


# Brief Announcement: Anonymous Distributed Localisation via Spatial Population Protocols

Leszek Gąsieniec ✉ 🏠 


Department of Computer Science, University of Liverpool, UK

Łukasz Kuszner ✉ 🏠 

Faculty of Mathematics, Physics and Informatics, University of Gdańsk, Poland

Ehsan Latif ✉ 🏠 

School of Computing and Department of Mathematics, Science, and Social Studies Education,  
University of Georgia, Athens, GA, USA

Ramviyas Parasuraman ✉ 🏠 

School of Computing, University of Georgia, Athens, GA, USA

Paul Spirakis ✉ 🏠 

Department of Computer Science, University of Liverpool, UK

Grzegorz Stachowiak ✉ 🏠 

Institute of Computer Science, University of Wrocław, Poland

---

## Abstract

In the *distributed localization problem* (DLP),  $n$  anonymous robots (agents)  $A_0, \dots, A_{n-1}$  begin at arbitrary positions  $p_0, \dots, p_{n-1} \in S$ , where  $S$  is a Euclidean space. Initially, each agent  $A_i$  operates within its own coordinate system in  $S$ , which may be inconsistent with those of other agents. The primary goal in DLP is for agents to reach a consensus on a unified coordinate system that accurately reflects the relative positions of all points,  $p_0, \dots, p_{n-1}$ , in  $S$ . Extensive research on DLP has primarily focused on the feasibility and complexity of achieving consensus when agents have limited access to inter-agent distances, often due to missing or imprecise data. In this paper, however, we examine a minimalist, computationally efficient model of distributed computing in which agents have access to all pairwise distances, if needed. Specifically, we introduce a novel variant of population protocols, referred to as the *spatial population protocols* model. In this variant each agent can memorise one or a fixed number of coordinates, and when agents  $A_i$  and  $A_j$  interact, they can not only exchange their current knowledge but also either determine the distance  $d_{ij}$  between them in  $S$  (distance query model) or obtain the vector  $\vec{v}_{ij}$  spanning points  $p_i$  and  $p_j$  (vector query model). We present here a *leader-based localisation protocol with distance queries*.

**2012 ACM Subject Classification** Theory of computation → Distributed computing models

**Keywords and phrases** Population Protocols, Distributed Localisation

**Digital Object Identifier** 10.4230/LIPIcs.SAND.2025.19

**Related Version** *Version with proofs and more results:* <https://arxiv.org/abs/2411.08434> [4]

**Funding** *Paul Spirakis:* Supported by the EPSRC grant EP/P02002X/1.

*Grzegorz Stachowiak:* Supported by the NCN grant 2020/39/B/ST6/03288.

## 1 Introduction

Location services are vital for modern computing paradigms like pervasive computing and sensor networks. Manual configuration and GPS, though effective for determining node locations, are impractical in large-scale or obstructed environments. Recent network localisation approaches use beacon nodes with known positions to help other nodes estimate their



© Leszek Gąsieniec, Łukasz Kuszner, Ehsan Latif, Ramviyas Parasuraman, Paul Spirakis, and Grzegorz Stachowiak;

licensed under Creative Commons License CC-BY 4.0

4th Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2025).

Editors: Kitty Meeks and Christian Scheideler; Article No. 19; pp. 19:1–19:5



Leibniz International Proceedings in Informatics

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

locations via distance measurements. Key challenges include ensuring unique localisability, managing computational complexity, and addressing deployment issues. In the *distributed localisation problem* (DLP),  $n$  anonymous robots (agents)  $A_0, \dots, A_{n-1}$  begin at arbitrary positions  $p_0, \dots, p_{n-1} \in S$ , where  $S$  is an Euclidean space. Initially, each agent  $A_i$  operates within its own coordinate system in  $S$ , which may be inconsistent with those of other agents. The primary goal in DLP is for agents to reach a consensus on a unified coordinate system that accurately reflects the relative positions of all points,  $p_0, \dots, p_{n-1}$ , in  $S$ .

