



HAL
open science

Know your audience: Communication model and computability in anonymous networks

Bernadette Charron-Bost, Patrick Lambein-Monette

► To cite this version:

Bernadette Charron-Bost, Patrick Lambein-Monette. Know your audience: Communication model and computability in anonymous networks. 2023. hal-04334359

HAL Id: hal-04334359

<https://hal.science/hal-04334359>

Preprint submitted on 10 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Know your audience

Communication model and computability in anonymous networks

BERNADETTE CHARRON-BOST, CNRS, DI ENS École Normale Supérieure, France

PATRICK LAMBEIN-MONETTE, Unaffiliated, France

Distributed function computation is the problem, for a networked system of n autonomous agents, to collectively compute the value $f(\omega_1, \dots, \omega_n)$ of some input values $\omega_1, \dots, \omega_n$, each initially private to one agent in the network. As is commonplace in the field of distributed computing, the question of *which functions f are computable* is exquisitely sensitive to minute model assumptions. Here, we study and organize results pertaining to distributed function computation in anonymous¹ networks, both for the static and the dynamic case, under a communication model of directed and synchronous message exchanges, but with varying assumptions in the degree of awareness or control that a single agent has over its *outneighbors*, i.e., the set of agents hearing from it in a given round.

Our main argument is three-fold. First, in the “blind broadcast” model, where in each round an agent merely casts out a unique message without any knowledge or control over its addressees, the computable functions are those that only depend on the *set* $\{\omega_1, \dots, \omega_n\}$ of the input values, but not on their multiplicities or relative frequencies in the input. Second, in contrast, when we assume either that **a)** in each round, the agents know how many outneighbors they have; **b)** all communications links in the network are bidirectional; or **c)** the agents may address each of their outneighbors individually, then the set of computable functions grows to contain all functions that depend on the *relative frequencies* of each value in the input – such as the *average* $\frac{\omega_1 + \dots + \omega_n}{n}$ – but **not** on their *multiplicities* – thus, not the *sum* $\omega_1 + \dots + \omega_n$. Third, however, if one or several agents are distinguished as *leaders*, or if the cardinality of the network is known, then under any of the above three assumptions it becomes possible to recover the complete multiset $[\omega_1, \dots, \omega_n]$ and thus compute any function of the distributed input as long as it is invariant under permutation of its arguments. In the case of dynamic networks, we also discuss the impact of multiple connectivity assumptions.

1 INTRODUCTION

The aim of many multi-agent systems is to get all agents to compute a common value, which is a function of the values observed or sensed by each agent. Some typical examples are minimum-finding, the computation of the average and of the sum of the agent values. The computation of such functions naturally arises in a wide range of practical situations, including sensor networks, distributed optimization, or distributed control in autonomous systems. Towards this purpose, the agents repeatedly alternate between internal computations and communicating with each other. Here, our objective is to understand the fundamental limitations and capabilities for function computation that are inherent to the communication model assumed when considering a multi-agent network system.

Our abstract model captures common requirements for a variety of different settings, including the case of wireless sensor networks. We consider a networked system with a fixed set of agents and communication links that may vary over time. Our basic connectivity assumption is of a *finite diameter*, i.e., any pair of agents can communicate, possibly indirectly, over a period of time that is *uniformly bounded* throughout the execution. We model agents as automata

¹When agents are given unique identifiers, the question of computability is rather shallow, as essentially all functions are computable given sufficient connectivity.

Authors’ addresses: [Bernadette Charron-Bost](#), CNRS, DI ENS and École Normale Supérieure, 75005, Paris, France, charron@di.ens.fr; [Patrick Lambein-Monette](#), Unaffiliated, Paris, France, patrick@lambein.name.



This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

interacting over reliable communication channels. These automata are *deterministic* – in particular, randomization is *not* allowed – and identical – i.e., each agent runs the same local algorithm.

The network is anonymous: agents do not possess unique identifiers, and nothing initially distinguishes any two agents apart from, possibly, their input value. In addition, agents have limited or no knowledge of the network. In particular, they are not assumed to know the size of the network; at best, they have an upper bound over it. They are also unaware of the structure of the network or its diameter. Concerning the memory of each agent, it cannot be bounded and must grow with the size of the network, given that the class of functions under consideration includes the sum, the average... However, for our positive results, we will be looking for finite-state solutions: if possible, an agent should only use bounded memory in any execution. We will also be looking for *self-stabilizing* algorithms, i.e., algorithms which tolerate arbitrary initializations, and for algorithms that tolerate *asynchronous starts* (which is obviously the case of self-stabilizing algorithms).

We then consider four classical models for inter-agent communication. In the lowest-level model – namely, the *simple broadcast* model, an agent “blindly” sends a message, without knowing by whom, or by how many, this message will be received; the content of the message is entirely determined by the local state of the agent, and is the same for every recipient. This model can be enriched in two ways, either with the feature of *symmetric communications* – an agent j receives a message from another agent i if and only if i itself receives a message from j – or with *outdegree awareness* – i.e., when an agent broadcasts a message, it knows in advance how many other agents will receive this message. In the latter case, the content of the message is still the same for every recipient, but is no longer determined solely by the local state of the sender, as it can also depend on its current outdegree.

Symmetric communications arise in many natural systems – such as the popular Hegselmann-Krause model used to study the dynamics of opinion formation [19] – and is a basic feature of the pairwise interactions of the celebrated *population protocols* model [2]. Outdegree awareness is often implicitly assumed by distributed algorithms designed for engineering systems, such as the credit-recovery algorithm for termination detection [27] or the Push-Sum algorithm used for decentralized optimization problems [23].

A final communication model is given by *output port awareness*, in which case each agent is aware of output ports corresponding to each communication links. This amounts to a local output labelling: the outgoing links of an agent i have unique labels, and the messages sent by i over a given link may depend on the corresponding label. In this model, it is no longer the case that agent i sends the same message to each of its neighbors. We note in passing that this model is only meaningful in the context of *static* networks.

We will consider a general definition of what it means for a networked system to “compute” a function: given a metric space (X, δ) and a function of arbitrary arity $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, each agent holds an output that must asymptotically converge in (X, δ) to the value taken by f when its arguments correspond to the agents’ input values. In particular, when X is endowed with the discrete metric, all agents must eventually settle on the desired value, but they are *not* required to become aware when their outputs stabilize. In contrast, when $X = \mathbb{R}^k$ is equipped with the Euclidean metric, the outputs need only converge towards the desired outcome but are not required to ever stabilize on its exact value; this is actually an extremely common metric choice in the literature on distributed control. In all cases, regardless the metric on X , our definition of function computation contains no *termination* requirement: an agent is *not* supposed to return a value that is the desired value exactly, or to within a given precision.

For the simple broadcast model, a function f is computable if and only if f is *set-based*, i.e., if its value only depends on the *set* of its arguments. Set-based functions include, for example, the minimum and the maximum, but not the sum or the average. This result holds for any metric δ , for static and dynamic networks, and assuming or not that the

network size is known. The simple gossip algorithm clearly computes the set of input values, and thus any set-based function. For the impossibility result, it was shown by Hendrickx and Tsitsiklis for an arbitrary network size [21], but the stronger version under the assumption of a known network size had been previously shown by Boldi and Vigna [6].

In this paper, we characterize the functions that are computable under the assumptions of symmetric communications, output port awareness, or simply outdegree awareness. In particular, we show that these communication models are actually equivalent in terms of function computability.

Summary and contributions. First, we provide a general model of computation in anonymous networks with the features described above and which encompasses the various communication models that we have just introduced.

Then, we tackle the case of static networks: we prove that with either output port awareness, symmetric communications, or outdegree awareness, a function f is computable if and only if f is *frequency-based*, i.e., its value only depends on the set of its arguments and their *frequencies*. In particular, computing the average of initial values is possible while it is impossible with the simple broadcast model. However, computing the sum of initial values remains impossible. This result holds in any of the three communication models under consideration, for any metric δ , and assuming or not that a bound on the network size is known. Our approach for both positive and negative results exploits the notion of *graph fibration*, which originated from homotopy theory and has been used first in order to characterize the classes of anonymous networks in which leader election is possible [1, 5, 32].

The impossibility results all use the fundamental *lifting lemma* [8] stating that all the agents in the same *fibre*, i.e., with similar in-neighborhood, have the same behavior if they start in the same state. In fact, our impossibility proofs are a formalization in terms of graph fibration of the argument used in [20].

The first step for our positive results is the distributed algorithm that computes the *minimal base* of the network [8]. Then we show that the minimal base allows for computing the cardinalities of the fibres up to some common factor. For that, each agent solves a homogeneous system, which we prove to be of rank one in each of the three communication models under consideration. While the latter property clearly holds in the cases of output port awareness and symmetric communications, it requires a more sophisticated argument for the model with outdegree awareness, which has led us to develop a method *à la* Perron-Frobenius for matrices whose diagonal entries may be negative. The result is a self-stabilizing and finite-state algorithm for computing a frequency-based function in a static network that is linear in time in the number of agents.

As a consequence, when the network size is known or if the network is able to appoint a leader, our approach allows for computing any *multiset-based* function of the initial values, that is to say any function invariant under permutation. Hence, each of these two assumptions considerably increases the computational power in the case of output port awareness, symmetric communications, or outdegree awareness, while they leave it unchanged in the simple broadcast model [6].

Table 1 summarizes the computability results for static networks in our considered communication models and under various assumptions of centralized help (knowledge of n or of a bound over n , presence of a leader).

For dynamic networks, we first observe that the above impossibility results also hold in this case since it encompasses static networks. For the positive results, we develop a different method based on the stochastic analysis of consensus algorithms derived from statistical physics, namely the *Metropolis* and the *Push-Sum* algorithms. These algorithms allow for asymptotically computing the average of initial values in dynamic networks. They are not efficient in time and are not self-stabilizing, but they tolerate asynchronous starts and use no persistent memory. We provide a concise

Table 1. Computable functions in static, strongly connected networks of n anonymous agents

	SIMPLE BROADCAST	OUTDEGREE AWARENESS	SYMMETRIC COMMUNICATIONS	OUTPUT PORT AWARENESS
no centralized help	set-based Hendrickx <i>et al.</i> [20]	frequency-based Theorem 4.1, eq. (1)	frequency-based Theorem 4.1, eq. (4)	frequency-based Theorem 4.1, eq. (3)
a bound over n is known	set-based Boldi & Vigna [6]	frequency-based Corollary 4.2, eq. (1)	frequency-based Corollary 4.2, eq. (4)	frequency-based Corollary 4.2, eq. (3)
n is known	set-based ^a Boldi & Vigna [6]	multiset-based Corollary 4.3, eq. (1)	multiset-based Corollary 4.3, eq. (4)	multiset-based Corollary 4.3, eq. (3)
one leader	set-based ^b Boldi & Vigna [6]	multiset-based Corollary 4.4, eq. (1)	multiset-based Corollary 4.4, eq. (4)	multiset-based Corollary 4.4, eq. (3)

^a This is in fact for $n \geq 4$; for smaller networks, the topology always allows for the recovery of the multi-set of input values, as Jérémie Chalopin pointed out to us.

