

Brief Announcement: The Virtue of Self-Consistency

Fabian Frei ✉ 

CSAIL, Massachusetts Institute of Technology, Cambridge, MA, USA

Koichi Wada ✉ 

Hosei University, Koganei, Tokyo, Japan

Abstract

We show that self-consistency can be a crucial property for autonomous mobile robots.

Specifically, we consider the task of gathering three robots, placed adversarially in distinct locations in the Euclidean plane, in a single point. We assume the natural scheduler ROUNDROBIN, which activates the robots in turns. An activated robot perceives all robot locations in an adversarially scaled, rotated, and mirrored Cartesian coordinate system with itself at the origin and then moves wherever it wants. We show that this task cannot be solved in the default robot model (without any consistency guarantees and no multiplicity detection) but becomes feasible if we assume self-consistency (i.e., no changes between the different activations of the same robot) of either the unit length (i.e., no scaling) or the compass (i.e., no rotating) by providing explicit algorithms.

2012 ACM Subject Classification Computer systems organization → Robotic autonomy

Keywords and phrases Autonomous Mobile Robots, Distinct Gathering, Round Robin, Disorientation, Self-Consistency

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.57

Funding *Fabian Frei*: This research was funded in part by grant number 218001 awarded to Fabian Frei by the Swiss National Science Foundation.

Koichi Wada: Supported by JSPS KAKENHI No. 25K03079.

Acknowledgements We thank the reviewers for their helpful comments and suggestions. The work by Fabian Frei was done in part while at CISPA Helmholtz Center for Information Security, Saarbrücken.

1 Model and Motivation

The computational power of autonomously acting, simple, mobile robots has been the object of intense research in the field of distributed computing. Ever since Suzuki and Yamashita's seminal work [13], a large amount of work has been dedicated to the research of theoretical models of such autonomous mobile robots [1, 2, 3, 6, 10, 11, 12]. The standard textbook by Flocchini et al. [7] provides an excellent overview of the basic models and literature. In the standard model, each robot is represented as a point in the two-dimensional plane with limited capabilities. In particular, the robots are assumed to be *oblivious* (have no persistent memory), *anonymous* (have no IDs), *uniform* (run identical algorithms) [7]. This basic model is often called OBLLOT, which is an abbreviation of “*oblivious robot*.”

The robots operate in *Look-Compute-Move* (*LCM*) cycles. A robot performs such a cycle whenever it is activated by the so-called scheduler. For the scheduler ROUNDROBIN, which we consider here, the robots are activated in turns, keeping the same, previously unknown, order forever. Whenever a robot is activated, it performs a full *LCM* cycle before the next robot is activated. Each cycle is composed of the following three successive operations: In the *Look* operation, a robot obtains a snapshot of the plane showing the positions of the other robots without multiplicity in a Cartesian coordinate system with the



© Fabian Frei and Koichi Wada;

licensed under Creative Commons License CC-BY 4.0

39th International Symposium on Distributed Computing (DISC 2025).

Editor: Dariusz R. Kowalski; Article No. 57; pp. 57:1–57:7

Leibniz International Proceedings in Informatics



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

observing robot at the origin. In the *Compute* operation, it then executes its algorithm, which is identical for all robots, using the snapshot as input. Finally, it moves towards the computed destination in the *Move* operation.

We assume the robots to be *transparent*, i.e., not obstructing each other's views, and *collision-resistant*, i.e., passing each other without consequence.

By default, the choice of the Cartesian coordinate system is fully volatile; that is, any robot's coordinate system may be rotated, scaled, and mirrored adversarially before each activation. If the adversary cannot scale and rotate, we speak of a *consistent unit length* and *consistent compass*, respectively. If the adversary can scale and rotate between different robots but not between different activations of the same robot, we say that the unit length and compass are *self-consistent*, respectively.

We also distinguish between the standard model of *rigid movement* and the weaker assumption of *nonrigid movement*, which models various adverse situations such as robots overheating or running out of energy during their movements. The standard assumption for mobile robots executing Look-Compute-Move cycles is that they always reach their target. This is called *rigid movement*, as opposed to *nonrigid movement*, where a robot's movement may be stopped at any time by the adversary, as long as the robot has moved some minimum distance δ during its current cycle. The minimum movement distance δ is *unknown* to the robots. Note that the knowledge of a common δ would indeed provide the robots with a consistent unit length.

