life. A more realistic model assumes the network to be “congested” and limits the size of every message to $O(\log n)$ bits, where n is the number of processes [38].²

General computation. A recent innovation in the study of anonymous dynamic networks with leaders was the introduction of *history trees* in [15], which led to an optimal deterministic solution to the *Generalized Counting problem*³ in the non-congested network model. This problem is “complete” for a large class of functions called *multiset-based functions*, which in turn are the only computable functions in this model. The theory of history trees was extended in [17] to *leaderless* networks, providing optimal algorithms for the *Frequency problem*.⁴ This problem is complete for the class of *frequency-based* functions, which are the only computable functions in leaderless systems. Thus, the computational landscape for the non-congested network model is fully understood, and optimal linear-time algorithms are known for anonymous dynamic systems with and without leaders. No previous research exists on the congested network model, except for a recent preprint that gives a Counting algorithm in $\tilde{O}(n^{2T(1+\epsilon)+3})$ rounds for networks with leaders [25]. Note that its running time is exponential in the dynamic disconnectivity T and becomes $\tilde{O}(n^{5+\epsilon})$ for connected networks.

² This $O(\log n)$ limit on message sizes does not imply that the processes have a-priori information about n . The size limit is not explicitly given to the processes, and it is up to the algorithm to automatically prevent larger messages from being sent.

³ In the Generalized Counting problem, each process starts with a certain input, and the goal is to determine how many processes have each input. That is, each process has to compute the *multiset* of all inputs.

⁴ In the Frequency problem, the goal is to determine the percentage of processes that have each input.

1.1 Contributions and Techniques

Contributions. In this paper, we provide a state-of-the-art general algorithmic technique for T -union-connected anonymous dynamic *congested* networks, with and without leaders. The resulting algorithms run in $O(Tn^3)$ rounds, where n is the (initially unknown) total number of processes.

In Section 4 we give a basic and slightly inefficient Counting algorithm that applies to a limited setting but already contains all of the key ideas of our technique. In Section 5 we sketch its correctness (the technical details of the proof are found in the arXiv version linked in the title page), and in Section 6 we discuss optimizations and extensions of the basic algorithm to several other settings. This includes the computation of all multiset-based functions. Section 7 concludes the paper with some directions for future research.

Technical background. Informally, a *history tree* is a way of representing the history of a network in the form of an infinite tree. Each node in a history tree represents a set of anonymous processes that are “indistinguishable” at a certain round, where two processes become “distinguishable” as soon as they receive different sets of messages (see Section 3).

The theory of history trees developed in [15, 17] yields optimal general algorithms for anonymous dynamic networks with and without leaders, assuming the network is not congested. The idea is that processes can work together to incrementally construct the history tree by repeatedly exchanging and merging together their respective “views” of it. Once they have a sufficiently large portion of the history tree (i.e., a number of “levels” proportional to n), each process can locally analyze its structure and perform arbitrary computations on the multiset of input values originally assigned to the processes.

Challenges. Unfortunately, implementing the above idea requires sending messages containing entire “views” of the history tree. The size of a view is $\Theta(n^3 \log n)$ bits in the worst case, and is therefore unsuitable for the congested network model [15]. There is a major difficulty in dealing with this problem deterministically, which stems from the lack of unique IDs combined with the dynamic nature of the network.

It is worth noting that the “naive” approach of breaking down large messages into smaller pieces to be sent in different rounds does not work. Indeed, it is not clear how the original message can then be reconstructed, because the pieces carry no IDs and a process’ neighbors may change at every round. This may result in messages from different processes being mixed up and non-existent messages being reconstructed.

Methodology. Our main contribution is a general method that allows history trees to be transmitted reliably and deterministically between anonymous processes in a dynamic *congested* network with a leader. To overcome the fundamental issues outlined above, we devised a basic protocol combining different techniques, as well as a number of extensions, including leaderless ones. Although the techniques introduced in this paper are self-contained and do not rely on the results of [15], they effectively allow us to reduce the congested network model to the non-congested one, making it possible to apply the Counting algorithm in [15] as a “black box”.

Firstly, we developed a method for dynamically assigning temporary (non-unique) IDs to processes; this method is an essential part of the history tree transmission algorithm. In fact, the nodes of our history trees are now augmented with IDs, meaning that each node represents the set of processes with a certain ID. When processes with equal IDs get disambiguated, they get new IDs.

The transmission of history trees occurs level by level, one edge at a time. Since the total ordering between IDs induces a total ordering on the history tree's edges, the processes can collectively transmit sets of edges with a method reminiscent of *Token Dissemination* [28].

Essentially, all processes participate in a series of broadcasts; the goal of each broadcast is to transmit the next “highest-value” edge to the whole network. The problem is that no upper bound on the *dynamic diameter* of the network is known, and so there is no way of knowing how many rounds it may take for all processes to receive the edge being broadcast.

We adopt a self-stabilizing approach to ensure that all messages are successfully broadcast. We give a communication protocol based on acknowledgments by the leader, where failure to broadcast a message alerts at least one process. Alerted processes start broadcasting error messages, which eventually cause a reset of the broadcast that caused the error. A mechanism that dynamically estimates the diameter of the network guarantees that no more than $O(\log n)$ resets are performed.

Finally, in order to achieve a cubic running time, we do not construct the history tree of the actual network, but a more compact history tree corresponding to a *virtual network*. The virtual network is carefully derived from the real one in such a way as to amortize the number of edges in the resulting history tree and further reduce the final worst-case running time by a factor of n .

2 Previous Work

Non-congested networks. The study of computation in anonymous dynamic networks has been mainly devoted to two fundamental problems: The *Counting problem* in networks with a leader [11, 12, 13, 21, 22, 23, 24, 26] and the *Average Consensus problem* in leaderless networks [5, 8, 9, 10, 24, 33, 35, 37, 41, 43]. This research thread produced a series of algorithms for these problems; the underlying technique used is a local averaging or mass-distribution method coupled with refined termination strategies.

A radically different technique based on history trees was recently used to optimally compute arbitrary multiset-based functions in $3n$ rounds in networks with a leader [15]. This approach was successfully extended to multi-leader, leaderless, and disconnected networks [17].