^b Even though Boldi and Vigna [6] do not consider networks with a leader, their impossibility argument can be adapted to this case.

and streamlined proof of the convergence of the Push-Sum algorithm to the average of the initial values in a dynamic network with a finite diameter.

In the case an upper bound on the network size is available, we obtain the same characterization of computable functions in dynamic networks with symmetric communications or outdegree awareness as in the static case, and this characterization holds for any metric, in particular for the discrete metric (exact computation). When no bound on the network size is known, these algorithms only achieve an approximate computation of the frequencies of the initial values, and thus of frequency-based functions if they satisfy some continuity property with respect to the frequencies of their arguments. Typical examples of such functions, which we call *continuous in frequency*, are the average function and the *threshold frequency predicates* with a non-rational threshold. We have thus proven that, in dynamic networks, the frequency-based condition which is necessary for approximate computability, is nearly sufficient in the sense that it simply needs to be enriched with continuity in frequency.

The results pertaining to dynamic networks are collected in Table 2.

Table 2. Computable functions in dynamic networks of n anonymous agents with finite diameter

	SIMPLE BROADCAST	OUTDEGREE AWARENESS	SYMMETRIC COMMUNICATIONS
no centralized help	set-based Hendrickx <i>et al.</i> [20]	?	frequency-based Di Luna & Viglietta [26]
a bound over n is known	set-based Hendrickx <i>et al.</i> [20]	frequency-based Corollary 5.3	frequency-based CB & LM [11]
n is known	set-based Hendrickx <i>et al.</i> [20]	multiset-based Corollary 5.4	multiset-based CB & LM [11]
one leader	set-based Hendrickx <i>et al.</i> [20]	?	multiset-based Di Luna & Viglietta [25]

Related works. There is a very large literature on computability in multi-agent systems, but most of it focuses on computing functions whose values may depend on the network topology, and not only on the initial values. Moreover, a common requirement is that all agents become aware that they produce the desired outputs. We refer the reader to [1, 18, 28, 34] for some fundamental results in this setting.

The biologically-inspired *population protocols* model has some common features with our model, namely a fixed set of anonymous agents with pairwise interactions² and no requirement of termination awareness. However, agents in this model are finite-state, and the fairness condition on interactions, despite implying that every pair of agents communicate infinitely often, does *not* require a bounded dynamic diameter. A spectacular result is the characterization of the predicates that are computable by population protocols: Angluin et al. [2, 3] proved that the class of computable predicates is exactly the class of predicates definable in Presburger arithmetic.

Closest to our work in the static case are [6] and [20]. The first paper characterizes the class of functions that are computable with simple broadcast or with symmetric communications when the network size is known. The second paper gives an almost characterization of the computable functions in a static network with symmetric communications when no bound on the network size is known: similarly to our results in the dynamic case, Hendrickx et al. proved that the frequency-based condition is sufficient only for an *approximate* computation of the frequencies of the initial values. Hence, we have solved the open question proposed in [20] of an exact characterization of computable functions with symmetric communications, and have extended it to communication models with output port awareness and outdegree awareness.

For dynamic networks, our positive result has to be compared with the remarkable algorithm proposed by Di Luna and Viglietta [25, 26] which allows for an exact computation of any frequency-based function in the case of symmetric communications. Their algorithm is linear in time, but it uses an infinite number of states and an infinite bandwidth in each of its executions. Moreover, it is not self-stabilizing and does not even tolerate asynchronous starts.

2 THE COMPUTING MODEL

2.1 Networked systems

We consider a networked system with a fixed and finite set of agents denoted $1, \dots, n$. Computation proceeds in *synchronized rounds*, which are communication closed in the sense that no message received in round t is sent in a round different from t . In round t ($t = 1, 2, \dots$), each agent successively (a) sends messages at the beginning of round t , (b) receives some messages, and (c) undergoes an internal transition to a new state. Communications that occur at round t correspond to the (directed) graph $\mathbb{G}(t) = ([n], E_t)$ where $[n] = \{1, \dots, n\}$ and $(i, j) \in E_t$ if and only if the agent j receives the message sent by i to j at round t . Hence, the agent $i \in [n]$ corresponds to the vertex i in each graph $\mathbb{G}(t)$, and we will sometimes refer to an agent or a vertex as a *node* of the network. We assume a self-loop at each vertex in each graph $\mathbb{G}(t)$, since an agent can communicate with itself instantaneously. The *network* is thus modeled by the *dynamic graph* \mathbb{G} , i.e., the infinite sequence of graphs $\mathbb{G} = (\mathbb{G}(t))_{t \geq 1}$, with the same set of vertices.

Each agent may possess more or less information about the network it belongs to. The knowledge of certain informations thus corresponds to some constraints on the network, and so to restrict to a non-empty subset of networks, called a *network class*. Since we consider *anonymous networks*, network classes are assumed to be closed under graph isomorphisms. As an example, the network class where the number of agents is known to be n is captured by the set of dynamic graphs with n vertices. Similarly, the class of *symmetric networks* corresponds to the set of dynamic graphs

²This corresponds to a dynamic network with symmetric communications and vertices of degree zero or one.

with bidirectional edges – that is, at each round t , $(i, j) \in E_t$ if and only if $(j, i) \in E_t$. We also consider the class of networks \mathbb{G} with a finite *dynamic diameter* – that is, there exists some positive integer D such that, for every $t \in \mathbb{N}_{>0}$, the graph product³ $\mathbb{G}(t) \circ \dots \circ \mathbb{G}(t + D - 1)$ is the complete graph. In other words, for every pair of vertices i, j and from every round t , there is a *dynamic path* of length at most D connecting i to j ; the smallest such integer D is called the *dynamic diameter* of \mathbb{G} . It measures connectivity *over time* and generalizes the diameter of static graphs. Note that with a dynamic diameter $D \geq 2$, some intermediate graphs in any period of length D may be disconnected (e.g., with only self-loops).

2.2 Algorithms, communication models, and executions

An *algorithm* \mathcal{A} is given by a set Q of local states, a subset $Q_0 \subseteq Q$ of initial states, a set of messages \mathcal{M} , a sending function, and a transition function.

The transition function determines the state after a transition: the new state is computed on the basis of the current state and the collection of messages that have been received. That corresponds to a transition function $\delta : Q \times \mathcal{M}^\oplus \rightarrow Q$, where \mathcal{M}^\oplus denotes the set of finite multi-sets over the set \mathcal{M} .

The messages to be sent by an agent depend on its current state and on its out-neighborhood, which can be compactly described by a local output labelling: if the outdegree is d^- , then the output ports are labelled with the numbers $1, \dots, d^-$. In this communication model, called *output port awareness*, a sending function is thus of the type $\sigma : Q \times \mathbb{N}_{>0} \rightarrow \bigcup_{k \in \mathbb{N}_{>0}} \mathcal{M}^k$. If q is the state of an agent and d^- is its outdegree, the message sent by the agent on the port labelled by $\ell \in [d^-]$ is the ℓ -th entry of $\sigma(q, d^-) \in \mathcal{M}^{d^-}$, denoted $\sigma(q, d^-)[\ell]$.

We can weaken this model by considering only the sending functions satisfying

$$\forall q \in Q, \forall k \in \mathbb{N}_{>0}, \forall \ell, \ell' \in [k]^2, \sigma(q, k)[\ell] = \sigma(q, k)[\ell'].$$

In this communication model called *outdegree awareness*, communications are isotropic – a sender sends the same messages to all its recipients – and sending functions are actually of type $\sigma : Q \times \mathbb{N}_{>0} \rightarrow \mathcal{M}$.

A further weakening of the communication model consists in requiring messages to depend only on the current state of the sender. In this *simple broadcast* model, sending functions satisfy the following graph-invariant property

$$\forall q \in Q, \forall k \in \mathbb{N}_{>0}, \forall \ell \in [k], \sigma(q, k)[\ell] = \sigma(q, 1)[1]$$

and thus correspond to functions of the type $\sigma : Q \rightarrow \mathcal{M}$.

Hence, these different notions of sending functions yield three communication models – namely, simple broadcast, outdegree awareness, and output port awareness. In the latter model, from an algorithmic point of view, we remark that labelling messages with output ports is only useful in the context of a static network and fixed output port labellings.

This description may be completed with a fourth model, called the model of *symmetric communications*, which corresponds to the restriction of the simple broadcast model to the class of networks with bidirectional links. For this model, in the case of static networks, each agent can determine its outdegree at the end of the receiving phase in the first round, since it is equal to its indegree, i.e., the number of messages it has just received. In other words, symmetric communications implies outdegree awareness in the class of static networks. This is no longer the case with dynamic networks, since the in/outdegree in round t is not yet available at the time of emission in the simple broadcast model, and may differ from one round to the next.

³Recall that the product $G = G_1 \circ G_2$ of two directed graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ with the same set of vertices V is the directed graph $G = (V, E)$ where $E = \{(j, i) \in V^2 : \exists k \in V, (i, k) \in E_1 \wedge (k, j) \in E_2\}$.

An *execution* of an algorithm \mathcal{A} in the dynamic graph \mathbb{G} proceeds as follows: In each round $t = 1, 2, \dots$, each agent applies the sending function σ to generate the message to be sent on each of its output port, then it receives the messages sent by its incoming neighbors in the graph $\mathbb{G}(t)$, and finally applies the transition function δ to its current state and the multi-set of messages it has just received to go to a next state.

An execution of \mathcal{A} in a network with n agents thus corresponds to an infinite sequence of *global states* C^0, C^1, C^2, \dots , where a global state is defined as a mapping $C : [n] \rightarrow \mathcal{Q}$. The sequence of global states is entirely determined by the initial global state C^0 and the dynamic graph \mathbb{G} . In the rest of the paper, we adopt the following notation: given an execution of \mathcal{A} , the value at the end of round t of any variable x_i , local to the agent i , is denoted by $x_i(t)$, and $x_i(0)$ is the initial value of x_i in this execution.