We consider the fundamental and extensively studied task of gathering all robots in a single, arbitrary point. Together with our assumption that the robots start from distinct locations, this task is also called *Distinct Gathering*, in contrast to *Self-Stabilizing Gathering* without this assumption. A recent survey summarizes results for both Distinct Gathering and Self-Stabilizing Gathering in the Euclidean plane [9].

It has been a widespread and natural assumption of the field that the purely robot-internal property of self-consistency cannot help them coordinate. We prove that this is not the case by showing that Distinct Gathering of three robots under ROUNDROBIN is an unsolvable problem in the default model without consistency guarantees but feasible as soon as we assume either a self-consistent unit or a self-consistent compass. This shows that it is important to be careful and precise in model descriptions when it comes to assumptions about self-consistency or the lack thereof.

2 Related Work and Contribution

A quarter century ago, Suzuki and Yamashita [13] introduced in their seminal paper the robot model that served as the basis for much of the research on mobile robots during the following years. The key difference to what we present here is their assumption that robots can detect how many robots occupy any location. They show how gathering can be achieved under this assumption [13, Thm. 3.4]. The main idea is to use the geometry of the initial configuration to create an arbitrary point of multiplicity, and then let all robots move to the point of highest multiplicity. But depending on the given sensors, the robots might not be able to detect multiplicity. Moreover, multiplicity makes it rather simple to achieve a gathering, such that most problems in this settings were resolved quickly. It has thus become the standard assumption of the subsequent papers that robots cannot detect multiplicity. As previously stated, this is also our assumption here.

The remaining question is therefore whether the task remains feasible without multiplicity detection. While gathering $n = 2$ robots is trivial under ROUNDROBIN, it remained a prominent open question whether Distinct Gathering is still possible for $n \geq 3$ robots. This

question was first officially raised at DISC 2006 [4], tackled repeatedly [5, 14], and finally resolved last year, as announced at DISC 2024 [8]: The problem is indeed infeasible for $n = 3$ robots, which matches the established conjecture by Défago et al. [5, Conj. 2], but surprisingly becomes solvable again for $n \geq 4$ robots.

For the purposes of this paper, however, the most interesting result is indeed the infeasibility of Distinct Gathering under ROUNDRBIN for $n = 3$ robots in the default model without any additional assumptions introduced above.

The last and best attempt to find an algorithm solving the problem by introducing a minimal set of assumptions is due to Terai et al. [14, Sect. 4, Alg. 10]. They managed to find an algorithm for the problem if the algorithm is equipped with some internal memory (referred to as internal lights) and at the same time a consistent unit distance. The results announced in this paper show that the aforementioned result still made several unnecessary assumptions. Specifically, we can show that solving the problem requires neither internal nor external memory nor a consistent unit distance.

For the first of our two results, which is formally stated in Theorem 1, we provide an algorithm (Algorithm 1) that works with only the additional assumption that the robots have a self-consistent unit length. In particular, each robot may be adversarially assigned its own unit distance for any given initial configuration. As mentioned above, it is quite surprising that a purely self-consistent property can help the robots coordinate among each other at all. The conventional assumption has been that some shared knowledge or perception is required for the robots to improve their capabilities.

Motivated by the first result, we have also investigated the case where not the unit length, but instead a sense of direction, which we call a compass (essentially a vector of variable length always pointing in the same direction), is given to the robots in a self-consistent way. As it turns out, the problem is even solvable under this assumption, without any self-consistency for the unit length. This is stated in Theorem 2. Due to the space constraints, we cannot provide the full algorithm in this case and can only sketch the main idea. The techniques for the two proofs of Theorems 1 and 2 are quite dissimilar. Indeed, it seems impossible to simulate either of the two algorithms using the other one.