Congested networks. As for the congested model, the only paper that has ever studied deterministic algorithms for anonymous dynamic networks, to the best of our knowledge, is the recent preprint [25], which solves the Counting problem in $\tilde{O}(n^{2T(1+\epsilon)+3})$ rounds. As usual, T is the dynamic disconnectivity of the network and ϵ is an arbitrarily small positive constant. By comparison, our main algorithm has a running time of $O(Tn^3)$ rounds; hence, its dependence on T is linear (as opposed to exponential) and, for connected networks (i.e., when $T = 1$), the improvement is a factor of $\Theta(n^{2+\epsilon} \log^k n)$.

Most previous research efforts on congested networks in the dynamic setting have focused on randomized algorithms or processes with unique IDs [18, 20, 28]. In this context, a problem similar to Counting is *Token Dissemination*, where each process starts with a token and the goal is for every process to collect all tokens. In connected networks, this problem is solved in $O(n^2)$ rounds by a simple *token-forwarding* algorithm (i.e., no manipulation is done on tokens other than storing, copying, and individually transmitting them) [28]. Interestingly, solving the Token Dissemination problem by token-forwarding algorithms requires at least $\Omega(n^2/\log n)$ rounds [18].

It is worth remarking that the randomized algorithm in [28] only solves the Counting problem approximately and assumes a-priori knowledge of an upper bound on n . Moreover, this algorithm only works with high probability. These are three key differences that make our

contribution preferable. Furthermore, assuming processes to have unique IDs (or randomly generating unique IDs) as in [28] defeats the purpose of safeguarding user privacy, which is a motivation of our work.

3 Definitions and Fundamentals

Computation model. A *dynamic network* is modeled by an infinite sequence $\mathcal{G} = (G_t)_{t \geq 1}$, where $G_t = (V, E_t)$ is an undirected multigraph whose vertex set $V = \{p_1, p_2, \dots, p_n\}$ is a system of n *anonymous processes* and E_t is a multiset of edges representing *links* between processes. Hence, there may be multiple links between two processes, or even from a process to itself.⁵

If there is a constant $T \geq 1$ such that, for every $t \geq 1$, the multigraph $G_t^* = (V, \bigcup_{i=t}^{t+T-1} E_i)$ is connected, the network is said to be *T-union-connected*, and the smallest such T is its *dynamic disconnectivity*.⁶

Each process p_i starts with an *input*, which is assigned to it at *round 0*. It also has an internal state, which is initially determined by its input. At each *round* $t \geq 1$, every process composes a message (as a function of its internal state) and sends it to its neighbors in G_t through all its incident links.⁷ In the congested network model, only messages of $O(\log n)$ bits can be sent.⁸ By the end of round t , each process reads all messages coming from its neighbors and updates its internal state according to a local algorithm \mathcal{A} . Note that \mathcal{A} is the same for all processes, and is a deterministic function of the internal state and the multiset of messages received in the current round.

The input of each process also includes a *leader flag*. In Section 4, we will assume that the leader flag of exactly one process is set (this process is the *unique leader*); in Section 6, we will discuss the *leaderless* case, where none of the processes has the leader flag set.

A process may return an *output* at the end of a round, which must be a function of its current internal state. A process may also *terminate* execution after returning an output. An algorithm \mathcal{A} solves the *Counting problem* if executing \mathcal{A} at every round eventually causes all processes to simultaneously output n and terminate. The (worst-case) *running time* of \mathcal{A} , as a function of n , is the maximum number of rounds it takes for \mathcal{A} to solve the problem, across all possible dynamic networks of size n and all possible input assignments.


History trees. *History trees* were introduced in [15] 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 the processes. 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 processes

⁵ Each self-loop in G_t represents a single link, hence a single message being sent and received by the same process.

⁶ A similar parameter for dynamic networks is the *dynamic diameter* D , defined as the maximum number of rounds it may take for a message to be broadcast from a process to all other processes. The parameters T and D are related by the inequalities $T \leq D \leq T(n-1)$, which are best possible [17].

⁷ Contrary to static networks, where processes may be allowed to send different messages to different neighbors, dynamic networks usually require processes to “broadcast” the same message through all incident links due to the lack of unique port numbers. This is the case, for instance, in wireless radio communications, where messages are sent to all processes within communication range, and the anonymity of the network prevents destinations from being specified.

⁸ We may want to assume that the total number of links in a network multigraph G_t (counted with their multiplicities) is bounded by a polynomial in n . In fact, when dealing with congested networks, we explicitly make this assumption, as it ensures that the multiplicity of a link can always fit in a single $O(\log n)$ -sized message.



■ **Figure 1** The first rounds of a dynamic network with $n = 9$ processes and the corresponding levels of the history tree. Level L_t consists of the nodes at distance $t + 1$ from the root r , which represent indistinguishable processes after the t th communication round. There are no leaders in the network, but each process has an input from the set $\{A, B, C\}$. Only the nodes in L_0 have explicit labels; all labels of the form a_i and b_i were added for the reader’s convenience, and indicate classes of indistinguishable processes (in contrast, the nodes of the *virtual history tree* introduced in Section 4 do have IDs). Note that the two processes in b_4 are still indistinguishable at the end of round 2, although they are linked to the distinguishable processes b_5 and b_6 . This is because such processes were in the same class a_5 at round 1. The subgraph induced by the nodes in the green blob is the *view* of the two processes in b_1 . We remark that a history tree does not contain any explicit information about how many processes each node represents.

that are *indistinguishable* at the end of round t (with the exception of L_{-1} , which contains a single node r representing all processes). The definition of distinguishability is inductive: At the end of round 0, two processes are distinguishable if and only if they have different inputs. At the end of round $t \geq 1$, two processes are distinguishable if and only if they were already distinguishable at round $t - 1$ or if they have received different multisets of messages at round t . (We refer to “multisets” of messages, as opposed to sets, because multiple copies of identical messages may be received; each message has a *multiplicity*.)

Each node in level L_0 has a label indicating the input of the processes 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 processes represented by the *parent node* v . The *red multi-edges* represent communications between processes. 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 process represented by v' receives m (identical) messages from processes represented by v .

The *view* of a process p at round $t \geq 0$ as the subgraph of the history tree which 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).

Applications of history trees. As proved in [15], the view of any process at round $3n$ contains enough information to determine the multiset of the initial inputs (i.e., how many processes have each input value). Thus, once a process is able to locally construct a view spanning $3n$ levels of a history tree, it can immediately do any computation on the inputs, and in particular determine n and solve the Generalized Counting problem.

History trees were adopted in [15, 17], where it is shown how processes can construct their views of the history tree in real time. For the algorithm to work, each process is required to repeatedly send its current view to all its neighbors at every round, merging it with all the views it receives from them.