We may consider the more general model of executions with *asynchronous starts* [12], where each agent is activated in an arbitrary round. Whether the basic network is static or not, this execution model can be handled by a simple dynamic graph with inactive agents being modeled as isolated vertices. Regarding eventual convergence properties, a *self-stabilizing* algorithm [15] – i.e., an algorithm that works for an arbitrary initialization – obviously tolerates asynchronous starts. In contrast, in the self-stabilizing model, an agent cannot measure the time elapsed since it started the computation, while it can easily do it in the execution model with asynchronous starts. In this sense, self-stabilization is more restrictive than tolerance to asynchronous starts.

2.3 Computability in a metric space

Let Ω be a non-empty set and let (X, δ) be a metric space. Observe that the topology induced by δ is always coarser than the discrete metric δ_0 defined by

$$\delta_0(x, y) := \begin{cases} 0 & \text{if } x = y \\ 1 & \text{otherwise.} \end{cases}$$

If $X = \mathbb{R}^k$, then we may also consider the Euclidean distance

$$\delta_2(\mathbf{x}, \mathbf{y}) := \sqrt{(x_1 - y_1)^2 + \dots + (x_k - y_k)^2}.$$

Let $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$ be a function of arbitrary arity, and let \mathcal{A} be an algorithm with a set of local states of the form $\Omega \times X$ and such that its transition function does not modify the first state component in Ω . The first component of the local state of agent i is i 's input value, and the second one corresponds to the value in X of an output variable denoted x_i .

We say that the algorithm \mathcal{A} *δ -computes the function f in the network class C* if, in every execution of \mathcal{A} with a network $\mathbb{G} \in C$ composed of the agents $1, \dots, n$ and with the input values v_1, \dots, v_n in Ω , all sequences $(x_i(t))_{t \in \mathbb{N}}$ converge with respect to the distance δ to the same value $x^* = f(v_1, \dots, v_n)$.

The function f is said to be *δ -computable in the network class C* if there exists an algorithm that δ -computes f in C . Therefore, if f is δ -computable in a network class C , then it is δ -computable in any subclass of C . However, if f is δ -computable in two network classes C_1 and C_2 , then f may be not δ -computable in $C_1 \cup C_2$.

Since the discrete metric δ_0 defines the finest topology on X , if \mathcal{A} δ_0 -computes a function f , then it δ -computes f for any distance δ on X . Moreover, there exists a round after which all the variables x_i are equal to x^* , and \mathcal{A} is said to *compute f in finite time*. If $X = \mathbb{R}^k$ and \mathcal{A} δ_2 -computes f is the Euclidean distance δ_2 , then \mathcal{A} computes f asymptotically or approximately.

In this paper, we focus on the class of the *frequency-based* functions [20] whose output values only depend on the set of input values and their frequencies. For a formal definition, let us first introduce some additional definitions. We say that a function $\nu : \Omega \rightarrow \mathbb{Q}_{\geq 0}$ is a *frequency function* if it is positive for a finite set of values and $\sum_{\omega \in \Omega} \nu(\omega) = 1$. Given a vector $\mathbf{v} \in \Omega^n$, the *frequency function in \mathbf{v}* , denoted $\nu_{\mathbf{v}}$, is defined by:

$$\nu_{\mathbf{v}} : \omega \in \Omega \mapsto \frac{|\mathbf{v}^{-1}(\omega)|}{n} \in \mathbb{Q}$$

where $|\mathbf{v}^{-1}(\omega)|$ is the multiplicity of the value ω in \mathbf{v} .

Conversely, for every frequency function ν on Ω , there exist vectors whose frequency functions are equal to ν : For instance, let us consider a total ordering $\omega_1, \dots, \omega_\ell$ on ν 's support, irreducible representants p_k/q_k of the positive rational numbers $\nu(\omega_k)$, and the vector $\langle \nu \rangle$ where values in ν 's support occur in order $\omega_1, \dots, \omega_\ell$ and with multiplicities $p_1q/q_1, \dots, p_\ell q/q_\ell$, where $q = \text{lcm}(q_1, \dots, q_\ell)$. This vector of size q will be denoted by $\langle \nu \rangle$ in the sequel.⁴

The vectors with the same frequency function ν are said to be *ν -frequenced* and *equivalent in frequency*. A function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$ is then *frequency-based* if f takes the same value on all vectors equivalent in frequency, that is,

$$\forall \mathbf{v}, \mathbf{w} \in \Omega^n \times \Omega^m, \nu_{\mathbf{v}} = \nu_{\mathbf{w}} \implies f(\mathbf{v}) = f(\mathbf{w}).$$

Typically, the average function, whose value for a vector $\mathbf{v} \in \mathbb{R}^n$ is denoted $\bar{v} = \frac{\sum_{i \in [n]} v_i}{n}$, is frequency-based.

We will similarly consider the class of *set-based* functions, which take the same value on all vectors with the same support, and the class of *multiset-based* functions – also referred to as “symmetric” functions – which take the same value on all vectors that are permutations of one another. The three classes clearly obey the inclusion:

$$\text{set-based} \subseteq \text{frequency-based} \subseteq \text{multiset-based};$$

examples of functions of each class include, respectively, the maximum, the average, and the sum of the entries of a vector.

3 GRAPH FIBRATIONS

A (directed, multi-)graph G is composed of a non-empty set of vertices $V_G = [n_G]$, where $[n_G] := \{1, \dots, n_G\}$ and a set E_G of edges defined by two functions $s_G, t_G : E_G \rightarrow V_G$, which specify the source and the target vertices of each edge.

A *valued graph* in Ω is a graph G together with a function $\nu : V_G \rightarrow \Omega$, called a *valuation*. A *colored graph*, with a set of colors C , is a graph G together with a coloring function $c : E_G \rightarrow C$. In the case where Ω and C are singletons, the introduction of values and colors makes no difference, i.e., vertices are actually not valued and edges not colored.

Outdegree awareness and output port awareness then respectively correspond to a valuation and a coloring of any arbitrary graph G . The valued graph, where each vertex i of G is valued with its outdegree d_i^- , is denoted by G_{od} . Similarly, G_{op} denotes any colored graph resulting from a coloring of edges in G with local output port labellings.

Our first results with static networks rely on the notion of *fibration*, which originates in homotopy theory, and has been developed in order to characterize the classes of static anonymous networks in which leader election is solvable [5, 34] and functions are computable [6, 33].

A *graph morphism* $\varphi : G \rightarrow H$ is a pair of functions $\varphi_V : V_G \rightarrow V_H$ and $\varphi_E : E_G \rightarrow E_H$ that commute with the source and target functions, i.e., $s_H \circ \varphi_E = \varphi_V \circ s_G$. In the valued and colored cases, we also require $\nu_G = \nu_H \circ \varphi_V$ if ν_G and ν_H

⁴Actually, the vector $\langle \nu \rangle$ depends on the total ordering on ν 's support and is defined up to some permutation.

are valuations of the graphs G and H , and $c_G = c_H \circ \varphi_E$ if c_G and c_H are colors of G and H . When no confusion may arise, subscripts will be dropped. A *graph isomorphism* is a graph morphism that is additionally bijective, i.e., both functions φ_V, φ_E are bijective.

A *fibration* between (valued, colored) graphs G and B is a morphism $\varphi : G \rightarrow B$ such that for every edge $e \in E_B$ and for each $i \in V_G$ with $\varphi(i) = t(e)$, there exists a *unique* edge \tilde{e}^i verifying $\varphi(\tilde{e}^i) = e$ and $t(\tilde{e}^i) = i$. We restrict fibrations to be *epimorphisms* – that is, φ_V and φ_E are both surjective. Under this condition, fibrations preserve strong connectivity. All graph isomorphisms are fibrations, but the converse is generally false.

If $\varphi : G \rightarrow B$ is a fibration, then B is called the *base* of the fibration. The *fibres* over a vertex i of the base B is the set of vertices in G that are mapped to i , and is denoted by $\varphi^{-1}(i)$. A fibration $\varphi : G \rightarrow B$ induces an equivalence relation between the vertices of G , whose classes are precisely the fibres of φ : when two vertices j and k are in the same fibre, they have similar in-neighborhoods – that is, there is a bijective correspondence between the egres of G edges coming to j and those coming to k .

3.1 Lifting lemma

Impossibility results are based on the fundamental *Lifting lemma* [8, Lemma 2], which relates the behaviors of the same algorithm on two different networks. For a formal statement of this lemma, we first observe that the notion of execution of an algorithm defined in Section 2.2 for communication graphs formed with simple edge graphs naturally extends to multi-graphs. This extension works in the broadcast model as well as in the communication models with outdegree awareness and output port awareness. Then we introduce some additional notation: if $\varphi : G \rightarrow B$ is a fibration and $C : V_B \rightarrow \mathcal{Q}$ is a global state of the vertices in B , then we obtain a global state C^φ of the vertices in G by copying the state of a vertex of B fibrewise. Formally, we let

$$\forall i \in V_G, C_i^\varphi := C_{\varphi(i)}.$$

Similarly, any valuation ν of the vertices in B can be lifted into the valuation ν^φ of the vertices in G . From the restriction of the function f to Ω^{n_G} , we thus define the n_B -arity function f^φ as $f^\varphi(\nu) := f(\nu^\varphi)$. We now recall the statement of the Lifting lemma.

LEMMA 3.1 (LIFTING LEMMA). *Let $\varphi : G \rightarrow B$ be a fibration. Then, for every algorithm \mathcal{A} and every computation C_0, C_1, \dots of \mathcal{A} on B , the sequence $(C_0)^\varphi, (C_1)^\varphi, \dots$ is a computation of \mathcal{A} on G .*

By unicity of the limit (with respect to any distance δ), the Lifting lemma imposes strong constraints on the type of computable functions in a class of networks that can be “collapsed” by fibration onto other ones in this class.

LEMMA 3.2. *Let $\varphi : G \rightarrow B$ be a fibration. If a function f is δ -computed by some algorithm \mathcal{A} on both graphs G and B , then $f^\varphi = f$.*

In particular, if the function f is δ -computable in a class that contains both G and B , then the premise of Lemma 3.2 hold and $f^\varphi = f$. By applying Lemma 3.2 to fibrations that are graph isomorphisms, and because the network classes under consideration are closed under graph isomorphisms, we obtain that the value of a computable function only depends on the *multi-set* of its arguments, and this in any network class considered in this work – that is, only symmetric functions are computable in our model.