► **Theorem 1.** *Distinct Gathering under ROUNDRBIN is possible for 3 robots with a self-consistent unit length and the otherwise default assumptions detailed in the introduction, even if movement may be nonrigid.*

Proof Sketch. The algorithm (Algorithm 1) relies on a distance classification scheme. Specifically, the continuum of positive real numbers $\mathbb{R}_{>0}$ is partitioned into two disjoint infinite unions of intervals:

$$\begin{aligned} \blacksquare \mathbb{D}_{\text{golden}} &= \bigcup_{k \in 2\mathbb{Z}} [\varphi^{k+1}, \varphi^k), \\ \blacksquare \mathbb{D}_{\text{gather}} &= \bigcup_{k+1 \in 2\mathbb{Z}} [\varphi^{k+1}, \varphi^k), \end{aligned}$$

where $\varphi = \frac{2}{1+\sqrt{5}} = 0.618\dots$ is the golden ratio. These intervals are logarithmically spaced and serve to discretize distance observations in a self-consistent manner.

The main challenge addressed by the algorithm is indeed the lack of consistency in perceived distances: each robot uses its own arbitrary unit scale. Hence, we cannot directly rely on any single robot's perception to define globally consistent behavior. The key idea is to enforce a structure on the observed distances that causes predictable transitions between configurations.

The algorithm operates on configurations of three robots, moving through the following high-level phases:

1. **Initialization and Contraction:** Starting from any arbitrary configuration (collinear or non-collinear), the robots execute moves that eventually reduce the diameter of the configuration without ever increasing it. Through repeated activations and based on the

classification of the observed triangle's side lengths, the algorithm gradually transforms any initial triangle into one with a unique shortest side. This prepares the configuration for controlled symmetry breaking.

2. **Controlled Symmetry Breaking:** Once a triangle with a unique shortest side is achieved, the robot not adjacent to this side is eventually activated and may attempt to move toward the other robots. However, this action is taken only when the perceived length of the shortest side falls into $\mathbb{D}_{\text{golden}}$. This restriction ensures that the two-point configuration created is well-structured and predictable in terms of how the observing robot will behave.
3. **Two-Point to Three-Point Cycling:** When a two-point configuration is created, the robot that initiated the move is now collocated with one of the other robots, while the third robot is at a different location. The key insight is that based on the perceived distance (again classified into $\mathbb{D}_{\text{golden}}$ or $\mathbb{D}_{\text{gather}}$), robots either remain stationary or move in small, controlled steps to create specific collinear three-point configurations. This cycling between two-point and three-point formations allows the robots to coordinate indirectly through geometry and ensures that enough variations of distance configurations are explored.
4. **Eventual Rigid Movement:** Another critical component of the algorithm is that it ensures a bounded decrease of the configuration diameter after a finite number of steps. Because the system allows only non-increasing diameter movements, and the distance thresholds used in the classification scheme are strictly decreasing with powers of φ , the configuration's diameter eventually becomes smaller than any possible non-rigid movement threshold. This guarantees that we can thereafter assume all robot moves to be rigid.
5. **Convergence to Gathering:** After sufficiently many transitions between specific two-point and three-point configurations, the robots arrive at a configuration where all three perceive distances from $\mathbb{D}_{\text{gather}}$ and are activated in one of a bounded number of scenarios. Each of these is explicitly handled in the algorithm and shown to lead to gathering within a finite number of steps, regardless of the scheduler's choices or initial scale mismatches. ◀

► **Theorem 2.** *Distinct Gathering under ROUNDROBIN is possible for 3 robots with a self-consistent compass and the otherwise default assumptions detailed in the introduction, even if movement may be nonrigid.*

Proof Sketch. The algorithm, which we have to defer to the future full paper due to the space constraint, solves the Distinct Gathering problem under ROUNDROBIN for three robots with self-consistent compasses, even when faced with nonrigid movement.

The overall strategy of the algorithm is structured in two main phases. The first phase achieves a collinear configuration in which all robots possess a well-defined and consistent sense of direction along the shared line. A robot derives such a sense of direction from its self-consistent compass vector if this vector is not orthogonal to the line formed by the robot positions. We can prove that such a collinear configuration with a self-consistent sense of direction for all robots eventually emerges from any initial configuration. Furthermore, we can show that the diameter of the configuration decreases over time as long as nonrigid movements occur, thus guaranteeing that rigid movement can eventually be assumed without loss of generality.