This approach is not feasible in the congested network model, because a view at round t has size $\Theta(tn^2 \log n)$ in the worst case, since there may be $\Theta(n^2)$ red edges in each level. In the following, we will develop a strategy whereby processes can construct a history tree one red edge at a time. The core idea is that the nodes of the history tree are assigned unique IDs, and therefore a single red edge can be encoded in only $O(\log n)$ bits as a pair of IDs and a multiplicity.

4 Basic Counting Algorithm

In this section we describe our deterministic Counting algorithm for congested anonymous dynamic networks in its most basic version. The algorithm assumes the network to be connected ($T = 1$), to have a unique leader, and execution terminates in $O(n^3 \log n)$ rounds with the leader reporting the total number of processes in the system, n . The complete pseudocode is found in the arXiv version of this paper.

In Section 6, we will optimize this basic algorithm, making it terminate in $O(n^3)$ rounds. We will also extend the algorithm in several directions, for instance by making all processes (as opposed to the leader only) simultaneously output n and terminate. Furthermore, we will show how to not only count the number of processes, but also compute arbitrary (multiset-based, cf. [15]) functions, assuming that input values are assigned to the processes. Finally, we will extend the algorithm to leaderless networks and T -union-connected networks.

4.1 Algorithm Outline

The only input given to each process is whether or not it is the leader. Each process also has some private memory which is used to permanently store information in the form of internal variables.

Virtual history tree (VHT). The overall goal of the algorithm is for the processes to implicitly agree on the first $O(n)$ levels of a particular history tree, called *virtual history tree (VHT)*, which corresponds to a dynamic network \mathcal{N} of n processes. Once the construction of each new level of the VHT is complete (refer to Section 4.5), the leader locally runs the Counting algorithm from [15] on the VHT. If this algorithm successfully returns a number (as opposed to “Unknown”), the leader outputs it; otherwise, the construction of a new level of the VHT is initiated.

Virtual network (\mathcal{N}). The dynamic network $\mathcal{N} = (N_1, N_2, \dots)$ represented by the VHT is in fact a *virtual network*, in the sense that none of the multigraphs N_t necessarily coincides with any multigraph of links actually occurring in the real communication network $\mathcal{G} = (G_1, G_2, \dots)$. However, each N_t is obtained by carefully adding and removing links from some G_{i_t} (see Figure 2). This manipulation has the purpose of reducing the size of the resulting VHT by a factor of n (see Section 4.4).

Temporary IDs. To cope with the fact that processes are anonymous and information can only be sent in small chunks of size $O(\log n)$, each process has a *temporary ID* stored in a local variable called MyID. Each node v in the VHT also has an ID, indicating that v

represents all processes having that ID. Thus, a *red-edge triplet* of the form $(ID1, ID2, Mult)$ can be used to unambiguously represent a red edge of multiplicity $Mult$ between the nodes of the VHT whose IDs are $ID1$ and $ID2$. Since a red-edge triplet has size $O(\log n)$, it can be included in a single message. Note that the variable $MyID$ of each process may be modified over time as the VHT acquires more nodes.

Broadcast phases. The construction of the VHT is carried out level by level, and is done through several *broadcast phases*, which are indirectly coordinated by the leader (see Section 4.3). At first, each process knows the red edges incident to its corresponding node of the VHT. Then, ideally every two broadcast phases, the whole network learns a new red edge of the VHT. The broadcast phases continue until all processes know all red edges in the level (see Section 4.6).

Estimating the diameter. In order to guarantee the success of a broadcast phase, all processes must keep sending each other information for a certain number of rounds, which depends on the *dynamic diameter* of the network, and is $n - 1$ in the worst case [28]. Since the processes do not initially possess any information at all, they can only make estimations on the dynamic diameter. The current estimate is stored by each process in the variable $DiamEstimate$. Its value is initially 1, and it is doubled every time the processes detect an error in a broadcast.

Error phases. Detecting broadcasting errors and consistently reacting to them is by no means a trivial task, and is discussed in Section 4.7. Our broadcasting technique ensures that, if some red-edge triplet fails to be broadcast to the entire network and does not become part of the local VHT of all processes, at least one process becomes aware of this fact. Such a process enters an *error phase*, sending a high-priority message at every round containing the level number at which the error occurred. Error messages supersede the regular ones and eventually reach the leader.

Reset phases. When the leader finally receives an error message, it initiates a *reset phase*, whose goal is to force the whole network to restore a previous state of the VHT and continue from there. This is achieved by broadcasting a high-priority reset message. Since the error must have occurred because $DiamEstimate$ was too small, its value is doubled at the end of the reset phase.

Note that there is no obvious way for the leader to tell if any level of the VHT is actually missing some parts: At any time, there may be processes in an error phase unbeknownst to the leader. One of the challenges of our method is to ensure that the leader will not terminate with an incorrect guess on n due to the VHT being incomplete.

4.2 Communication and Priority

Counting rounds. The processes have to implicitly synchronize with one another to start and finish each broadcast phase at the same time. Part of the synchronization is achieved by the function $SendAndReceive$, which is called by each process at every communication round. This function simply sends a given message to all neighbors, collects all messages coming from the neighbors, and increments the internal variable $CurrentRound$. Since communications in the network are synchronous, all processes always agree on the value of $CurrentRound$.

Message types. The processes use messages of various types to share information with one another. Each message has a *label* describing its type, as well as at most three additional integer parameters. As it will turn out in the analysis of the algorithm, each parameter has size $O(\log n)$ bits.⁹ The message types and their parameters are as follows.

- **Null message.** Label: “Null”. No parameters.
- **Level-begin message.** Label: “Begin”. Parameters: ID.
- **Level-end message.** Label: “End”. No parameters.
- **Done message.** Label: “Done”. Parameters: ID.
- **Red-edge message.** Label: “Edge”. Parameters: ID1, ID2, Mult.
- **Error message.** Label: “Error”. Parameters: ErrorLevel.
- **Reset message.** Label: “Reset”. Parameters: ResetLevel, StartingRound, NewDiam.

Priority. Messages have *priorities* that determine how they are handled during a broadcast. The priority of a message is defined as follows:

Null < Begin < End < Done < Edge < ... < Reset $k + 1$ < Error k < Reset k < ... < Error 1 < Reset 1

That is, the message with lowest priority is the Null message, followed by all possible Level-begin messages, then the Level-end message, etc. The priority of a Level-begin message is independent of its parameter. For all other message types, however, the priority is also a function of the parameters.