A network of agents' unique localisability is determined by specific combinatorial properties of its graph and the number of *anchors* (agents aware of their real location). For example, *graph rigidity theory* [2, 5, 6] provides a necessary and sufficient condition for unique localisability [2]. Specifically, a network of agents located in the plane is uniquely localisable if and only if it has at least three anchors and the network graph is *globally rigid*. However, unless a network is dense and regular, global rigidity is unlikely. Even without global rigidity, large portions of a network may still be globally rigid, though the positions of remaining nodes will remain indeterminate due to multiple feasible solutions. The decision version of this problem, often referred to as the *graph embedding* or *graph realisation* problem, requires determining whether a weighted graph can be embedded in the plane so that distances between adjacent vertices match the edge weights, a problem known to be strongly NP-hard [14]. Furthermore, this complexity holds even when the graph is globally rigid [2]. In *sensor networks*, where nodes measure distances only within a communication range  $r$ , the network is best represented as a *unit disk graph*. Here, two nodes are adjacent if and only if their distance is  $\leq r$ . The corresponding decision problem, known as *unit disk graph reconstruction*, requires determining whether a graph can be embedded in the plane such that distances between adjacent nodes match edge weights, while distances between non-adjacent nodes exceed  $r$ . This problem is also NP-hard [1], indicating that no efficient algorithm can solve localisation in the worst case unless  $P = NP$ . Furthermore, even for instances with unique reconstructions, no efficient randomised algorithm exists to solve this problem unless  $RP = NP$  [1].

Distributed localisation is also crucial in *robotic systems*, enabling robots to autonomously determine their spatial position within an environment – a fundamental requirement for applications such as navigation, mapping, and multi-robot coordination [18]. Accurate localisation allows robots to interact more effectively with their surroundings and with each other, facilitating tasks from autonomous driving to warehouse automation and search-and-rescue operations [10, 12]. Localisation approaches generally fall into two broad categories: centralised and distributed systems [19]. Centralised localisation systems offer high accuracy but may struggle with scalability and robustness [9]. In contrast, distributed localisation systems allow each robot to perform localisation computations independently or in collaboration with neighbouring robots, enhancing adaptability and resilience, although this may come at the cost of increased complexity [7, 8]. Within distributed systems, leader-based localisation mechanisms involve one or more designated robots that serve as reference points or coordinators for localisation [3], which can streamline computations but may create single points of failure. Leaderless localisation, where all robots contribute equally to position estimation without relying on specific leader nodes, is advantageous in decentralised applications where flexibility and fault tolerance are paramount [7, 15]. Both methods have been explored using probabilistic [17], geometric [11], and graph-based models [7], with leaderless approaches gaining traction due to their robustness in large-scale and dynamically changing settings. Various methods leverage tools such as Kalman filters [13], particle filters [9], and graph rigidity theory [16] to enhance localisation accuracy and efficiency in complex environments.

## 1.1 Spatial population protocols

In this paper, we explore a minimalist, computationally efficient model of distributed computing, where agents have probabilistic access to pairwise distances. Our focus is on achieving *anonymity* while maintaining high time efficiency and minimal use of network resources, including limited local storage (agent state space) and communication. To meet these goals, we introduce a new variant of population protocols, referred to as the *spatial population protocols* model, specified later in this section.

While population protocols provide an elegant and resilient framework for randomised distributed computation, they lack spatial embedding. To address this limitation, we introduce a new *spatial* variant of population protocols that extends the transition function to include basic geometric queries. In particular, in this model each agent can memorise one or a fixed number of coordinates, and during an interaction of two agents  $A_i$  and  $A_j$ , in addition to exchange of their current knowledge, the agents can determine:

- (1) the distance  $d_{ij}$  separating them in  $S$ , in *distance query model*, and
- (2) vector  $\vec{v}_{ij}$  spanning points  $p_i$  and  $p_j$ , in *vector query model*.

## 2 Leader-based localisation in distance query model

In this section, we discuss a localisation protocol with predefined input stabilising in  $o(n)$  time, i.e., after this time labels of all agents become stable. This protocol is non-self-stabilising. We assume that one of the agents starts as the *leader* of the population. If the identity of the leader is not known, the localisation protocol can be preceded by one of the leader election protocols discussed in the introduction. The agents' positions  $p_0, \dots, p_{n-1}$  are distributed in a  $k$ -dimensional Euclidean space  $S$ , where  $k$  is a fixed integer. It is assumed that any  $k + 1$  agents' positions span the entire space. For example, in two-dimensional space, this assumption guarantees that no three points are collinear. Although our algorithms can be adapted to handle an arbitrary distribution of agents' positions, the time guarantees of such adaptations would be weaker. The state of an agent can accommodate a fixed number of agent positions and distances.