LEMMA 3.3. *If a function f is δ -computed by some algorithm \mathcal{A} in an anonymous network class, then f is a multiset-based function.*

Thus, the arguments of a computable function f are actually multi-sets in Ω^{\oplus} . If the support of a vector $\mathbf{v} \in \mathbb{R}^n$ is the set $\{\omega_1, \dots, \omega_k\}$ and if μ_ℓ denotes the multiplicity of ω_ℓ in \mathbf{v} , then we have

$$f(\mathbf{v}) = f(\omega_1^{\mu_1}, \dots, \omega_k^{\mu_k}),$$

where $\omega_\ell^{\mu_\ell}$ denotes the sequence $\omega_\ell, \dots, \omega_\ell$, of length μ_ℓ .

3.2 Minimum base

A graph G is said to be *fibration prime* if every fibration from G is an isomorphism – that is, if G cannot be collapsed onto a smaller graph by a fibration. Every graph G has exactly one fibration prime basis (up to some isomorphism), called the *minimum base* of G .

Boldi and Vigna [8] constructed a self-stabilizing algorithm which distributively computes, in finite time, the minimum base of the underlying graph of a static and strongly connected network. More precisely, when running this algorithm over the static network G , each agent builds a graph at the end of each round that, from round $n + D$ onwards, is guaranteed to be the minimum base of G , where n is the number of vertices in G and D is its diameter. Then they proposed a finite-state variant of this algorithm that fits our model and stabilizes with an overhead of $O(D \log D)$ rounds when compared to the infinite-state version.

4 COMPUTABILITY IN STATIC NETWORKS

In this section, we establish our main theorem for static networks, namely that a function of arbitrary arity is computable if and only if its output value only depends on the frequencies of input values. This result holds whatever the distance is and even when agents only know their outdegrees.

THEOREM 4.1. *Let (X, δ) be a metric space. In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is frequency-based;
- (ii) f is δ -computable in the class of static strongly connected networks.

4.1 Proof of the impossibility result

The main argument in the proof of the impossibility result in Theorem 4.1 already appears in [6, 20]. We recall it in order to make the paper self-contained and to show that the result holds whatever the distance is and in any of the three communication models under consideration.

Proof of (ii) \Rightarrow (i).

Let \mathbf{v} and \mathbf{w} be two Ω -vectors, of respective lengths n and m , with the same frequency functions $v_{\omega} = v_{\mathbf{w}}$. They share the same support $\{\omega_1, \dots, \omega_\ell\}$, and we let $v_{\omega}(\omega_k) = \frac{p_k}{q_k}$ with $\gcd(p_k, q_k) = 1$. Then the integer $p = \text{lcm}(q_1, \dots, q_\ell)$ divides both n and m .

Let us now consider the bidirectional rings R^p , R_{op}^p , and R_{od}^p , of size p , with or without output port awareness and outdegree awareness. The mapping $i \in [n] \mapsto i \bmod p \in [p]$ induces a fibration $\varphi : R^n \rightarrow R^p$. This fibration preserves both the deterministic coloration of the outgoing links and the outdegree valuation of the vertices. Similarly, we define a fibration $\psi : R^m \rightarrow R^p$, which is also a fibration of the colored and valued rings R_{op}^m and R_{od}^m .

Let \mathcal{A} be an algorithm that δ -computes a function f in the class of static strongly connected networks, with or without output port or outdegree awareness. In particular, \mathcal{A} δ -computes f on R^n (resp. R_{op}^n and R_{od}^n). Lemma 3.2 shows that \mathcal{A} δ -computes the function f^φ on R^p (resp. R_{op}^p and R_{od}^p). Similarly, \mathcal{A} δ -computes the function f^ψ on R^p (resp. R_{op}^p and R_{od}^p). We thus obtain that $f^\varphi = f^\psi$, that is for every vector $\mathbf{u} \in \Omega^p$, $f(\mathbf{u}^\varphi) = f(\mathbf{u}^\psi)$.

The vector $\mathbf{u} = (\omega_1^{p_1/q_1}, \dots, \omega_k^{p_k/q_k}) = \langle v_\mathbf{v} \rangle$ is of size p . Then \mathbf{u}^φ and \mathbf{u}^ψ are respectively obtained by permuting the entries of \mathbf{v} and \mathbf{w} . Lemma 3.3 shows that $f(\mathbf{v}) = f(\mathbf{u}^\varphi)$ and $f(\mathbf{w}) = f(\mathbf{u}^\psi)$. The equality $f^\varphi = f^\psi$ leads to $f(\mathbf{v}) = f(\mathbf{w})$ as required.

4.2 Positive result with outdegree awareness

We now devise an algorithm that computes any frequency-based function f in the class of static strongly connected networks with outdegree awareness in linear time. Since the discrete metric δ_0 defines the finest topology on X , our algorithm also δ -computes f for any distance δ on X . For that, we use the self-stabilizing algorithm of Boldi and Vigna [8] that constructs distributively the minimum base of the underlying network G . Their algorithm relies firstly on the inductive construction by each agent of its *view* in the graph G , and secondly on a partial function \mathcal{B} that allows an agent to extract from a truncated and possibly incorrect version of its view a candidate for the minimum base of G . More precisely, in each round t , each agent i builds an in-tree T_i^t and a multi-graph $\mathcal{B}(T_i^t)$ that is guaranteed to be the minimum base of the graph G (up to some isomorphism) from round $n + D$, where n is the number of vertices in G and D is its diameter. The algorithm is then made finite-state with a loss of less than $D \log(1 + D)$ rounds in the stabilization time. Moreover, it straightforwardly adapts to any valued version of G .

Proof of (i) \Rightarrow (ii) in the case of outdegree awareness.

Let $G_{\mathbf{v}, d^-}$ be a double-valued graph of size n and of finite diameter D . The first valuation \mathbf{v} forms a vector in Ω^n ; we let S denote \mathbf{v} 's support. The second valuation $d^- \in \mathbb{N}^n$ is the valuation of the vertices with their outdegrees. As above explained, at each round $t \geq n + D + D \log(1 + D)$ of Boldi and Vigna's algorithm in the graph $G_{\mathbf{v}, d^-}$, each agent builds the minimum base $B_{\mathbf{w}, b}$ of $G_{\mathbf{v}, d^-}$. Let $\varphi : G_{\mathbf{v}, d^-} \rightarrow B_{\mathbf{w}, b}$ be the corresponding fibration and let $m = |V_B|$.

All the vertices of a fibre share the same entry in \mathbf{v} and the same outdegree:⁵

$$\forall i \in V_B, \forall k \in \varphi^{-1}(i), \quad w_i = v_k \quad \text{and} \quad b_i = d_k^-.$$

Clearly, the m fibres satisfy the following equalities:

$$b_i |\varphi^{-1}(i)| = \sum_{j \in V_B} d_{i,j} |\varphi^{-1}(j)|, \quad (1)$$

where $d_{i,j}$ denotes the number of edges in B from i to j . The minimum base $B_{\mathbf{w}, b}$ entirely determines the square matrix \mathbf{M} of size m defined by

$$M_{i,j} = \begin{cases} d_{i,j} & \text{if } i \neq j \\ d_{i,i} - b_i & \text{if } i = j. \end{cases}$$

Let us now consider the linear system $\mathbf{M} \mathbf{z} = \mathbf{0}$ where $\mathbf{0}$ is the zero vector of size m . Observe that the system $\mathbf{M} \mathbf{z} = \mathbf{0}$ has a positive solution \mathbf{z} whose entries are given by $z_i = |\varphi^{-1}(i)|$.

Define the matrix $\mathbf{P} := \mathbf{M} + \alpha \mathbf{I}$, where \mathbf{I} is the identity matrix and α is an arbitrary real number satisfying

$$\alpha > - \min_{1 \leq i \leq m} M_{i,i}.$$

⁵Observe that b_i may be different from the outdegree of the vertex i in the graph B .

The matrix P is non-negative, and its diagonal entries are all positive. Hence, there is a self-loop at each vertex of the associated graph G_P associated⁶ to P . Moreover, G_P coincides with the support of the graph $B_{w,b}$, except possibly for the self-loops. Therefore, this graph is strongly connected, i.e., the matrix P is irreducible.

The Perron-Frobenius theorem then applies to the matrix P : the spectral radius ϱ_P of P is an eigenvalue of P of algebraic and geometric multiplicity one, and associated with an eigenvector \mathbf{x} with positive entries. Hence, $\lambda = \varrho_P - \alpha$ is a real eigenvalue of M , and every (complex) eigenvalue of M other than λ has a real part less than λ . The eigenvalue λ has algebraic and geometric multiplicity one with the positive eigenvector \mathbf{x} .

If $\mathbf{y} \in (\mathbb{R}_{>0})^m$ is a non-zero eigenvector of the matrix M for a real eigenvalue μ , then \mathbf{y} is also an eigenvector of the matrix P for the real eigenvalue $\mu + \alpha$. As just shown, we have:

$$\mu + \alpha \leq \varrho_P.$$

Let y_i be a non-zero entry of the vector \mathbf{y} . Since P is non-negative, the i -th entry of $P\mathbf{y}$ satisfies:

$$(P\mathbf{y})_i = (\mu + \alpha) y_i \geq P_{i,i} y_i$$

which implies that $\mu + \alpha > 0$. We now show that \mathbf{y} is a positive vector. For that, we use the strong connectivity of the graph G_P , and prove by finite induction over ℓ , $0 \leq \ell \leq m - 1$, that the entry y_j is positive whenever j is at distance ℓ from i in G_P .

- (1) The base case $\ell = 0$ is by definition of i .
- (2) Inductive step: let $\ell \geq 1$ and assume that the entries of the vector \mathbf{y} for all vertices at distance $\ell - 1$ from i are positive. The vertex j has an outgoing neighbor k at distance $\ell - 1$ from i , and the j -th entry of $P\mathbf{y}$ satisfies:

$$(P\mathbf{y})_j = (\mu + \alpha) y_j \geq P_{j,k} y_k.$$

Since both $P_{j,k}$ and $\mu + \alpha$ are positive, the inductive assumption implies that y_j is positive.