As for Done messages, different ID parameters yield different priorities. Thus, priority induces a total ordering on the set of all possible Done messages; the precise ordering is irrelevant, as long as all processes implicitly agree on it. All Done messages have greater priority than Null, Level-begin and Level-end messages. The same goes for Red-edge messages: Different parameters yield different priorities, and all processes agree on the priority function. All Red-edge messages have greater priority than all Done messages, and lower priority than all Error and Reset messages.

An Error message (respectively, a Reset message) with a smaller **ErrorLevel** (respectively, **ResetLevel**) has a greater priority. Moreover, the priorities of Error and Reset messages are interleaved: The priority of an Error message with **ErrorLevel** = k is strictly between the priority of a Reset message with **ResetLevel** = $k + 1$ and the priority of a Reset message with **ResetLevel** = k .

4.3 Broadcasting Data

During the execution of the algorithm, a non-leader process may have a particular piece of information that it wishes to send to the leader. Similarly, the leader may have some information that it wishes to share with all processes in the network. Both operations are performed via a *broadcast* spanning several rounds.

The broadcast technique used to construct the VHT is implemented as follows. It is assumed that all processes participating in the broadcast are *synchronized*, i.e., they start at the same round and continue broadcasting for the same number of rounds (which is equal to **DiamEstimate**). Each process is also assumed to have the information it wishes to share packed in a message of the appropriate type (see Section 4.6), which is passed to the function **BroadcastPhase** as the argument **Message**.

⁹ As already remarked, the algorithm spontaneously creates $O(\log n)$ -sized messages without any a-priori knowledge on n .

At each broadcast round, each process sends its message to all its neighbors. Then it examines the messages received from the neighbors, as well as its own message, and keeps only the message with highest priority, discarding all others (function `BroadcastStep`). This is the message the process will send in the next round, and so on.¹⁰

Ideally, if the broadcast is continued for a sufficiently large number of rounds, all processes participating in the broadcast will eventually obtain the message having the highest priority among the ones initially owned by the processes (note that this may be an Error or Reset message, as well).

4.4 Defining the Virtual Network

Definition. Recall that the VHT is the history tree of the virtual network $\mathcal{N} = (N_1, N_2, N_3, \dots)$. We will now define N_t by induction on t . That is, assuming that the multigraphs N_1, N_2, \dots, N_{t-1} are already known, we will construct N_t based on the multigraph G_{i_t} , which represents the real communication network \mathcal{G} at a selected round i_t (Section 4.6 explains how i_t is selected).

Since the first $t - 1$ rounds of the virtual network are known, we can construct the first levels of the VHT up to level L_{t-1} . By definition of history tree, each node $v \in L_{t-1}$ represents a class P_v of processes that are indistinguishable after the first $t - 1$ “virtual rounds” modeled by the communication networks N_1, N_2, \dots, N_{t-1} .

Consider the simple undirected graph $H = (V, E)$, where $V = L_{t-1}$ and $\{u, v\} \in E$ if and only if $u \neq v$ and there is at least one link in G_{i_t} between a process in P_u and one in P_v . Recall that we are assuming that G_{i_t} is connected, and therefore H is connected, as well (disconnected networks will be discussed in Section 6). Let $S = (V, E')$ be an arbitrary spanning tree of H .

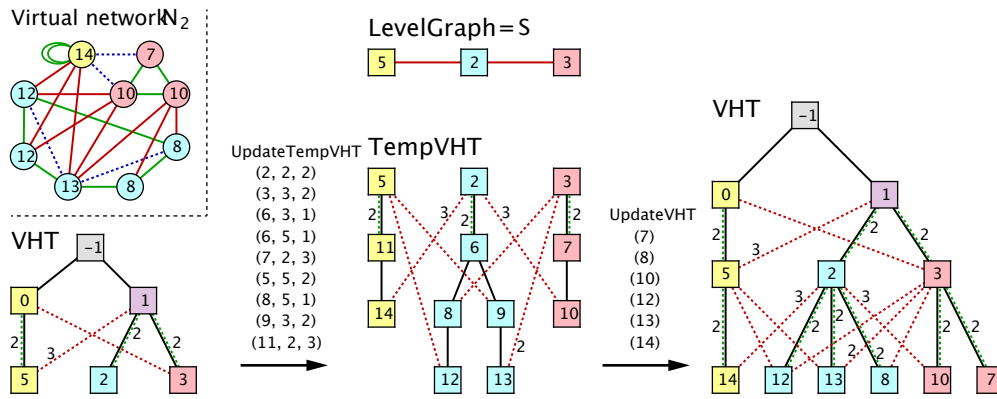
Now, N_t is the network having the following links (refer to Figure 2):

- For every edge $\{u, v\} \in E'$ of S , the multigraph N_t contains all the links in G_{i_t} having an endpoint in P_u and an endpoint in P_v (with the respective multiplicities).
- For every $v \in V$, the multigraph N_t contains a cycle C_v spanning all the processes in P_v . In the special case where P_v contains exactly two processes p_1 and p_2 , C_v is a double link between p_1 and p_2 . If P_v contains a single process p , C_v is a double self-loop on p .

Note that, in every case, C_v induces a 2-regular multigraph on P_v . The purpose of these (possibly degenerate) cycles C_v is to ensure that N_t is connected. On the other hand, the purpose of using a spanning tree S , as opposed to the full graph H , is to reduce the size of the VHT.

Implementation. In the actual algorithm, N_t is defined only implicitly in a distributed manner (because the ultimate goal is merely the construction of the VHT). The implicit construction of N_t starts at round i_t , where each process (not in an error phase) invokes the function `SetUpNewLevel` and sends a `Begin` message containing its own ID to each neighbor. As a result, each process learns the IDs of all its neighbors in G_{i_t} , as well as their multiplicities, and stores this information as a list of pairs of the form `(ID, Mult)` in the internal variable `ObsList`. It should be noted that a process discards all `Begin` messages from processes with the same ID as its own. These are replaced by the single pair `(MyID, 2)`, which accounts for the two edges of the cycle C_v incident to the process in N_t .

¹⁰Note that our broadcasting strategy implements a token-forwarding mechanism, and the sequence of all broadcast phases is akin to a Token Dissemination algorithm. It is known that any such algorithm must have a worst-case running time of at least $\Omega(n^2/\log n)$ rounds [18].



