



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/243146/>

Version: Accepted Version

Proceedings Paper:

Soma, K., Talamali, M.S., Miyauchi, G. et al. (2026) Split over n resource sharing problem: are fewer capable agents better than many simpler ones? In: Groß, R., Becker, A.T., Di Caro, G., Haghighat, B., Ani Hsieh, M., Abu-Aisheh, R., Talamali, M.S. and Dorigo, M., (eds.) UNSPECIFIED 15th International Conference, ANTS 2026, 08-10 Jun 2026, Darmstadt, Germany. Lecture Notes in Computer Science, 16515 LNCS. Springer Nature Switzerland, pp. 438-446. ISBN: 9783032261229.

https://doi.org/10.1007/978-3-032-26123-6_39

© 2026 The Authors. Except as otherwise noted, this author-accepted version of a journal article published in Swarm Intelligence is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Split Over n Resource Sharing Problem: Are Fewer Capable Agents Better Than Many Simpler Ones?

Karthik Soma^{1,2*}, Mohamed S. Talamali², Genki Miyauchi², Giovanni Beltrame¹, Heiko Hamann³, and Roderich Groß^{2,4}

¹ MIST lab, Polytechnique Montréal, Montreal, Canada

`karthik.soma@polymtl.ca`

² School of Electrical and Electronic Engineering, The University of Sheffield, Sheffield, UK

³ Department of Computer and Information Science, University of Konstanz, Konstanz, Germany

⁴ Department of Computer Science, Technical University of Darmstadt, Darmstadt, Germany

Abstract. In multi-agent systems, should limited resources be concentrated into a few capable agents or distributed among many simpler ones? This work formulates the split over n resource sharing problem where a group of n agents equally shares a common resource (e.g., monetary budget, computational resources, physical size). We present a case study in multi-agent coverage where the area of the disk-shaped footprint of agents scales as $1/n$. A formal analysis reveals that the initial coverage rate grows with n . However, if the speed of agents decreases proportionally with their radii, groups of all sizes perform equally well, whereas if it decreases proportionally with their footprints, a single agent performs best. We also present computer simulations in which resource splitting increases the failure rates of individual agents. The models and findings help identify optimal distributiveness levels and inform the design of multi-agent systems under resource constraints.

1 Introduction

Artificial systems must operate with limited resources. These resources could be concentrated into a single agent. Alternatively, they could be divided among multiple agents. For example, consider a system designer with a limited monetary budget to purchase robots for a given mission. If opting for a single-agent system, they could afford a high-spec robot featuring advanced actuation, sensing, and computational resources. However, if opting for a system comprising many agents, they could only afford robots of substantially lower cost (per unit) featuring lower quality actuation, sensing, and computational resources.

This paper formulates the aforementioned dilemma as the *split over n resource sharing problem* (see Fig. 1), where n agents are assumed to each possess

* Corresponding author.




Fig. 1. Split over n resource sharing problem. A finite resource is equally shared among n robots; What is the optimal n ? Different values of n might correspond to different robot platforms, for example, Turtlebot4 [1], e-puck2 [9], and Kilobot [14].

an equal share of a limited resource. The problem addresses a fundamental question in swarm intelligence: given a fixed resource, what is the optimal level of distributiveness?

The split over n resource sharing problem differs from the vast literature in the field of swarm intelligence that focuses on the optimal group size under the assumption of an effectively unlimited supply of identical robots. For example, will 5, 50, or 500 robots of a particular design yield the best performance? Such studies [10,11] provide models and answers for questions about optimal group size or swarm density (robots/area). These models show an initial increase of system performance with increasing number of robots and robot density that can even be superlinear, that is, doubling the number of robots results in more than doubled system performance [10]. There is an optimal number of robots (and density) at which system performance is maximized. Increasing the number of robots beyond this critical value leads to performance degradation due to congestion or interference [16].

This paper investigates the split over n resource sharing problem in the context of a spatial coverage task [6,3,13], where agents must visit all locations of the environment collectively. Spatial coverage tasks are fundamental to many real-world applications, such as vacuuming floors and mapping. As coverage depends on both footprint and mobility, it serves as a benchmark for this study.

The contributions of this work are threefold. First, we formulate the split over n resource sharing problem, the solution to which may inform optimal level of distributiveness in multi-robot system design assuming a constant resource budget. Second, we present a multi-agent coverage case study incorporating practical miniaturization constraints, such as reduced mobility and higher failure rates. Third, we formally model the system and reveal conditions under which the optimal level of distributiveness is minimal, maximal, or irrelevant.

2 Formulation of the Split Over n Resource Sharing Problem

We assume a periodic environment of size $p \times p$ arbitrary units. It is populated by n identical, disk-shaped agents. Let $A > 0$ denote the shared resource, representing the combined footprint of all agents. Hence, the resource (i.e., footprint)




Fig. 2. (a) The four velocity profiles as functions of n . (b) Area covered by a disk-shaped agent of radius r moving forward with velocity V .

per agent is A/n . Each agent has a continuous position in 2-D space and an orientation. It propels by alternating between (i) moving forward with velocity v for δ time steps, and (ii) rotating instantaneously by angle θ (in rad). δ is drawn from a power-law distribution and turning angles θ from a wrapped Cauchy distribution to realize a random walk [7], given by:

$$P_\alpha(\delta) \propto \delta^{-(\alpha+1)}, \quad 0 < \alpha \leq 2. \quad f_w(\theta; \mu, \rho) = \frac{1}{2\pi} \frac{1 - \rho^2}{1 + \rho^2 - 2\rho \cos(\theta - \mu)}$$

