



HAL
open science

Modèles de communication et calculabilité dans les réseaux anonymes

Bernadette Charron-Bost, Patrick Lambein-Monette

► **To cite this version:**

Bernadette Charron-Bost, Patrick Lambein-Monette. Modèles de communication et calculabilité dans les réseaux anonymes. 27èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (Algotel-Cores25), Alain Cournier; Stéphane Devismes; Geoffrey Defalque; Yoann Dieudonné; Franck Petit; Léo Robert, Jun 2025, Saint Valery-sur-Somme, France. <hal-05029829v2>

HAL Id: hal-05029829

<https://hal.science/hal-05029829v2>

Submitted on 22 Apr 2025

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.



HAL Authorization

Modèles de communication et calculabilité dans les réseaux anonymes

Bernadette Charron-Bost¹ and Patrick Lambein-Monette¹

¹ CNRS, DI ENS, École Normale Supérieure, F-75230 Paris, France

Le problème du calcul distribué de la fonction f consiste, pour un réseau de n agents autonomes, à ce que chaque agent calcule la valeur $f(\omega_1, \dots, \omega_n)$ lorsque $\omega_1, \dots, \omega_n$ sont les valeurs d'entrée des agents, chacune étant initialement privée à un agent. Dans cet article, nous étudions cette question dans les réseaux synchrones anonymes où la communication se fait par échanges de messages, avec différentes hypothèses sur le degré de connaissance ou de contrôle que chaque agent a sur ses voisins. Dans le modèle de diffusion "à l'aveugle" où, à chaque tour, un agent se contente d'envoyer un message unique sans aucune connaissance ni contrôle sur ses destinataires, Boldi et Vigna ont prouvé que les fonctions calculables sont exactement celles dont la valeur ne dépend que de l'ensemble $\{\omega_1, \dots, \omega_n\}$ des arguments. Ce résultat est encore valable même s'il existe un leader ou si les agents connaissent la taille du réseau. Ici, nous examinons trois modèles de communication, plus forts mais tout aussi naturels, régis par les principes suivants : (1) à chaque tour, les agents connaissent le nombre de leurs voisins sortants ; (2) tous les liens de communication dans le réseau sont bidirectionnels ; ou (3) les agents peuvent s'adresser individuellement à chacun de leurs voisins sortants. Nous montrons que, dans les réseaux statiques et pour chacun de ces trois modèles de communication, la classe des fonctions calculables est exactement celle composée des fonctions dont la valeur ne dépend que de la fréquence relative de chaque valeur d'entrée ω_i . Si, en outre, le réseau possède un leader ou si les agents connaissent la taille du réseau, chaque agent peut récupérer en fait le multi-ensemble complet $\{\{\omega_1, \dots, \omega_n\}\}$ et peut donc calculer n'importe quelle fonction invariante sous permutation de ses arguments. Autrement dit, une connaissance, même minimale, sur les destinataires des messages augmente considérablement la calculabilité en distribué. Dans une dernière partie, nous étudions comment étendre les résultats obtenus pour les réseaux statiques aux réseaux dynamiques.

Keywords: anonymous networks, synchronous networks, communication model, function computability

1 Introduction

Function computation is a fundamental problem of distributed computing in networks. With unique identifiers, the question of *which* functions can be computed is essentially trivial: by running a straightforward flooding algorithm, any function can be computed, as long as it is invariant under permutation of its arguments. However, in anonymous networks of deterministic agents, function computability becomes a non-trivial question of much importance, as function computation is an extremely commonplace building block of many distributed applications. Here, our objective is to understand the fundamental limitations and capabilities for function computation by synchronous networked systems of anonymous, deterministic agents, in either static or dynamic networks. More precisely, we study the question of which functions are computable with varying assumptions in the degree of awareness or control each agent has over its outneighbors. While Tables 1 and 2 give a high-level summary of our argument, we refer the reader to our preprint [CL23] for precise definitions of our models, methods, results, as well as their proofs.