■ **Figure 2** Construction of the virtual network N_2 and level L_2 of the virtual history tree (VHT). The real network G_{i_2} consists of the $n = 9$ processes in the upper-left picture, connected by the red and blue links (not the green ones). Before the construction starts, the IDs in the network are 5, 2, and 3. Same-colored processes have equal IDs when the construction of L_2 starts; the labels indicate their IDs when the construction finishes. Accordingly, level L_1 of the VHT has three nodes with IDs 5, 2, 3. The graph H is a triangle on these three nodes, while its spanning tree S is the same as **LevelGraph**. Therefore, to construct N_2 we remove the blue edges from G_{i_2} and keep the red ones (see Section 4.4). For example, after the triplets $(6, 3, 1)$ and $(6, 5, 1)$ are accepted, any elements of **ObsList** corresponding to the two blue edges incident to the yellow process are deleted and are never broadcast. We also have to add the green edges, which represent the cycles C_v . Note that failing to add them would result in a disconnected network, because the process with ID 7 would be isolated. For clarity, the edges of the (temporary) VHT are colored red or green to match the edges of the virtual network that they represent (although technically they are all red edges).

The choice of links to be included in N_t (which directly reflects on the red edges included in the VHT) is guided by the construction and maintenance of an *auxiliary graph* stored in each process' internal variable **LevelGraph**. The auxiliary graph is a graph on $V = L_{t-1}$ and starts with no edges at all; it gradually acquires more edges until it becomes the spanning tree S (as defined above). This is carried out as part of the function **UpdateTempVHT**, which is invoked every time a red-edge triplet $(ID_1, ID_2, Mult)$ is selected to become a new red edge of the VHT at the end of a broadcast phase. Among other operations (described in Section 4.5), this function adds an edge to the auxiliary graph which connects the two nodes corresponding to ID_1 and ID_2 . Then the function **PreventCyclesInLevelGraph** is called, which deletes all the pairs in **ObsList** whose selection would cause the creation of a cycle in the auxiliary graph. This guarantees that, eventually, **LevelGraph** will be a tree (representing S).

4.5 Constructing the Virtual History Tree (VHT)

Initialization. The VHT is initialized by the function **InitializeVariables**. At first, the VHT only contains level L_{-1} (a single root node whose ID is -1) and level L_0 (two nodes with IDs 0 and 1, representing the leader and the non-leader processes, respectively). Accordingly, the leader initializes its own **MyID** variable to 0 and the non-leaders to 1.

Temporary VHT. Then, for all $t \geq 1$, the level L_t of the VHT is constructed based on the virtual network N_t as in Figure 2. When the construction begins, all processes (that are not in an error phase) acquire new pairs of the form $(ID, Mult)$ at the same round i_t and store them in **ObsList**, as explained in Section 4.4. The construction of a new level is not carried

49:12 Efficient Computation in Congested Anonymous Dynamic Networks

out directly on the VHT, but in a separate *temporary VHT* stored in the internal variable `TempVHT`. Initially, this is just a copy of level L_{t-1} of the VHT, i.e., a set of nodes, each with a distinct ID, and no edges.

Adding red edges. In order to determine how `TempVHT` should be updated, several broadcast phases are performed, allowing the processes to share their red-edge triplets with one another. Note that a process can transform each element of `ObsList` of the form $(ID, Mult)$ into a red-edge triplet by simply adding its own ID as the first element. Every time a red-edge triplet $(ID1, ID2, Mult)$ has been broadcast by a process to the entire network and has been “accepted” (see Section 4.6 for details), the function `UpdateTempVHT` is called. The result is that the node v of `TempVHT` whose ID is $ID1$ gets a new child v' with a new unique ID. Pictorially, a black edge is created connecting v' to v . Also, a red edge with multiplicity $Mult$ is added to `TempVHT`, connecting the new node with the node whose ID is $ID2$.

Updating IDs. By definition, the red-edge triplet $(ID1, ID2, Mult)$ indicates that *some* processes whose ID is $ID1$ have received exactly $Mult$ messages from processes whose ID is $ID2$. These processes, which were previously represented by the node v , are now represented by v' . Therefore, every process whose ID is $ID1$ that has the pair $(ID2, Mult)$ in its local `ObsList` removes it from the list and modifies its own ID from $ID1$ to the ID of v' .¹¹

Updating the VHT. When a process has no more red-edge triplets to share (that is, its `ObsList` is empty), it broadcasts a `Done` message containing its current ID. Recall that `Done` messages have lower priority than `Red-edge` messages (see Section 4.2). When the final result of a broadcast is a `Done` message containing a certain ID, all processes assume that *some* processes with that ID have sent all their red-edge triplets. Therefore, the node of the temporary VHT with that ID is ready to be added to the VHT.

To this end, the function `UpdateVHT` is called with the ID contained in the `Done` message. This function creates a new node v in level L_t of VHT corresponding to the node v' of `TempVHT` whose ID is the one passed to the function (refer to Figure 2). The node v gets the same ID as v' and becomes a child of the node $u \in L_{t-1}$ having the same ID as the root v'' of the tree containing v' . Then, v takes all the red edges found along the path from v' to v'' .

Finalizing the level. When the `ObsList` of a process is empty and the VHT already contains a node with its ID, the process broadcasts an `Level-end` message, which has lower priority than `Done` and `Red-edge` messages. When the result of a broadcast is an `Level-end` message, the construction of the level is finished.

4.6 Main Loop

The entry point of the algorithm is the function `Main`. After initializing the internal variables as already explained, the function goes through a loop that constructs the VHT level by level. At several points in this loop there may be errors that cause some of the levels to be undone and execution to resume from the beginning of the loop; we will discuss errors in Section 4.7. Next, we will describe an ideal error-free execution.

¹¹To clarify, the nodes of the VHT have unique IDs. Each node represents a class of “indistinguishable” processes, all of which have the same ID as the node itself. When processes disambiguate in the virtual network, they obtain different IDs. However, it is not necessarily true that all processes will have distinct IDs eventually. For example, if the network is the (static) complete graph, all non-leader processes will always have the same ID (which is incremented at every virtual round).

Level initialization. With each iteration of the main loop, a new level of the VHT is constructed. At the beginning, the function `SetUpNewLevel` is executed by all processes: This marks a selected round G_{i_t} as defined in Section 4.4 (the non-trivial fact that all processes always call this function at the same round will become apparent from our analysis of the algorithm, in Section 5).