We adopt a symmetric model of communication, which means that when agents  $A_u$  and  $A_v$  interact, they both gain access to each other's states as well as the distance  $d_{uv}$ . The transitions assigned to the leader are distinct from those of the other agents, and the leader also serves as the initiator of the entire process. Initially, the state of each agent  $A_u$  stores a label  $x_u$  (representing a hypothetical position in  $S$ ) and its color  $C[A_u]$ . We assume that at the beginning, the leader is coloured **green**, and each non-leader's colour is set to **blue**. Finally, the leader's label (position in  $S$ ) is set at the origin of the coordinate system, i.e., this label is used as the anchor in the localisation process.

### 2.1 Localisation via multi-contact epidemic

The localisation protocol presented in this section consists of two parts (see the formal description of Algorithm 1). In the first part of the protocol, the labels of  $k + 1$  agents (including the leader) become stable (positions of these agents become fixed), and these agents become **green**. The counter  $i$ , initially set to 1, of agents with stable (**green**) labels is maintained by the leader. In the second part the labels of all remaining (**blue**) agents become stable. And this is done by contacting with  $k + 1$  different **green** agents. And once the agent is positioned, it becomes **green**. We refer to this (*multilateration*) process as  $(k + 1)$ -contact epidemic.

Positioning of each of the first  $k$  agents has to be approved by the leader. More precisely, after an aspiring to be **green** agent  $A_v$  interaction with all  $i < k$  **green** agents is concluded, to become **green**  $A_v$  must meet the leader to get approved. And when this happens, the leader updates the counter of **green** agents, and the new **green** agent  $A_v$  is ready to calculate its projection onto the subspace spanned by its  $i$  **green** predecessors and the leader, as well as its Euclidean distance from this subspace. Namely, the first  $i$  coordinates of  $A_v$ 's label are determined by this projection, and the  $(i + 1)$ -th coordinate (in the newly formed dimension) is equal to its distance from the aforementioned subspace. When positioning the remaining agents, we use the fact that interactions with  $k + 1$  green agents allow for the unambiguous determination of an agent's position.

■ **Algorithm 1** Positioning in  $k$  dimensions.

---

```

begin
   $i = 0$ ;  $C[\text{leader}] = \text{green}$ ;
  while  $i < k$  do
    during an interaction between agents  $A_u$  and  $A_v$ ;
    if  $(A_u \text{ is leader}) \wedge (C[A_v] = \text{blue})$  then
      if  $A_v$  interacted with all  $i$  green agents then
        position  $A_v$  and stabilise  $h_v$ ;
         $C[A_v] = \text{green}$ ;
         $i = i + 1$ 
      end
    end
  end
  while not all agents are positioned do
    during an interaction between agents  $A_u$  and  $A_v$ ;
    if  $(C[A_u] = \text{green}) \wedge (C[A_v] = \text{blue})$  then
      if  $A_v$  interacted with  $k + 1$  green agents then
        position  $A_v$  and stabilise  $h_v$ ;
         $C[A_v] = \text{green}$ 
      end
    end
  end
end

```

---

► **Theorem 1.** *Algorithm 1 stabilises labels of all agents in  $k$  dimensions in parallel time  $O(n(\log n/n)^{1/(k+1)})$ .*

### 3 Concluding remarks

In this paper, we introduce a novel variant of spatial population protocols and explore its applicability to the distributed localization problem. In the full version of this paper [4], we explore faster leader-based localisation in one-dimensional space using the model with distance queries, alongside superfast localisation in the model with vector queries. Any meaningful advances in this problem could pave the way for developing faster and more robust lightweight communication protocols suitable for real-world applications. It could also provide insights into the limitations of what can be achieved in such systems.