Hence, we can choose a positive real number ε small enough in order to have the componentwise inequality:

$$\varepsilon \mathbf{x} \leq \mathbf{y}.$$

Since the matrix P is non-negative, we obtain the following inequality for every positive integer k :

$$\varepsilon P^k(\mathbf{x}) = \varepsilon (\varrho_P)^k \mathbf{x} \leq P^k(\mathbf{y}) = (\mu + \alpha)^k \mathbf{y}.$$

It follows that $\varrho_P \leq \mu + \alpha$, and thus $\mu = \varrho_P - \alpha = \lambda$.

Since the system $M\mathbf{z} = \mathbf{0}$ has a positive solution whose entries are equal to the cardinalities of the fibres, we deduce that $\lambda = 0$, and the set of solutions of this system, namely $\ker M$, is a subspace of dimension one.

In a second step of the algorithm, using Gaussian elimination over the Euclidean ring \mathbb{Z} (see e.g., [22]), each agent computes a positive integer vector $\mathbf{z} \in \mathbb{N}^m$ whose all entries are coprime and such that $\ker M = \mathbb{R}\mathbf{z}$. Subsequently, the agent computes $f(\tilde{\mathbf{v}})$ where $\tilde{\mathbf{v}} \in \Omega^P$ is a vector of size $p := \sum_{i=1}^m z_i$ and where each value v_i in S occurs with multiplicity z_i . Since we have just proved that there exists a positive integer k such that

$$\forall i \in V_B, |\varphi^{-1}(i)| = k z_i, \tag{2}$$

and f is a frequency-based function, each agent actually outputs the value $f(\tilde{\mathbf{v}}) = f(\mathbf{v})$.

⁶See Section 5.2 for a definition of the associated graph.

In each round t , each agent i builds the finite tree T_i^t and the (valued) multi-graph $\mathcal{B}(T_i^t)$, and then applies the Gaussian elimination method over the Euclidean ring to solve the linear system corresponding to $\mathcal{B}(T_i^t)$. The above recalled result by Boldi and Vigna [8] on the graphs $\mathcal{B}(T_i^t)$ implies that, in this way, each agent computes the value $f(\mathbf{v})$ no later than in round $n + D$.

4.3 Positive results with output port awareness and symmetric graphs

Each agent can easily retrieve its outdegree in a bidirectional network as well as when it is output port aware. The above algorithm with a preliminary phase of outdegree calculation thus allows agents to compute any frequency-based function in the models with symmetric communications or with output port awareness. We now present two variants of the algorithm for symmetric communications and with output port awareness accordingly, which directly compute frequency-based functions without pre-calculation of the outdegrees, leading to linear systems that can be easily solved without the use of Gaussian elimination.

Output port awareness. With output port awareness, any fibration is actually a *covering*, i.e., for any pair of vertices i and j in the same fibre, the outgoing edges of i and j are in one-to-one correspondence. This local isomorphism property gives a bijective correspondence between the whole neighborhoods of two vertices in the same fibre and, as a result, the cardinality of all fibres is the same (see e.g., [7]). In the case of output port awareness, eq. (1) is thus replaced by:

$$|\varphi^{-1}(i)| = |\varphi^{-1}(j)|. \quad (3)$$

Each agent builds the multi-graphs $\mathcal{B}(T_i^t)$ which are eventually equal to the minimum base $B_{\mathbf{w}}$ of the (colored and valued) graph $G_{\mathbf{v}}$. If p denotes the common cardinality of all the fibres, then $\mathbf{v} = (w_1^p, \dots, w_m^p)$. It follows that $f(\mathbf{v}) = f(\mathbf{w})$ since f is a frequency-based function. In this way, each agent can thus directly compute the value $f(\mathbf{v})$ from the construction of the valued multi-graph $B_{\mathbf{w}}$.

Symmetric communications. If the network G is bidirectional, then for any fibration $\varphi : G \rightarrow B$, we have:

$$d_{i,j} |\varphi^{-1}(j)| = d_{j,i} |\varphi^{-1}(i)| \quad (4)$$

where $d_{i,j}$ denotes the number of $i \rightarrow j$ edges in the multi-graph B .

Let m the number of vertices in the graph B . Up to some permutation of the vertices in B , we may assume that none of the degrees $d_{1,2}, \dots, d_{m-1,m}$ is zero, since B is strongly connected. Hence, the cardinalities of the fibres form a solution of a linear system of m equations and m variables whose set of solutions is thus of dimension at least one. Moreover, the positive integer vector \mathbf{z} defined by:

$$z_i := \begin{cases} d_{1,2} \times \dots \times d_{m-1,m} & \text{if } i = 1 \\ \frac{d_{2,1} \times \dots \times d_{i,i-1}}{d_{1,2} \times \dots \times d_{i-1,i}} z_1 & \text{if } i \neq 1 \end{cases}$$

is obviously a basis of the solution set. Consequently, if $\tilde{\mathbf{v}}$ denotes a vector with the same support as \mathbf{v} and where each value $v_i \in S$ occurs with multiplicity z_i , then $f(\tilde{\mathbf{v}}) = f(\mathbf{v})$, since the function f is frequency-based. eq. (4) thus yields an algorithm that directly computes the function f from the construction of the minimum base $B_{\mathbf{w}}$ in the case of a bidirectional network.

4.4 Computing with knowledge on the network size

As explained in Section 2.1, computability of a function when the number of agents is known means that for every positive integer n , there exists an algorithm which computes the function in the class C_n^s of static strongly connected networks with n agents. Similarly, the function is computable when an upper bound on the network size is known if for every positive integer N , there is an algorithm which computes the function in the network class $C_{\leq N}^s = C_1^s \cup \dots \cup C_N^s$.

A refinement of the proof in Section 4.1 shows that the impossibility result still holds when a bound on the size of the network is known. With the above positive results, we thus obtain the following corollary.

COROLLARY 4.2. *Let (X, δ) be a metric space. In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is frequency-based;
- (ii) f is δ -computable in every network class $C_{\leq N}^s$.

PROOF. We only need to prove that (ii) \Rightarrow (i). For that, we refine the argument in the impossibility proof in Theorem 4.1. We consider two Ω -vectors of length n and m , with the same frequency functions and continue the proof by replacing computability in the class of strongly connected networks with computability in the sub-network class $C_{\leq N}^s$ with $N = \max(n, m)$. \square

When the exact size of the network is known, deducing multiplicities from frequencies is straightforward, and computing frequency-based functions thus allows for computing multiset-based functions. Hence, knowing the size of the network considerably increases the computational power in the case of output port awareness, symmetric communications, or outdegree awareness, while leaving it unchanged in the simple broadcast model [6].

Any function which is computable in a network class closed under graph isomorphisms is necessarily invariant under permutation, i.e., is a multiset-based function. We thus obtain the following computability result when the network size is known.

COROLLARY 4.3. *Let (X, δ) be a metric space. In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is multiset-based;
- (ii) f is δ -computable in every network class C_n^s .

4.5 Computing with leaders

We now study the impact of having leaders on the computational power in a static network with either output port awareness, symmetric communications, or outdegree awareness. In the case of a unique leader, its fibre is of cardinality one, and so the linearity coefficient k in eq. (2) is equal to one. Hence, agents compute the cardinality of each fibre, and our algorithm thus computes any function that is multiset-based.

COROLLARY 4.4. *Let (X, δ) be a metric space. In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is multiset-based;

(ii) f is δ -computable in the class of static strongly connected networks with one leader.

Observe that this result can be easily extended to the case of ℓ leaders if ℓ is known of all agents: the vertices in a graph G corresponding to the leaders collapse onto some subset L_B of vertices in the minimum base B of G . The cardinality of each fibre $\varphi^{-1}(i)$ is then given by:

$$|\varphi^{-1}(i)| = \frac{\ell}{\sum_{j \in L_B} z_j} z_i, \quad (5)$$

where z is the positive integer vector whose all entries are coprime and such that $\ker M = \mathbb{R} z$. Hence, our algorithm, together with eq. (5), allows each agent to compute any multiset-based function.

5 DYNAMIC NETWORKS

We now investigate how to compute a frequency-based function with symmetric communications or outdegree awareness in dynamic networks with a finite dynamic diameter. In the case of symmetric communications, a remarkable recent algorithm due Di Luna and Viglietta [25, 26] exactly computes any frequency-based function in linear time in the dynamic diameter of the network, solving an important open question of computability. In particular, for a dynamic graph that is strongly connected in each round, their algorithm operates in linear time in the size of the network. Unfortunately, this algorithm is not self-stabilizing and does not even tolerate asynchronous starts. Moreover, this algorithm is based on the construction, by each agent, of an infinite *history tree*, and so uses an infinite number of states and an infinite bandwidth in each of its executions.

In this section, we propose to develop another method which consists in using consensus algorithms derived from statistical physics – namely, the *Metropolis* and the *Push-Sum* algorithms. These algorithms only achieve asymptotic convergence and their temporal complexity is non-linear. However, both tolerate asynchronous starts and use no persistent memory.

The Metropolis algorithm computes the average of initial values in the class of *symmetric* networks with a finite dynamic diameter and in the communication model of outdegree awareness, even under asynchronous starts. Its convergence rate has been showed to be quadratic [10] in the case of a dynamic network that is strongly connected in each round.⁷ A variant for the simple model of symmetric communications (without assuming outdegree awareness) has been proposed [11, 24], but its temporal complexity is in $O(n^4)$.

In the case of outdegree awareness, we develop another approach based on the *Push-Sum* algorithm. The algorithm was introduced in [23], where its correctness was shown in a probabilistic communication model with pairwise communications in the fully-connected graph. This result was then extended to arbitrary strongly connected graphs in [4]. Further, Nedić et al. [30] proved the correctness of Push-Sum in any dynamic network with a finite diameter. Below, we give a self-contained and streamlined convergence proof, and then describe how Push-Sum can be used to compute frequency-based functions.

⁷Using the *Lazy Metropolis* algorithm [30, 31], this result can be extended to the case of (symmetric) networks with a finite dynamic diameter.

5.1 The Push-Sum algorithm and the quot-sum function

The Push-Sum algorithm proceeds as follows: each agent i maintains three variables x_i , y_i , and z_i . The two variables y_i and z_i are initialized respectively to $v_i \in \mathbb{R}$ and $w_i \in \mathbb{R}_{>0}$, and they are updated as follows:

$$y_i(t) = \sum_{k \in \text{In}_i(t)} \frac{y_k(t-1)}{d_k^-(t)} \quad (6)$$

$$z_i(t) = \sum_{k \in \text{In}_i(t)} \frac{z_k(t-1)}{d_k^-(t)} . \quad (7)$$

The variable x_i is initialized to v_i/w_i and set to $x_i = y_i/z_i$ at the end of each round. Observe that by the very definition of its update rules, the Push-Sum algorithm requires output port awareness.