VHT broadcast. Then an inner loop is entered; the purpose of each iteration is for all processes to learn new information, which causes an update of the temporary VHT or the VHT itself. At first, each process calls the function `MakeVHTMessage`, which picks an element from `ObsList`, converts it into a red-edge triplet by adding its own ID to it, and wraps it in a Red-edge message. If `ObsList` is empty, a Done message containing the process' ID is generated instead. The resulting message is used in a first broadcast phase which, after a number of rounds equal to `DiamEstimate`, returns the highest-priority message circulating in the network. This message is stored in the variable `VHTMessage`.

Acknowledgment broadcast. Note that, in the presence of a faulty broadcast, different processes may end up having different versions of `VHTMessage`. To ensure that all processes (that are not in an error phase) update their local copies of the (temporary) VHT in a consistent way, an “acknowledgment” broadcast phase is performed. Its purpose is for the leader to inform all other processes of the *accepted message*; by definition, the accepted message is the leader's version of `VHTMessage`.

In the acknowledgment phase, all non-leader processes broadcast a low-priority Null message, while the leader broadcasts the accepted message. The message resulting from the broadcast is then stored in the variable `AckMessage`.

Updating the level. Now, each process compares the contents of the two messages `VHTMessage` and `AckMessage`. If they are the same (hence not Null), then this is indeed the accepted message coming from the leader, and the data therein is used to update the temporary VHT or the VHT. Specifically, if `AckMessage` is a Red-edge message, the red-edge triplet therein is used to update the temporary VHT. If it is a Done message, the VHT is updated instead.

Finalizing the level. When a process is done broadcasting red-edge triplets and the VHT already contains a node representing it, the process broadcasts a low-priority Level-end message. As soon as the accepted message is the Level-end message, the level is considered complete and the inner loop is exited.

Counting processes. Now the leader extracts its own view from the VHT and locally runs the Counting algorithm from [15] on it (function `CountFromView`). By “extracts its own view” we mean that it makes a copy of the VHT and deletes all nodes that are not on a shortest path from the root to the deepest leader's node.

If the value returned by `CountFromView` is a number, this is taken as the correct number n of processes in the virtual network \mathcal{N} (and hence in the real network \mathcal{G}) and becomes the leader's output. Otherwise, a new iteration of the main loop starts, a new level of the VHT is constructed, and so on.

4.7 Handling Errors and Resets

Detecting errors. As mentioned, a broadcast phase lasts a number of rounds indicated by the variable `DiamEstimate`. If this number happens to be smaller than the dynamic diameter of the network, the broadcast may be unsuccessful, in the sense that not all processes may end up agreeing on the same highest-priority message. Fortunately, the protocol described in Section 4.6 makes error detection very simple: At the end of every acknowledgment broadcast, if the contents of a process' variables `VHTMessage` and `AckMessage` are different, the broadcast was unsuccessful. Any process that detects this event enters an error phase.

Error phase. When a process enters an error phase, it runs the function `BroadcastError`; as it turns out, only non-leader processes ever execute this function. During an error phase, a process continually broadcasts an Error message containing the index of the level of the VHT where the error was detected. This message is replaced by any other message of higher priority received by the process during the error phase (see Section 4.2). As soon as the process receives a Reset message for a level of smaller or equal depth than the one indicated by the current Error message (i.e., a Reset message of higher priority), it ends the error phase and enters a reset phase.

Another situation where a non-leader process enters an error phase is when it receives an Error message from some other process. In this case, it calls the function `HandleError` and starts broadcasting a new Error message containing the smaller between the index of the current level and the one contained in the received Error message. This causes higher-priority error messages to propagate through the network and eventually reach the leader.

Reset phase. This phase is initiated by the leader at the end of a broadcast phase (or after sending a Begin message) in case it received an Error message. The leader first waits for $2 \cdot \text{DiamEstimate} + 1$ rounds, sending only Null messages. This is to ensure that all non-leader processes finish any broadcast phase and enter an error phase. In turn, this prevents any possible conflicts between different broadcast phases before and after a reset.

Then the leader creates a Reset message containing the index of the VHT level where the error occurred, the current round number, and the new estimate on the dynamic diameter, i.e., twice the current one, since the error occurred because the estimate was too small.

Now the leader calls the function `BroadcastReset`, where it broadcasts the Reset message; any process that receives this message starts broadcasting it as well (provided that it has higher priority than the Error message it is currently broadcasting, as usual). The broadcast continues for a total number of rounds equal to the new dynamic diameter estimate. All processes that receive the Reset message are able to synchronize with one another and finish the reset phase at the same round, thanks to the information contained in the message itself.

At the end of the reset phase, the reset is actually performed: Every process that got the Reset message deletes the most recent levels of the VHT up to the level where the error occurred (this information is contained in the Reset message) and reverts its ID to the one it had at the beginning of the construction of that level. The variable `DiamEstimate` is also updated with the new estimate. At this point, the network is ready to resume construction of the VHT.

5 Proof of Correctness

We will now sketch a proof of correctness for the Counting algorithm in Section 4. The interested reader may refer to the arXiv version for a detailed and technically rigorous proof.

We say that two processes *agree* on a variable at a certain round if their local instances of that variable are equal. As it turns out, all processes that are not in an error phase must always agree on the variable `DiamEstimate`: this can be proved by induction on the number of reset phases that the leader has completed.

In turn, this implies that all processes (that are not in an error phase) are able to implicitly synchronize their broadcast phases with the leader, since these phases have a duration of `DiamEstimate` rounds. As a consequence, all processes (not in an error phase) also agree on the variables `CurrentLevel`, `VHT`, `TempVHT`, `LevelGraph`, and `NextFreshID`, because these variables are updated either at the beginning or at the end of a broadcast or reset phase.

Thus, we can unambiguously refer to phases and variables with no explicit mention of any individual process, because these phases and variables are the same throughout the whole network (excluding the processes in an error phase).

The goal of the algorithm is to construct the first levels of the VHT, i.e., the history tree of the virtual network \mathcal{N} . However, information on the virtual network is shared by all processes through several broadcast phases, and there is no guarantee that every broadcast phase will have the correct outcome (some data may fail to be communicated if `DiamEstimate` is smaller than the network's actual dynamic diameter). Thus, we distinguish between the *ideal VHT*, i.e., the “correct” history tree corresponding to \mathcal{N} , and the *effective VHT*, i.e., the one that is actually being constructed by the processes, which may be missing some information.