We define three classes of functions of arbitrary arity, i.e., of type $f : \bigcup_{n \geq 1} \Omega^n \rightarrow X$. First, *set-based* functions, whose value only depend on the set of its input values. Second, *frequency-based* functions, which can depend on the *relative frequencies* of the input. Third, *multiset-based* functions, which can depend on the multiset of the input. A prototypical example of a set-based function is the maximum, an example of a frequency-based function is the average, and an example of a multiset-based function is the sum. As these examples illustrate, the inclusion between the three classes is strict.

In terms of communication, we consider four distinct models. In the base case of *blind broadcast*, in a round t , an agent i prepares and sends a message $m_i(t)$ which is delivered to each of i 's round t outneighbors. This message is entirely determined by the current state of the agent. The second model, *outdegree awareness*, only differs in one way: the message $m_i(t)$ may now also depend on agent i 's outdegree in the round t communication graph. Transmissions in these two models are *isotropic*, in the sense that, in any given round, any two messages from the same sender are the same. In contrast, in our third model of *output port awareness*, each agent i has access to different ports, each linking to a different outneighbor, and sends one message per port in each round. Finally, our fourth model of *symmetric communications* corresponds to the restriction of the blind broadcast model, to the class of networks with bidirectional links.

We then adopt the classical notions of algorithm in an anonymous network and of synchronous execution, for each of these communication models (see [CL23]). We say that the algorithm \mathcal{A} (*exactly*) *computes* a function f in the network class \mathcal{C} if, in every execution of \mathcal{A} in a network in \mathcal{C} composed of n agents and with the input values v_1, \dots, v_n , all agents eventually adopt the value $f(v_1, \dots, v_n)$. The function f is said to be *computable in the network class \mathcal{C}* if there exists an algorithm that computes f in \mathcal{C} .

2 Computability in static networks

Our main result for static networks is an exact characterization of computable functions in the class of strongly connected networks when assuming either output port awareness, symmetric communication links, or outdegree awareness. Here, our approach for both positive and negative results exploits the notion of *graph fibration* in homotopy graph theory [Ang80, YK89, BSV⁺96].

Theorem 1 *In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbf{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

1. f is frequency-based;
2. f is computable in the class of static strongly connected networks.

The impossibility result relies on the fact that all agents in the same *fib*re, i.e., with similar in-neighborhood, have the same behavior if they all start in the same state. For the positive results, we provide a unifying explanation why frequency-based functions are computable with output port awareness, symmetric communications, or outdegree awareness: we show that in each case, agents can compute the cardinalities of the fibres up to some common factor by solving a homogeneous system, which we prove to be always of rank one. While the latter property clearly holds with 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 novel algebraic result *à la* Perron-Frobenius for matrices whose diagonal entries may be negative. In each communication model, the result is a finite-state and linear-time algorithm for computing a frequency-based function in a static network.

With a refinement of the proof of the impossibility result in Theorem 1, we obtain the same computability result when a bound on the number of agents is known.

Corollary 1 *In any of the three communication models with either output port awareness, symmetric communications, or outdegree awareness, and for any function $f : \bigcup_{n \in \mathbf{N}_{>0}} \Omega^n \rightarrow X$, the following assertions are equivalent:*

1. f is frequency-based;
2. for every positive integer N , f is computable in the class of strongly connected networks with at most N nodes.

When the exact size of the network is known, our approach allows for computing any *multiset-based* function – i.e., any function invariant under permutation.

Corollary 2 *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:*

1. f is multiset-based;
2. for every positive integer n , f is computable in the class of strongly connected networks with exactly n nodes.

In the case of a network with one leader, its fibre is reduced to itself. Hence, our method actually allows for finding exactly the cardinality of each fibre, and so computes any multiset-based function.

Corollary 3 *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:*

1. f is multiset-based;
2. f is computable in the class of static strongly connected networks with one leader.

Therefore, each of the two above assumptions, namely the knowledge of the network size and the existence of a leader, considerably increases the computational power in the case of output port awareness, symmetric communications, or outdegree awareness, while they leave it unchanged in the blind broadcast model [BV97].