The main result of this section is that, under the assumption of a network with a finite dynamic diameter, the Push-Sum algorithm computes the *quot-sum function* defined by

$$\begin{aligned} qs : \bigcup_{n \in \mathbb{N}_{>0}} (\mathbb{R} \times \mathbb{R}_{>0})^n &\rightarrow \mathbb{R} \\ ((v_1, w_1), \dots, (v_n, w_n)) &\mapsto \frac{\sum_{k \in [n]} v_k}{\sum_{k \in [n]} w_k} . \end{aligned}$$

In other words, for each agent i , the quotient $x_i(t)/y_i(t)$ asymptotically converges to the quot-sum of the initial values.

5.2 Preliminaries

We first introduce some notation. Let n be a positive integer, $\boldsymbol{v} \in \mathbb{R}^n$ a real vector, and $\boldsymbol{A} \in \mathbb{R}^n \times \mathbb{R}^n$ a real square matrix, both of size n . The vector \boldsymbol{v} or the matrix \boldsymbol{A} is said to be *non-negative* (resp. *positive*) if all its entries are non-negative (resp. positive). The *graph associated to* a non-negative matrix \boldsymbol{A} is the directed graph $G_{\boldsymbol{A}} = ([n], E_{\boldsymbol{A}})$, where $E_{\boldsymbol{A}}$ is the set of edges defined as

$$E_{\boldsymbol{A}} := \{(j, i) \in [n]^2 : A_{i,j} > 0\}.$$

A vector is *stochastic* if it is non-negative and its entries sum to 1; a *matrix* is in turn (row-) *stochastic* if each of its rows is a stochastic vector; correspondingly, a matrix is *column-stochastic* if each of its *columns* is a stochastic vector. Importantly, the sum of entries of a vector is left invariant by any column-stochastic matrix \boldsymbol{A} – namely, $\sum_{i \in [n]} (\boldsymbol{A} \boldsymbol{v})_i = \sum_{i \in [n]} v_i$.

Let $\boldsymbol{A} = (\boldsymbol{A}(t))_{t \in \mathbb{N}_{>0}}$ be a sequence of non-negative square matrices of size n . The entry of $\boldsymbol{A}(t)$ at the i -th row and j -th column is denoted by $A_{i,j}(t)$. For any positive real number α , the matrix $\boldsymbol{A}(t)$ is said to be α -*safe* if all its positive entries lie in the interval $[\alpha, +\infty)$. The graph associated to the matrix $\boldsymbol{A}(t)$ is denoted $\mathbb{G}_{\boldsymbol{A}}(t)$, and the sequence $\mathbb{G}_{\boldsymbol{A}}(1), \mathbb{G}_{\boldsymbol{A}}(2), \dots$ forms a *dynamic graph* $\mathbb{G}_{\boldsymbol{A}}$ over the set of nodes $[n]$. If $t, t' \in \mathbb{N}_{>0}$, with $t' \geq t$, we let $\boldsymbol{A}(t' : t)$ denote the backward product $\boldsymbol{A}(t') \times \dots \times \boldsymbol{A}(t)$. Hence, the graph associated to the matrix $\boldsymbol{A}(t' : t)$ is the (forward) product $\mathbb{G}(t : t') = \mathbb{G}(t) \circ \dots \circ \mathbb{G}(t')$.

LEMMA 5.1. *Let $(\boldsymbol{A}(t))_{t \in \mathbb{N}_{>0}}$ be a sequence of α -safe column-stochastic matrices of size n with positive diagonal entries, and let $\boldsymbol{v} \in \mathbb{R}^n$ be a non-negative vector. If the associated dynamic graph $\mathbb{G}_{\boldsymbol{A}}$ has a finite dynamic diameter D , then for all $i \in [n]$ and all $t \geq D$, we have*

$$\alpha^D \sum_{k \in [n]} v_k \leq v_i(t) \leq \sum_{k \in [n]} v_k$$

where $\boldsymbol{v}(t) := \boldsymbol{A}(t : 1) \boldsymbol{v}$.

PROOF. By induction, we easily check that, for any $t \in \mathbb{N}$, all the entries of $\boldsymbol{v}(t)$ are positive. Moreover, the sum of the entries in $\boldsymbol{v}(t)$ is invariant with t . It follows that, for every integer $t \in \mathbb{N}$, we have

$$v_i(t) \leq \sum_{k \in [n]} v_k(t) = \sum_{k \in [n]} v_k.$$

For the lower bound, we start by observing that since all the positive entries of the matrices $\boldsymbol{A}(1), \boldsymbol{A}(2), \dots$ are at least equal to α , the positive entries of every product of t matrices in this sequence are at least equal to α^t . In particular, if $t \geq D$, then all the positive entries of the matrix $\boldsymbol{A}(t : t - D + 1)$ are at least equal to α^D . Moreover, the graph associated to the product matrix $\boldsymbol{A}(t : t - D + 1)$ is the graph $\mathbb{G}(t - D + 1 : t)$ equal to the complete graph, since D is the dynamic diameter of \mathbb{G} . Therefore, all the entries of the matrix $\boldsymbol{A}(t : t - D + 1)$ are positive, and thus at least equal to α^D .

Pick $t \geq D$. We have $\boldsymbol{v}(t) = \boldsymbol{A}(t : t - D + 1) \boldsymbol{v}(t - D)$, and thus

$$v_i(t) = \sum_{k \in [n]} A_{i,k}(t : t - D + 1) v_k(t - D) \geq \alpha^D \sum_{k \in [n]} v_k(t - D).$$

Since the sum of the entries in $\boldsymbol{v}(t)$ is constant, we obtain $v_i(t) \geq \alpha^D \sum_{k \in [n]} v_k$. \square

5.3 Push-Sum for computing the quot-sum function

We are now in position to prove that the Push-Sum algorithm computes the quot-sum of initial values.

THEOREM 5.2. *The Push-Sum algorithm computes the quot-sum function in the class of networks with a finite dynamic diameter. More precisely, in any execution of Push-Sum with a network of dynamic diameter D , all the output variables are within ε of the quot-sum of the initial values in $O\left(n^{2D} \log \frac{1}{\varepsilon}\right)$ rounds.*

PROOF. We first consider an execution of the Push-Sum algorithm in a dynamic network \mathbb{G} with n agents and a finite dynamic diameter D , and synchronous starts at round one for all agents. Observe that if $\boldsymbol{A}(t)$ is the square matrix defined from the directed graph $\mathbb{G}(t)$ by

$$A_{i,j}(t) = \begin{cases} 1/d_j^-(t) & \text{if } (j, i) \in E(t) \\ 0 & \text{otherwise} \end{cases}$$

where $d_j^-(t)$ denotes the outdegree of j in $\mathbb{G}(t)$, then $\boldsymbol{A}(t)$ corresponds to the update rules for the variables y_i and z_i at round t , namely

$$\boldsymbol{y}(t) = \boldsymbol{A}(t) \boldsymbol{y}(t - 1) \text{ and } \boldsymbol{z}(t) = \boldsymbol{A}(t) \boldsymbol{z}(t - 1).$$

Each matrix $\boldsymbol{A}(t)$ is column-stochastic and $\frac{1}{n}$ -safe. Lemma 5.1 shows that for all $i \in [n]$ and $t \geq D$,

$$n^{-D} \sum_{k \in [n]} w_k \leq y_i(t) \leq \sum_{k \in [n]} w_k. \tag{8}$$

The vector $\boldsymbol{z}(t)$ is positive, and thus $\boldsymbol{x}(t) = [\text{diag}(\boldsymbol{z}(t))]^{-1} \boldsymbol{y}(t)$. It follows that

$$\boldsymbol{x}(t) = \boldsymbol{B}(t) \boldsymbol{x}(t - 1)$$

where $\boldsymbol{B}(t) := [\text{diag}(\boldsymbol{z}(t))]^{-1} \boldsymbol{A}(t) \text{diag}(\boldsymbol{z}(t - 1))$.

We easily check that all the entries of $\boldsymbol{B}(t)$ are non-negative, and $\boldsymbol{B}(t)$ is a stochastic matrix with a positive diagonal. Its associated graph is the same as $\boldsymbol{A}(t)$'s, namely $\mathbb{G}(t)$.

The next step of the proof consists in proving that the product matrix $B(t : 1)$ converges to a rank one matrix.⁸ Let us first observe that Dobrushin's ergodic coefficient [16, eq. (1.5)] of a stochastic matrix P , defined by

$$\delta(P) := 1 - \min_{i \neq j} \sum_{k \in [n]} \min(P_{i,k}, P_{j,k})$$

lies in the range $[0, 1]$ and satisfies the inequality

$$\delta(P) \leq 1 - n\alpha$$

when P is α -safe and its associated graph is fully-connected. A result by Seneta [17] combined with a straightforward argument of convex duality shows that for any stochastic matrix P , the ergodic coefficient $\delta(P)$ coincides with the matrix seminorm

$$\sup_{\delta(v) > 0} \frac{\delta(Pv)}{\delta(v)}$$

associated to the seminorm on \mathbb{R}^n defined by $\delta(v) = \max_{i \in [n]} v_i - \min_{i \in [n]} v_i$. Consequently, the ergodic coefficient δ is a matrix seminorm, and so is sub-multiplicative.

The sequence $(A(t))_{t \in \mathbb{N}_{>0}}$ is $\frac{1}{n}$ -safe, and the inequalities in eq. (8) show that every matrix product

$$B(t + D - 1 : t) = [\text{diag}(z(t + D - 1))]^{-1} A(t + D - 1 : t) \text{diag}(z(t - 1))$$

is n^{-2D} -safe. It follows that

$$\delta(B(t : 1)) \leq \left(1 - n^{-2D}\right)^{\lfloor t/D \rfloor}.$$

Because of the inequality $\log(1-a) \leq -a$, valid whenever $a \geq 0$, we obtain that if $t \geq D n^{2D} \log(\frac{1}{\varepsilon})$, then $\delta(B(t : 1)) \leq \varepsilon$ and $\delta(x(t)) \leq \varepsilon \delta(x(0))$.