Once a collinear configuration as described is reached, the second phase of the algorithm begins. This phase relies on the robots' ability to break symmetry using their self-consistent compasses. The robots restrict their movements to the line formed by them and behave

■ **Algorithm 1** An algorithm for Distinct Gathering under ROUNDROBIN for $n = 3$ robots that works with potentially inconsistent but always self-consistent perception of unit length and for nonrigid movement with an unknown minimum movement distance for each robot and otherwise the default assumptions detailed in the introduction. We denote by φ the golden ratio 0.618... solving the equation $\varphi + \varphi^2 = 1$. Moreover, we denote by $\mathbb{D}_{\text{golden}} = \cup_{k \in 2\mathbb{Z}}[\varphi^{k+1}, \varphi^k)$ and $\mathbb{D}_{\text{gather}} = \cup_{k+1 \in 2\mathbb{Z}}[\varphi^{k+1}, \varphi^k)$ the two sets of the bipartition $(\mathbb{D}_{\text{golden}}, \mathbb{D}_{\text{gather}})$ of the positive real numbers.

Input: A distinct configuration of robots in the Euclidean plane, observed in a Cartesian coordinate system with the observing robot at the origin.

Output: A target location that the observing robot moves towards on a straight line, not arriving at the target only when interrupted by nonrigid movement.

Algorithm:

- (1) **If** a *one-point configuration* (i.e., a gathering) is observed, **then** no movement is made.
- (2) **If** a *two-point configuration* is observed, **then** there are two possible behaviors and which one is chosen depends on the perceived distance d between the two locations:
 - A. **If** $d \in \mathbb{D}_{\text{gather}}$, **then** the opposite location is targeted.
 - B. **If** $d \in \mathbb{D}_{\text{golden}}$, **then** the activated robot targets the point between the two locations that is at perceived distance $\varphi^2 d$ from the observation point and thus at perceived distance φd from the opposite location.
- (3) **If** a *collinear three-point configuration* is observed, **then** the behavior depends on the observed relative distances:
 - A. **If** the ratio between the two distances from the middle robot to the outer robots is the golden ratio and the observing robot is not the middle one, the behavior depends again on the observed distance d between the observing robot and the middle robot:
 - a. **If** $d \in \mathbb{D}_{\text{golden}}$, **then** the robot stays where it is.
 - b. **If** $d \in \mathbb{D}_{\text{gather}}$, **then** the behavior depends on whether the middle robot is closer to the observing robot or the other one.
 - α . **If** the middle robot is closer to the non-observing robot, **then** the target is the middle robot.
 - β . **If** the middle robot is closer to the observing robot, **then** the trajectory goes past the middle robot to the location between the middle robot and the other non-observing robot that was from the middle robot a distance that is a factor φ of the distance to the other non-observing robot.
 - B. **If** the observing robot is the middle one, the behavior depends on the ratio between the two distances from the middle robot to the outer robots:
 - a. **If** the ratio is φ or $1/\varphi$, **then** the target is the second location with the ratio φ or $1/\varphi$ between the outer robots.
 - b. **If** the ratio is neither φ nor $1/\varphi$, **then** the activated robot moves towards an arbitrary one of the two points forming an isosceles right triangle with the outer two robots.
- (4) **If** a robot observes a *non-collinear three-point configuration*, **then** the behavior depends on whether the observing robot is incident to any of the shortest sides in the triangle formed by the three robot locations:
 - a. **If** the observing robot is incident to such a shortest side, **then** it targets the midpoint of this shortest side, breaking potential ties arbitrarily.
 - b. **If** the observing robot is not incident to such a shortest side, **then** there is only one shortest side and the behavior depends on the observed length d of this side:
 - α . **If** $d \in \mathbb{D}_{\text{golden}}$, **then** the target is any arbitrary one of the two endpoints of the unique shortest side.
 - β . **If** $d \in \mathbb{D}_{\text{gather}}$, **then** the robot stays where it is.

differently depending on their observed relative positions and compass directions. The correctness of this phase is verified through an exhaustive case analysis of all 48 possible directional configurations that they eventually lead to a gathering.