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

	BLIND BROADCAST	OUTDEGREE AWARENESS	SYMMETRIC COMMUNICATIONS	OUTPUT PORT AWARENESS
no centralized help	set-based Hendrickx <i>et al.</i> [HOT11]	frequency-based Theorem 1	frequency-based Theorem 1	frequency-based Theorem 1
a bound over n is known	set-based Boldi & Vigna [BV97]	frequency-based Corollary 1	frequency-based Corollary 1	frequency-based Corollary 1
n is known	set-based Boldi & Vigna [BV97]	multiset-based Corollary 2	multiset-based Corollary 2	multiset-based Corollary 2
one leader	set-based Boldi & Vigna [BV97]	multiset-based Corollary 3	multiset-based Corollary 3	multiset-based Corollary 3

3 Computability in dynamic networks

We now investigate function computability in dynamic networks, and confine ourselves to the models with symmetric links and with outdegree awareness, since output port awareness encompasses the latter. The property of strong connectivity must then be replaced by that of having a finite *dynamic diameter*, that is to say there exists an integer D such that, for every pair of vertices i, j and from every round, there is a *temporal path* of length at most D connecting i to j .

Obviously, the necessary condition in Theorem 1 continues to hold, i.e., only frequency-based functions are computable. For the sufficient condition, we propose a totally different approach to that developed for static networks, which is inspired by averaging methods in statistical physics: our method for dynamic graphs consists in directly computing the relative frequencies of the input values.

Various algorithms have been proposed in distributed control for the computation of the average according to the communication model, including the *MaxMetropolis* algorithm [CL22] for symmetric links and the *Push-Sum* algorithm [KDG03] for outdegree awareness. Depending on the communication model, we choose one of these algorithms and transform it for computing the relative frequency of each input

value, thanks to appropriate initializations (see [CL23] for a complete description). In all cases, our method requires infinite-precision values, but uses no persistent memory in the sense that state-update rules correspond to (discrete) differential equations of order one.

When agents know an upper bound on the network size, we obtain the same characterization of computable functions as in the static case.

Corollary 4 *With 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:*

1. f is frequency-based;
2. for every positive integer N , f is computable in the class of networks with at most N nodes.

When no bound on the network size is known, our method allows for asymptotically computing the frequency-based functions that satisfy some continuity property with respect to the relative frequencies of their arguments. We have thus proven that, in dynamic networks, the frequency-based condition which is necessary for approximate computability, is also nearly sufficient in the sense that it simply needs to be enriched with continuity in frequency. Our computability results are summarized in Table 2.

Our method for dynamic networks has to be compared with the linear-time algorithm proposed by Di Luna and Viglietta [LV23] which allows for an exact computation of frequency-based functions in the case of symmetric communications. Unfortunately, this algorithm does not fit our model since it does not work with asymmetric links, and does not tolerate asynchronous starts. Moreover, as it relies on the construction of an infinitely growing history tree, the algorithm requires both an infinite number of states and an infinite bandwidth, even in the case a bound on the network size is available.

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

	BLIND BROADCAST	OUTDEGREE AWARENESS	SYMMETRIC COMMUNICATIONS
no centralized help	set-based Hendrickx <i>et al.</i> [HOT11]	?	?
a bound over n is known	set-based Hendrickx <i>et al.</i> [HOT11]	frequency-based Corollary 4	frequency-based Charron-Bost <i>et al.</i> [CL22]
n is known	set-based Hendrickx <i>et al.</i> [HOT11]	multiset-based Corollary 4	multiset-based Charron-Bost <i>et al.</i> [CL22]
one leader	set-based Hendrickx <i>et al.</i> [HOT11]	?	?

References

- [Ang80] 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.
- [BSV⁺96] 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.
- [BV97] 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.

- [CL22] 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*, volume 221 of *LIPICs*, pages 10:1–10:16, 2022.
- [CL23] Bernadette Charron-Bost and Patrick Lambein-Monette. Know your audience. *CoRR*, abs/2311.17728, 2023.
- [HOT11] Julien M. Hendrickx, Alexander Olshevsky, and John N. Tsitsiklis. Distributed anonymous discrete function computation. *IEEE Trans. Autom. Control.*, 56(10):2276–2289, 2011.
- [KDG03] 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.
- [LV23] 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.
- [YK89] 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.