This shows that $\lim_{t \rightarrow \infty} x^+(t) - x^-(t) = 0$, where $x^-(t) := \min_{i \in [n]} x_i(t)$ and $x^+(t) := \max_{i \in [n]} x_i(t)$. Since each matrix $B(t)$ is stochastic, the sequences $(x^+(t))_{t \in \mathbb{N}}$ and $(x^-(t))_{t \in \mathbb{N}}$ are non-increasing and non-decreasing, respectively. Hence, the two sequences converge to the same limit, that we denote x^* , and all the sequences $(x_i(t))_{t \in \mathbb{N}}$ also converge to x^* . The convergence rate follows from the above.

Every sequence $(z_i(t))_{t \in \mathbb{N}}$ is bounded (see eq. (8)), and hence $\lim_{t \rightarrow \infty} y_i(t) - x^* z_i(t) = 0$. Since the sum of the entries in $y(t)$ and $z(t)$ are constant, this implies that $\sum_{k \in [n]} v_k = x^* \sum_{k \in [n]} w_k$. \square

Push-Sum is not a self-stabilizing algorithm (initializations of the variables y_i and z_i cannot be arbitrary), but it tolerates asynchronous starts. Clearly, an execution with the dynamic graph \mathbb{G} and the agents i starting at rounds s_i is similar to the execution where all the agents start at round one and with the dynamic graph $\tilde{\mathbb{G}}$ with n vertices and the set of edges defined by:

$$\tilde{E}_t = \{(i, j) \in E_t : i = j \vee t \geq \max(s_i, s_j)\}.$$

Observe that if \mathbb{G} has a finite dynamic diameter D , then $\max(s_i) + D$ is an upper bound on $\tilde{\mathbb{G}}$'s dynamic diameter.

5.4 Push-Sum for computing a frequency-based function

A Push-Sum based algorithm for computing the frequency function $v : v \in \Omega^n \rightarrow v_v \in \mathbb{Q}^\Omega$ can be easily derived from Theorem 5.2 (see Algorithm 1). The variables x_i , y_i , and z_i are now three arrays of dynamic size, each of which initially contains only one variable, indexed by the initial value of vertex i and equal to one. As soon as the vertex i becomes

⁸This convergence result can be directly deduced from a theorem [9, Theorem 3] established by Cao, Morse, and Anderson for dynamic graphs that are rooted with bounded delay. The proof that we develop here in the particular case of a finite dynamic diameter is simpler and provides a better bound on the convergence rate.

aware that some value $\omega \in \Omega$ is initially present in the network, upon the first receipt of some variables indexed by ω , it appends three variables to x_i , y_i , and z_i all indexed by ω , and respectively initialized to 0, 0, and 1. Then, vertex i starts to run the Push-Sum algorithm with the variables $y_i[\omega]$ and $z_i[\omega]$.

Algorithm 1: The Push-Sum algorithm for computing the frequency function

```

1 Input:  $v_i \in \Omega$ 
2 Initially:
3    $x_i[v_i] \leftarrow 0, y_i[v_i] \leftarrow 1, z_i[v_i] \leftarrow 1$ 
4 In each round:
5   send  $\langle y_i, z_i, d_i^- \rangle$  to all
6   receive  $\langle y_{j_1}, z_{j_1}, d_{j_1}^- \rangle, \dots, \langle y_{j_\ell}, z_{j_\ell}, d_{j_\ell}^- \rangle$  from the in-neighbors ▷  $\ell$  in-neighbors
7   for  $\omega$  appearing in the support of any vector  $y_i, y_{j_1}, \dots, y_{j_\ell}$  do
8     for  $k = 1, \dots, \ell$  do
9       if  $\omega$  is not in the support of  $y_{j_k}$  then
10         $y_{j_k}[\omega] \leftarrow 0, z_{j_k}[\omega] \leftarrow 1$ 
11       if  $\omega$  is not in the support of  $y_i$  then
12         $x_i[\omega] \leftarrow 0, y_i[\omega] \leftarrow 0, z_i[\omega] \leftarrow 1$ 
13     for  $\omega$  in the support of  $y_i$  do
14        $y_i[\omega] \leftarrow \sum_{k=1}^{\ell} y_{j_k}[\omega] / d_{j_k}^-$ 
15        $z_i[\omega] \leftarrow \sum_{k=1}^{\ell} z_{j_k}[\omega] / d_{j_k}^-$ 
16        $x_i[\omega] \leftarrow y_i[\omega] / z_i[\omega]$ 

```

For each value $\omega \in \Omega$, the execution of this algorithm corresponds to one instance of the Push-Sum algorithm initiated by the vertices whose initial value is ω . Since Push-Sum tolerates asynchronous starts, the frequency of the value ω is asymptotically computed in the corresponding variable in the x_i array if ω is the initial value of some agent, or is equal to zero if ω is not initially present in the network.

Observe that the frequency of the value ω in the vector $v \in \Omega^n$ – namely, $v_\omega(\omega)$ – is a rational number in the *finite* set

$$\mathbb{Q}_n = \left\{ \frac{p}{q} \in \mathbb{Q} : p \in \mathbb{N} \wedge q \in \mathbb{N}_{>0} \wedge 0 \leq p \leq q \leq n \right\}.$$

If the agent i knows an upper bound N on the network size, then i can determine the set $\mathbb{Q}_N \supseteq \mathbb{Q}_n$ and, at each round, it can compute the nearest rational number to $x_i[\omega](t)$ in \mathbb{Q}_N . Theorem 5.2 shows that these rational numbers are eventually all equal to $v_\omega(\omega)$. In this way, agent i computes the frequency function in any vector $v \in \Omega^n$ in finite time, and hence the value of $f(v)$, if f is a frequency-based function. Since two different numbers in \mathbb{Q}_N are at distance at least $1/N^2$, we obtain that the stabilization time of the algorithm is in $O(n^{2D} \log N)$.

Let $\mathcal{C}_{\leq N}$ denote the class of *dynamic* networks with at most N vertices and a finite dynamic diameter. Combined with the impossibility result in Corollary 4.2, we obtain the following characterization of computable functions in dynamic anonymous networks with a finite dynamic diameter, when a bound on the network size is known.

COROLLARY 5.3. *Let (X, δ) be a metric space. With outdegree awareness, for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is frequency-based;

(ii) f is δ -computable in every network class $C_{\leq N}$.

Since any function which is computable in a network class closed under graph isomorphisms is necessarily a multiset-based function, we derive the following computability result when the network size is known.

COROLLARY 5.4. *Let (X, δ) be a metric space. With outdegree awareness, for any function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

- (i) f is multiset-based;
- (ii) f is δ -computable in every network class C_n .

We now study how to use the Push-Sum algorithm for computing a function when no bound on the network size is available. For that, we first observe that at each round t , the positive rational numbers $x_i[\omega](t)$ may not correspond to a frequency vector, since their sum may be different from one. This is why each agent i maintains an additional variable \tilde{x}_i , set to

$$\tilde{x}_i[\omega] = \frac{x_i[\omega]}{\sum_{\omega'} x_i[\omega']}$$

to form a frequency function. Then, agent i can easily construct a vector $\langle \tilde{x}_i \rangle$ on Ω that is \tilde{x}_i -frequenced, i.e., whose frequency function is \tilde{x}_i .

This leads us to introduce the notion of δ -continuity in frequency for a frequency-based function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$. If $\nu(1), \nu(2), \dots$ is a sequence of vectors of arbitrary size such that, for every value $\omega \in \Omega$, the sequence of frequencies $\nu_{\nu(1)}[\omega], \nu_{\nu(2)}[\omega], \dots$ converges to some limit value $\nu^*[\omega]$ and those limit values form a frequency function ν^* , then the sequence $f(\nu(1)), f(\nu(2)), \dots$ converges in (X, δ) to $f(\langle \nu^* \rangle)$.

As an example, the average function is continuous in frequency with the classical metric on \mathbb{R} . Other examples are provided by the *threshold frequency predicates* $\Phi_r^\omega : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow \{0, 1\}$, where r is a real number in $[0, 1]$ and $\omega \in \Omega$, defined by

$$\Phi_r^\omega(\nu) := \begin{cases} 1 & \text{if } \nu_\omega \geq r \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, the function Φ_r^ω is continuous in frequency with the discrete metric on $\{0, 1\}$ if and only if r is *irrational*.

Algorithm 1, complemented with the variables \tilde{x}_i , then δ -computes any function f that is δ -continuous in frequency: the variables $f(\langle \tilde{x}_i \rangle)$ all tend to $f(\nu)$ (in the sense of the metric δ).

COROLLARY 5.5. *Let (X, δ) be a metric space. In the communication model with outdegree awareness, every frequency-based function $f : \bigcup_{n \in \mathbb{N}_{>0}} \Omega^n \rightarrow X$ that is δ -continuous in frequency is δ -computable in the class of dynamic networks of finite dynamic diameter.*

5.5 Computing with leaders

If there is a set of $\ell \geq 1$ leaders in the network, with ℓ known to all agents, a slight variant of the Push-Sum algorithm allows each agent i to compute any multiset-based function: its code is unchanged except if it is not a leader in which case its variables $z_i[\omega]$ are initially set to zero instead of one (cf. lines 3, 10, and 12 in Algorithm 1; it is then possible, on line 16, that $x_i[\omega]$ be equal to ∞ , but only for finitely many rounds). The variable $\ell x_i[\omega]$ tends to ω 's multiplicity, which is thus asymptotically computable. The frequency-based condition can thus be replaced by the multiset-based condition in Corollaries 5.3 and 5.5.

6 CONCLUDING REMARKS

In this paper, we have presented a panorama of function computability by anonymous networks, both static and dynamic, for the various communication models that are typically considered when studying message-passing distributed systems.

Three classes of functions stand out: *set-based* functions, whose output is determined by the *set* of values appearing in the input vector; *multiset-based* – or symmetric – functions, which are determined by the multiset of values; and the intermediate class of *frequency-based functions*, whose values may depend on the relative frequency of the input values, but not on their multiplicities.

A fundamental result states that in anonymous networks communicating through a simple *local broadcast* primitive – where an agent has no knowledge or control of the recipients of its messages – only set-based functions can be computed, both in the dynamic and in the static case. This holds even if we assume global symmetry breaking in the form of one or several agents being designated as leaders. Conversely, a simple flooding algorithm easily allows all agents to recover the set of all input values in finite time, and thus to compute any set-based function.

In hope of computing a larger class of functions, we have described three ways of augmenting this simple communication model: we can work under the assumption that communication links are symmetric; we can assume that senders are aware, ahead of emission, of how many other agents will receive each of their messages; or, in the case of a static network, we can assume that an agent may individually address each of its neighbors.