## References

- 1 J. Aspnes, D.K. Goldenberg, and Y.R. Yang. On the computational complexity of sensor network localization. In *Algorithmic Aspects of Wireless Sensor Networks: First International Workshop, ALGOSENSORS 2004, Turku, Finland, July 16, 2004. Proceedings*, volume 3121 of *Lecture Notes in Computer Science*, pages 32–44. Springer, 2004. doi:10.1007/978-3-540-27820-7\_5.
- 2 T. Eren, D.K. Goldenberg, W. Whiteley, Y.R. Yang, A.S. Morse, B.D.O. Anderson, and P.N. Belhumeur. Rigidity, computation, and randomization in network localization. In *Proceedings IEEE INFOCOM 2004, The 23rd Annual Joint Conference of the IEEE Computer and Communications Societies, Hong Kong, China, March 7-11, 2004*, pages 2673–2684. IEEE, 2004. doi:10.1109/INFCOM.2004.1354686.
- 3 R. Fareh, M. Baziyad, S. Khadraoui, B. Brahmi, and M. Bettayeb. Logarithmic potential field: A new leader-follower robotic control mechanism to enhance the execution speed and safety attributes. *IEEE Access*, 2023.
- 4 L. Gąsieniec, Ł. Kuszner, E. Latif, R. Parasuraman, P. Spirakis, and G. Stachowiak. Anonymous distributed localisation via spatial population protocols, 2024. arXiv:2411.08434.
- 5 B. Hendrickson. Conditions for unique graph realizations. *SIAM Journal on Computing*, 21(1):65–84, 1992. doi:10.1137/0221008.
- 6 B. Jackson and T. Jordán. Connected rigidity matroids and unique realizations of graphs. *J. Comb. Theory B*, 94(1):1–29, 2005. doi:10.1016/J.JCTB.2004.11.002.
- 7 E. Latif and R. Parasuraman. Dgorl: Distributed graph optimization based relative localization of multi-robot systems. In *International Symposium on Distributed Autonomous Robotic Systems*, pages 243–256. Springer, 2022.
- 8 E. Latif and R. Parasuraman. Multi-robot synergistic localization in dynamic environments. In *ISR Europe 2022; 54th International Symposium on Robotics*, pages 1–8. VDE, 2022.
- 9 E. Latif and R. Parasuraman. Instantaneous wireless robotic node localization using collaborative direction of arrival. *IEEE Internet of Things Journal*, 2023.
- 10 E. Latif and R. Parasuraman. Seal: Simultaneous exploration and localization for multi-robot systems. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5358–5365. IEEE, 2023.
- 11 E. Latif and R. Parasuraman. GPRL: Gaussian processes-based relative localization for multi-robot systems. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024.
- 12 M. Misaros, S. Ovidiu-Petru, D. Ionut-Catalin, and M. Liviu-Cristian. Autonomous robots for services—state of the art, challenges, and research areas. *Sensors*, 23(10):4962, 2023. doi:10.3390/S23104962.
- 13 A.P. Moreira, P. Costa, and J. Lima. New approach for beacons based mobile robot localization using Kalman filters. *Procedia Manufacturing*, 51:512–519, 2020.
- 14 J.B. Saxe. Embeddability of weighted graphs in k-space is strongly NP-hard. *17th Allerton Conf. Commun. Control Comput.*, 1979, pages 480–489, 1979.
- 15 S. Wang, Y. Wang, D. Li, and Q. Zhao. Distributed relative localization algorithms for multi-robot networks: A survey. *Sensors*, 23(5):2399, 2023. doi:10.3390/S23052399.
- 16 J. X. Guo, Hu, J. Chen, F. Deng, and T.L. Lam. Semantic histogram based graph matching for real-time multi-robot global localization in large scale environment. *IEEE Robotics and Automation Letters*, 6(4):8349–8356, 2021. doi:10.1109/LRA.2021.3058935.
- 17 M. Xu, N. Snderhauf, and M. Milford. Probabilistic visual place recognition for hierarchical localization. *IEEE Robotics and Automation Letters*, 6(2):311–318, 2020.
- 18 Y. Yue and D. Wang. *Collaborative Perception, Localization and Mapping for Autonomous Systems*, volume 2. Springer Nature, 2020.
- 19 F. Zafari, A. Gkelias, and K.K. Leung. A survey of indoor localization systems and technologies. *IEEE Communications Surveys & Tutorials*, 21(3):2568–2599, 2019. doi:10.1109/COMST.2019.2911558.