Together, these two phases ensure that the algorithm terminates with all robots gathered at a single location. ◀

3 Conclusion

We have explored the minimal assumptions required to make the Distinct Gathering of three robots under ROUNDROBIN feasible. The most recent algorithm attempting to minimize the assumptions still required rigidity of movement, a common unit distance, and additionally memory. We have provided an algorithm that reduces all of this to only a self-consistent unit length. We further have an algorithm that requires only a self-consistent compass instead of the self-consistent unit. Both results highlight the importance of the commonly neglected distinction between robot perceptions that are potentially not consistent but at least self-consistent and fully volatile perceptions without any such guarantees.

References

- 1 N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM Journal on Computing*, 36(1):56–82, 2006. doi:10.1137/050645221.
- 2 Z. Bouzid, S. Das, and S. Tixeuil. Gathering of mobile robots tolerating multiple crash faults. In *the 33rd Int. Conf. on Distributed Computing Systems*, pages 334–346, 2013.
- 3 M. Cieliebak, P. Flocchini, G. Prencipe, and N. Santoro. Distributed computing by mobile robots: Gathering. *SIAM Journal on Computing*, 41(4):829–879, 2012. doi:10.1137/100796534.
- 4 Xavier Défago, Maria Gradinariu, Stéphane Messika, and Philippe Raipin Parvédy. Fault-tolerant and self-stabilizing mobile robots gathering. In Shlomi Dolev, editor, *Proceedings of the 20th International Symposium on Distributed Computing (DISC 2006)*, volume 4167 of *Lecture Notes in Computer Science*, pages 46–60. Springer, 2006. doi:10.1007/11864219_4.
- 5 Xavier Défago, Maria Potop-Butucaru, and Philippe Raipin Parvédy. Self-stabilizing gathering of mobile robots under crash or Byzantine faults. *Distributed Comput.*, 33(5):393–421, 2020. doi:10.1007/s00446-019-00359-x.
- 6 B. Degener, B. Kempkes, T. Langner, F. Meyer auf der Heide, P. Pietrzyk, and R. Wattenhofer. A tight run-time bound for synchronous gathering of autonomous robots with limited visibility. In *23rd ACM SPAA*, pages 139–148, 2011.
- 7 Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. *Distributed Computing by Oblivious Mobile Robots*. Synthesis Lectures on Distributed Computing Theory. Morgan and Claypool Publishers, 2012.
- 8 Fabian Frei and Koichi Wada. Brief announcement: Distinct Gathering under Round Robin. *Leibniz International Proceedings in Informatics, LIPIcs*, 319, 2024. doi:10.4230/LIPIcs.DISC.2024.48.
- 9 Fabian Frei and Koichi Wada. Invited paper: Gathering oblivious robots in the plane. In Toshimitsu Masuzawa, Yoshiaki Katayama, Hirotugu Kakugawa, Junya Nakamura, and Yonghwan Kim, editors, *Stabilization, Safety, and Security of Distributed Systems – 26th International Symposium, SSS 2024, Nagoya, Japan, October 20–22, 2024, Proceedings*, volume 14931 of *Lecture Notes in Computer Science*, pages 39–54. Springer, 2024. doi:10.1007/978-3-031-74498-3_3.
- 10 T. Izumi, Z. Bouzid, S. Tixeuil, and K. Wada. Brief announcement: The BG-simulation for Byzantine mobile robots. In *25th DISC*, pages 330–331, 2011.

- 11 S. Kamei, A. Lamani, F. Ooshita, and S. Tixeul. Asynchronous mobile robot gathering from symmetric configurations without global multiplicity detection. In *18th SIROCCO*, pages 150–161, 2011.
- 12 S. Souissi, X. Défago, and M. Yamashita. Using eventually consistent compasses to gather memory-less mobile robots with limited visibility. *ACM Transactions on Autonomous and Adaptive Systems*, 4(1):1–27, 2009. doi:10.1145/1462187.1462196.
- 13 Ichiro Suzuki and Masafumi Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM J. Comput.*, 28(4):1347–1363, 1999. doi:10.1137/S009753979628292X.
- 14 Satoshi Terai, Koichi Wada, and Yoshiaki Katayama. Gathering problems for autonomous mobile robots with lights. *Theor. Comput. Sci.*, 941:241–261, 2023. doi:10.1016/j.tcs.2022.11.018.