Under all three models for static networks, and under the former two for dynamic networks, the class of computable functions is almost characterized as that of the frequency-based functions, up to some restrictions in the dynamic case, discussed hereafter. Moreover, under any of these models, breaking the symmetry by introducing one or several leaders allows for recovering the full multiset of input values, and thus for computing multiset-based functions, as does providing the agents with the size of the network, but *not* with a bound over the size of the network. For static networks, these results are collected in Table 1, and for dynamic networks they are presented in Table 2.

There are slight limitations to the above picture that we expect further works to address. First, our Push-Sum-based method for computing frequency-based functions only works if a bound over the number of agents is known by the agents, in order to turn an approximate result into an exact one. Otherwise, we must restrict ourselves to functions that are additionally *continuous in frequency*. Can we lift these restrictions and recover the same computability statement for the static and the dynamic case? What exactly characterizes continuity in frequency?

Another consideration, still regarding the dynamic case, concerns the connectivity assumption. The algorithms that we consider have not, in general, been shown to work under the relaxed assumption of a network that, while never becoming permanently split, do not have a finite dynamic diameter – besides from the Metropolis-based family of algorithms. The convergence of the latter, for a symmetric communication model, results from Moreau’s remarkable theorem [29, Theorem 1], which ensures the convergence of a wide family of algorithms. This weaker connectivity assumption is often considered when studying natural systems through a distributed lens [2, 3, 13, 14]. Which of our computability results continue to hold in this case? The recent algorithm designed by Di Luna and Viglietta [25, 26] for the case of symmetric networks, could conceivably work; what, however, can be said in the outdegree awareness model, where Moreau’s theorem does not apply?

Finally, what of self-stabilizing computation? Here again, neither Di Luna and Viglietta’s algorithm, nor Push-Sum, continue to work. Can either of them be fixed? If not, what can be said of self-stabilizing computation over dynamic networks, under the different communication models that we consider?

ACKNOWLEDGMENTS

We would like to thank Jérémie Chalopin, Louis de Monterno, Alex Olshevsky, and Michaël Thomazo for useful discussions related to this work.

REFERENCES

- [1] Dana Angluin. Local and global properties in networks of processors (extended abstract). In Raymond E. Miller, Seymour Ginsburg, Walter A. Burkhard, and Richard J. Lipton, editors, *Proceedings of the 12th Annual ACM Symposium on Theory of Computing, April 28-30, 1980, Los Angeles, California, USA*, pages 82–93. ACM, 1980. URL: <https://doi.org/10.1145/800141.804655>.
- [2] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Comput.*, 18(4):235–253, 2006. URL: <https://doi.org/10.1007/s00446-005-0138-3>.
- [3] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Comput.*, 20(4):279–304, 2007. URL: <https://doi.org/10.1007/s00446-007-0040-2>.
- [4] Florence Bénézit, Vincent D. Blondel, Patrick Thiran, John N. Tsitsiklis, and Martin Vetterli. Weighted gossip: Distributed averaging using non-doubly stochastic matrices. In *IEEE International Symposium on Information Theory, ISIT 2010, June 13-18, 2010, Austin, Texas, USA, Proceedings*, pages 1753–1757. IEEE, 2010. URL: <https://doi.org/10.1109/ISIT.2010.5513273>.
- [5] Paolo Boldi, Shella Shammah, Sebastiano Vigna, Bruno Codenotti, Peter Gemmel, and Janos Simon. Symmetry breaking in anonymous networks: Characterizations. In *Fourth Israel Symposium on Theory of Computing and Systems, ISTCS 1996, Jerusalem, Israel, June 10-12, 1996, Proceedings*, pages 16–26. IEEE Computer Society, 1996.
- [6] Paolo Boldi and Sebastiano Vigna. Computing vector functions on anonymous networks. In Danny Krizanc and Peter Widmayer, editors, *SIROCCO '97, 4th International Colloquium on Structural Information & Communication Complexity, Monte Verita, Ascona, Switzerland, July 24-26, 1997*, pages 201–214. Carleton Scientific, 1997.
- [7] Paolo Boldi and Sebastiano Vigna. Fibrations of graphs. *Discret. Math.*, 243(1-3):21–66, 2002. URL: [https://doi.org/10.1016/S0012-365X\(00\)00455-6](https://doi.org/10.1016/S0012-365X(00)00455-6).
- [8] Paolo Boldi and Sebastiano Vigna. Universal dynamic synchronous self-stabilization. *Distributed Comput.*, 15(3):137–153, 2002. doi:10.1007/s004460100062.
- [9] Ming Cao, A. Stephen Morse, and Brian D. O. Anderson. Reaching a Consensus in a Dynamically Changing Environment: A Graphical Approach. *SIAM Journal on Control and Optimization*, 47(2):575–600, 2008. URL: <http://epubs.siam.org/doi/10.1137/060657005>, doi:10.1137/060657005.
- [10] Bernadette Charron-Bost. Geometric bounds for convergence rates of averaging algorithms. *Inf. Comput.*, 285(Part):104909, 2022. doi:10.1016/j.ic.2022.104909.
- [11] Bernadette Charron-Bost and Patrick Lambein-Monette. Computing outside the box: Average consensus over dynamic networks. In James Aspnes and Othon Michail, editors, *1st Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2022, March 28-30, 2022, Virtual Conference*, volume 221 of *LIPICs*, pages 10:1–10:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.SAND.2022.10.
- [12] Bernadette Charron-Bost and Shlomo Moran. The firing squad problem revisited. *Theor. Comput. Sci.*, 793:100–112, 2019. doi:10.1016/j.tcs.2019.07.023.
- [13] Bernard Chazelle. The Total s-Energy of a Multiagent System. *SIAM Journal on Control and Optimization*, 49(4):1680–1706, 2011. URL: <http://epubs.siam.org/doi/10.1137/100791671>, doi:10.1137/100791671.
- [14] Bernard Chazelle. Natural algorithms and influence systems. *Commun. ACM*, 55(12):101–110, 2012. URL: <https://doi.org/10.1145/2380656.2380679>.
- [15] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *CACM*, 17(11):643–644, 1974.
- [16] Roland L. Dobrushin. Central limit theorem for nonstationary markov chains. II. *Theory of Probability and Its Applications*, 1:329–383, 1956. URL: <https://api.semanticscholar.org/CorpusID:119707237>.
- [17] Roland L. Dobrushin. Coefficients of ergodicity: Structure and applications. *Advances in Applied Probability*, 11(3):576–590, 1956. URL: <https://doi.org/10.2307/1426955>.
- [18] Faith E. Fich and Eric Ruppert. Hundreds of impossibility results for distributed computing. *Distributed Comput.*, 16(2-3):121–163, 2003. URL: <https://doi.org/10.1007/s00446-003-0091-y>.
- [19] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation*, 5(3):1–33, 2002.
- [20] Julien M. Hendrickx, Alexander Olshevsky, and John N. Tsitsiklis. Distributed anonymous discrete function computation. *IEEE Trans. Autom. Control.*, 56(10):2276–2289, 2011. doi:10.1109/TAC.2011.2163874.
- [21] Julien M. Hendrickx and John N. Tsitsiklis. Fundamental limitations for anonymous distributed systems with broadcast communications. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 9–16. IEEE, 2015. URL: <http://ieeexplore.ieee.org/document/7446980>, doi:10.1109/ALLERTON.2015.7446980.
- [22] Nathan Jacobson. *Basic Algebra*, volume 1. Freeman, 1980.
- [23] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 482–491. IEEE Computer Society, 2003. URL: <https://doi.org/10.1109/SFCS.2003.1238221>.

- [24] Patrick Lambein-Monette. *Average consensus in anonymous dynamic networks: An algorithmic approach. (Consensus de moyenne dans les réseaux dynamiques anonymes: Une approche algorithmique)*. PhD thesis, Polytechnic Institute of Paris, France, 2020. URL: <https://tel.archives-ouvertes.fr/tel-03168053>.
- [25] Giuseppe Antonio Di Luna and Giovanni Viglietta. Computing in anonymous dynamic networks is linear. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 1122–1133. IEEE, 2022. doi:10.1109/FOCS54457.2022.00108.
- [26] Giuseppe Antonio Di Luna and Giovanni Viglietta. Optimal computation in leaderless and multi-leader disconnected anonymous dynamic networks. In Rotem Oshman, editor, *37th International Symposium on Distributed Computing, DISC 2023, October 10-12, 2023, L'Aquila, Italy*, volume 281 of *LIPICs*, pages 18:1–18:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.DISC.2023.18.
- [27] Friedemann Mattern. Global quiescence detection based on credit distribution and recovery. *Inform. Processing Letters*, 30(4):195–200, 1989.
- [28] Shlomo Moran and Manfred K. Warmuth. Gap theorems for distributed computation. *SIAM J. Comput.*, 22(2):379–394, 1993. URL: <https://doi.org/10.1137/0222028>.
- [29] Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005. URL: <http://ieeexplore.ieee.org/document/1393134/>, doi:10.1109/TAC.2004.841888.
- [30] Angelia Nedic, Alex Olshevsky, and Michael G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proc. IEEE*, 106(5):953–976, 2018. doi:10.1109/JPROC.2018.2817461.
- [31] Alex Olshevsky. Linear time average consensus and distributed optimization on fixed graphs. *SIAM J. Control. Optim.*, 55(6):3990–4014, 2017. doi:10.1137/16M1076629.
- [32] Masafumi Yamashita and Tiko Kameda. Electing a leader when processor identity numbers are not distinct (extended abstract). In Jean-Claude Bermond and Michel Raynal, editors, *Distributed Algorithms, 3rd International Workshop, Nice, France, September 26-28, 1989, Proceedings*, volume 392 of *Lecture Notes in Computer Science*, pages 303–314. Springer, 1989. URL: https://doi.org/10.1007/3-540-51687-5_52.
- [33] Masafumi Yamashita and Tiko Kameda. Computing functions on asynchronous anonymous networks. *Math. Syst. Theory*, 29(4):331–356, 1996. doi:10.1007/bf01192691.
- [34] Masafumi Yamashita and Tsunehiko Kameda. Computing on anonymous networks: Part I- characterizing the solvable cases. *IEEE Trans. Parallel Distributed Syst.*, 7(1):69–89, 1996. URL: <https://doi.org/10.1109/71.481599>.