It is a byproduct of our protocol that, whenever a node is added to the effective VHT, then all red edges incident to it (and to nodes in the previous level of the ideal VHT) are also included. In particular, extracting the leader's view from the effective VHT yields a graph isomorphic to the view of the leader in the ideal VHT. This implies that the function `CountFromView` is always given as input a view of the history tree of a connected network of n processes. As proved in [15], this function must return either “Unknown” or n ; in particular, it returns n when the input view spans at least $3n$ levels.

Hence, our algorithm cannot return an incorrect number. It remains to prove that it does terminate within the desired number of rounds. Every time a faulty broadcast is detected and the leader is notified of it, a reset phase is initiated and `DiamEstimate` is doubled. It follows that no more than $O(\log n)$ reset phases ever occur, and `DiamEstimate` is always $O(n)$; this is also the duration of each broadcast phase.

Recall that links are not included in the virtual network if they cause cycles to appear in `LevelGraph` (see Section 4.4). Although this strategy may not guarantee that every level of the VHT has $O(n)$ red edges, it does amortize the total number of red edges over several levels. In fact, the first $O(n)$ levels of the ideal VHT only contain $O(n^2)$ red edges (as opposed to $O(n^3)$, which may be the case for the history tree of a generic network).

Summarizing, it takes $O(n^2)$ broadcast phases to construct the first $3n$ levels of the VHT, resulting in a total of $O(n^3)$ rounds. The construction process is partially undone and repeated at most $O(\log n)$ times due to resets, and so the final running time is $O(n^3 \log n)$.

It is straightforward to prove that the algorithm works in the congested network model, because all the data contained in a message are binary representations of values that are polynomial in n , e.g., round numbers, estimates of the dynamic diameter, or temporary IDs of processes. Since each message contains at most three of these values, its size is $O(\log n)$.

6 Extensions and Improvements

Simultaneous termination. In Section 4 we showed how the leader can determine n . If we want all processes, not just the leader, to simultaneously output n and terminate, we can use the following protocol. As soon as the leader knows n , it broadcasts a message of

49:16 Efficient Computation in Congested Anonymous Dynamic Networks

maximum priority containing n and the current round number c . Any process that receives this message keeps forwarding it until round $c + n$. By this round they all have it, and thus can output n and terminate simultaneously.

General computation. If each process is assigned an $O(\log n)$ -sized input (other than the leader flag), we can also perform general computations on the multiset of inputs. We only have to adapt the algorithm of Section 4 to work with history trees whose level L_0 contains one node per input occurring in the system, and then run the algorithm in [15] on this extended history tree. This can easily be done with the techniques developed in Section 4: We can construct L_0 like any other level, except that we use broadcast phases to transmit inputs instead of red edges.

Optimized running time. We can slightly speed up our algorithm at the cost of some additional bookkeeping. The idea is to refine the error and reset mechanism by resuming the construction of the VHT not from the *level* that caused an error, but from the *broadcast phase* that caused it.

The processes have to remember the order in which the leader accepted Red-edge and Done messages, as well as the Begin messages received at every begin round. When a process detects an error, it broadcasts an Error message containing not the current level number, but the number of messages that the leader has accepted up to that time. The advantage is that the reset phase can rewind the (temporary) VHT exactly to the desired point without erasing entire levels. Furthermore, by looking up the Begin messages received in the appropriate begin round, each process is also able to reconstruct its local `ObsList` at the desired time.

If a broadcast phase causes an error, the leader receives a notification after less than n rounds, and so it only has to undo the work done in $O(n)$ rounds. Since there can be at most $O(\log n)$ resets, the total time spent in reset phases or doing work that later gets undone is $O(n \log n)$ rounds. Thus, the new algorithm runs in $O(n^3) + O(n \log n) = O(n^3)$ rounds.

Leaderless computation. It is known that doing any non-trivial computation in a leaderless network requires some extra assumptions, for example the knowledge of an upper bound D on its dynamic diameter [17]. However, knowing D allows processes to reliably broadcast messages in phases of D rounds. This immediately yields an extension of our algorithm to leaderless networks in $O(Dn^2)$ rounds, where no acknowledgment phases or error and reset phases are needed.

Disconnected networks. We can extend our algorithm to T -union-connected networks, assuming the parameter T is known. As discussed in [17], the idea is to divide the sequence of rounds into blocks of size T . Within every block, each process keeps sending the same message and stores all incoming messages. At the end of a block, each process runs the algorithm from Section 4 (or its optimization) pretending that all the stored messages arrived in a single round. This is equivalent to running the algorithm on the dynamic network $\mathcal{G}^* = (G_1^*, G_{T+1}^*, G_{2T+1}^*, \dots)$, which is always connected (G_t^* is defined in Section 3). The running time is simply $O(Tn^3)$ rounds.

7 Concluding Remarks

In this paper we have extended the theory of history trees by introducing the tools necessary for the distributed construction and transmission of history trees in the congested network model. This resulted in a new state of the art for general computation in disconnected anonymous dynamic congested networks, with or without leaders.

Our history tree construction technique leads to general algorithms whose running time is cubic in the size of the network. An immediate open problem is whether this running time can be reduced. Since our algorithm broadcasts information using a token-forwarding approach, and by virtue of the $\Omega(n^2/\log n)$ lower bound of [18], we believe that it would be unlikely to achieve a better running time without a radical change in the technique used.

Understanding whether the Counting problem has as a super-linear lower bound in congested networks is of special importance, because it would mark a computational difference between anonymous dynamic networks in the congested and non-congested models.

It would be interesting to do a thorough fine-grained tradeoff analysis of our algorithm. For instance, it is not difficult to show that, if messages have size $O(n \log n)$, the running time of our algorithm can be reduced to $O(n^2)$.

References

- 1 D. Angluin, J. Aspnes, and D. Eisenstat. Fast Computation by Population Protocols with a Leader. *Distributed Computing*, 21(3):61–75, 2008.
- 2 J. Aspnes, J. Beauquier, J. Burman, and D. Sohier. Time and Space Optimal Counting in Population Protocols. In *Proceedings of the 20th International Conference on Principles of Distributed Systems (OPODIS '16)*, pages 13:1–13:17, 2016.
- 3 J. Beauquier, J. Burman, S. Clavière, and D. Sohier. Space-Optimal Counting in Population Protocols. In *Proceedings of the 29th International Symposium on Distributed Computing (DISC '15)*, pages 631–646, 2015.
- 4 J. Beauquier, J. Burman, and S. Kutten. A Self-stabilizing Transformer for Population Protocols with Covering. *Theoretical Computer Science*, 412(33):4247–4259, 2011.
- 5 D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., USA, 1989.
- 6 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.
- 7 A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-Varying Graphs and Dynamic Networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012.
- 8 B. Charron-Bost and P. Lambein-Monette. Randomization and Quantization for Average Consensus. In *Proceedings of the 57th IEEE Conference on Decision and Control (CDC '18)*, pages 3716–3721, 2018.
- 9 B. Charron-Bost and P. Lambein-Monette. Computing Outside the Box: Average Consensus over Dynamic Networks. In *Proceedings of the 1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND '22)*, pages 10:1–10:16, 2022.
- 10 B. Chazelle. The Total s-Energy of a Multiagent System. *SIAM Journal on Control and Optimization*, 49(4):1680–1706, 2011.
- 11 G. A. Di Luna and G. Baldoni. Brief Announcement: Investigating the Cost of Anonymity on Dynamic Networks. In *Proceedings of the 34th ACM Symposium on Principles of Distributed Computing (PODC '15)*, pages 339–341, 2015.
- 12 G. A. Di Luna, R. Baldoni, S. Bonomi, and I. Chatzigiannakis. Counting in Anonymous Dynamic Networks Under Worst-Case Adversary. In *Proceedings of the 34th IEEE International Conference on Distributed Computing Systems (ICDCS '14)*, pages 338–347, 2014.
- 13 G. A. Di Luna, S. Bonomi, I. Chatzigiannakis, and R. Baldoni. Counting in Anonymous Dynamic Networks: An Experimental Perspective. In *Proceedings of the 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics (ALGOSENSORS '13)*, pages 139–154, 2013.

- 14 G. A. Di Luna, P. Flocchini, T. Izumi, T. Izumi, N. Santoro, and G. Viglietta. Population Protocols with Faulty Interactions: The Impact of a Leader. *Theoretical Computer Science*, 754:35–49, 2019.
- 15 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.
- 16 G. A. Di Luna and G. Viglietta. Brief Announcement: Efficient Computation in Congested Anonymous Dynamic Networks. In *Proceedings of the 42nd ACM Symposium on Principles of Distributed Computing (PODC '23)*, pages 176–179, 2023.
- 17 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.
- 18 C. Dutta, G. Pandurangan, R. Rajaraman, Z. Sun, and E. Viola. On the Complexity of Information Spreading in Dynamic Networks. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '13)*, pages 717–736, 2013.
- 19 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.
- 20 B. Haeupler and F. Kuhn. Lower Bounds on Information Dissemination in Dynamic Networks. In *Proceedings of the 26th International Symposium on Distributed Computing (DISC '12)*, pages 166–180, 2012.
- 21 D. R. Kowalski and M. A. Mosteiro. Polynomial Counting in Anonymous Dynamic Networks with Applications to Anonymous Dynamic Algebraic Computations. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP '18)*, pages 156:1–156:14, 2018.
- 22 D. R. Kowalski and M. A. Mosteiro. Polynomial Anonymous Dynamic Distributed Computing Without a Unique Leader. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP '19)*, pages 147:1–147:15, 2019.
- 23 D. R. Kowalski and M. A. Mosteiro. Polynomial Counting in Anonymous Dynamic Networks with Applications to Anonymous Dynamic Algebraic Computations. *Journal of the ACM*, 67(2):11:1–11:17, 2020.
- 24 D. R. Kowalski and M. A. Mosteiro. Supervised Average Consensus in Anonymous Dynamic Networks. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '21)*, pages 307–317, 2021.
- 25 D. R. Kowalski and M. A. Mosteiro. Efficient Distributed Computations in Anonymous Dynamic Congested Systems with Opportunistic Connectivity. *arXiv:2202.07167 [cs.DC]*, pages 1–28, 2022.
- 26 D. R. Kowalski and M. A. Mosteiro. Polynomial Anonymous Dynamic Distributed Computing Without a Unique Leader. *Journal of Computer and System Sciences*, 123:37–63, 2022.
- 27 F. Kuhn, T. Locher, and R. Oshman. Gradient Clock Synchronization in Dynamic Networks. *Theory of Computing Systems*, 49(4):781–816, 2011.
- 28 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.
- 29 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.
- 30 F. Kuhn and R. Oshman. Dynamic Networks: Models and Algorithms. *SIGACT News*, 42(1):82–96, 2011.
- 31 O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Naming and Counting in Anonymous Unknown Dynamic Networks. In *Proceedings of the 15th International Symposium on Stabilizing, Safety, and Security of Distributed Systems (SSS '13)*, pages 281–295, 2013.

- 32 O. Michail and P. G. Spirakis. Elements of the Theory of Dynamic Networks. *Communications of the ACM*, 61(2):72, 2018.
- 33 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.
- 34 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.
- 35 A. Olshevsky. Linear Time Average Consensus and Distributed Optimization on Fixed Graphs. *SIAM Journal on Control and Optimization*, 55(6):3990–4014, 2017.
- 36 A. Olshevsky and J. N. Tsitsiklis. Convergence Speed in Distributed Consensus and Averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.
- 37 A. Olshevsky and J. N. Tsitsiklis. A Lower Bound for Distributed Averaging Algorithms on the Line Graph. *IEEE Transactions on Automatic Control*, 56(11):2694–2698, 2011.
- 38 D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, USA, 2000.
- 39 N. Sakamoto. Comparison of Initial Conditions for Distributed Algorithms on Anonymous Networks. In *Proceedings of the 18th ACM Symposium on Principles of Distributed Computing (PODC ’99)*, pages 173–179, 1999.
- 40 T. Sharma and M. Bashir. Use of Apps in the COVID-19 Response and the Loss of Privacy Protection. *Nature Medicine*, 26(8):1165–1167, 2020.
- 41 J. N. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 1984.
- 42 M. Yamashita and T. Kameda. Computing on Anonymous Networks. I. Characterizing the Solvable Cases. *IEEE Transactions on Parallel and Distributed Systems*, 7(1):69–89, 1996.
- 43 Y. Yuan, G.-B. Stan, L. Shi, M. Barahona, and J. Goncalves. Decentralised Minimum-Time Consensus. *Automatica*, 49(5):1227–1235, 2013.