We consider Brownian random motion [7] ($\rho = 0$, $\mu = 0$, $\alpha = 2$) and investigate four velocity profiles that capture practical conditions where footprint miniaturization imposes limits on the maximum achievable velocity of smaller sized agents (see Fig. 2a):

1. *constant*: an agent's velocity is constant, V_0 , that is, irrespective of the group size (n);
2. *linear*: an agent's velocity is V_0 if $n = 1$, however, as n increases, it decreases proportionally until reaching zero;
3. *radius*: an agent's velocity is proportional to $\frac{1}{\sqrt{n}}$. Hence, it is proportional to the radius of its footprint, which is $\sqrt{\frac{A}{n\pi}}$;
4. *area*: an agent's velocity is proportional to $\frac{1}{n}$. Hence, it is proportional to the area of its footprint, which is $\pi \sqrt{\frac{A}{n\pi}} \sqrt{\frac{A}{n\pi}} = \pi \frac{A}{n\pi} = \frac{A}{n}$.

To determine the area covered, we discretize the environment into $m \times m$ square cells. At any given time, an agent senses all cells whose centers lie within its footprint.

The coverage percentage up to time t is given as:

$$c(t) = \frac{\sum_{i=1}^m \sum_{j=1}^m P_{i,j}(t)}{m^2} \cdot 100 \quad P_{i,j}(t) = \begin{cases} 1 & \text{if cell } (i,j) \text{ covered} \\ 0 & \text{otherwise} \end{cases}$$

Let t_f denote the first time at which the coverage reaches 100%, where $t_f = \inf\{t \in \mathbb{R}^+ \mid c(t) = 100\}$ and $\inf$ denotes the infimum, or the greatest lower bound. The objective is to find the optimal level of distributiveness (n): $\operatorname{argmin}_n t_f$.

Table 1. Initial coverage rates for various velocity profiles (from Fig. 2a).

Constant	Linear
$\frac{d\mathcal{P}_0}{dt} \propto 2\sqrt{\frac{An}{\pi}} \cdot V_0 \propto \sqrt{n}$	$\frac{d\mathcal{P}_0}{dt} \propto 2\sqrt{\frac{An}{\pi}} \cdot (V_0 - \gamma(n-1)) \propto a\sqrt{n} - bn^{3/2}$
Radius	Area
$\frac{d\mathcal{P}_0}{dt} \propto 2\sqrt{\frac{An}{\pi}} \cdot \frac{V_0}{\sqrt{n}} \propto \mathcal{P}_c$	$\frac{d\mathcal{P}_0}{dt} \propto 2\sqrt{\frac{An}{\pi}} \cdot \frac{V_0}{n} \propto \frac{1}{\sqrt{n}}$

3 Formal Analysis

We derive the initial coverage rates for the four velocity profiles, ignoring overlaps with other agents. Let $\mathcal{P}_0^f$ and $\mathcal{P}_0^i$ denote the area that an agent had covered after and prior to a given placement, respectively, and r denotes its radius. An agent covers new area at a rate proportional to its diameter and velocity (see Fig. 2b). Therefore, as $r = \sqrt{\frac{A}{n\pi}}$, the initial collective coverage rate $\frac{d\mathcal{P}_0}{dt}$ of all agents is:

$$\begin{aligned}
d\mathcal{P}_0 &\propto n \times (\mathcal{P}_0^f - \mathcal{P}_0^i) \\
d\mathcal{P}_0 &\propto n \times (V dt 2r + \pi r^2 - \pi r^2) \\
\frac{d\mathcal{P}_0}{dt} &\propto n \times 2r \times v(n) \propto 2 \times \sqrt{\frac{An}{\pi}} \times v(n)
\end{aligned}$$

Table 1 presents the initial coverage rates for all velocity profiles.

4 Simulation Results

This section uses computer simulations to evaluate how different levels of distributiveness affect coverage performance for disk-shaped agents (see Section 2) under no gradual motion, varying velocity profiles, failure rates, and collisions. Each configuration was tested across 30 different initial conditions. The cumulative coverage percentage $c(t)$ achieved by the agents was recorded using tools from the OpenCV library [2]. We use a no-split agent velocity: $V_0 = 0.005$ units/step, a linear velocity-profile slope: $\gamma = 4 \times 10^{-6}$, and an environment of size: $p \times p = 1 \times 1$ sq. units discretized into $m \times m = 1000 \times 1000$ cells. The total footprint is fixed at $A = \pi \times 0.1^2$ sq. units.

4.1 Effect of No Gradual Motion

We investigate the scenario in which agents can instantly appear in any random position in the environment without undergoing gradual movement. This experiment isolates the effect of resource splitting from other influences, such as velocity profiles and physical collisions altering the course. These experiments use the parameters described above, except the total number of cells is increased to $10^4 \times 10^4$. Moreover, coverage is reported once per *repositioning round*, that is, each time the entire group has been repositioned. As seen in Fig. 3a, the coverage trends are identical across all group sizes.




Fig. 3. (a) Coverage trends when agents instantly relocate to random positions (b) Coverage at 100th simulation step with various velocity profiles.

Coverage initially even match ideal rates, corresponding to perfect avoidance of revisiting explored regions (i.e., coverage rate of A per round). However, after a few rounds, the coverage achieved for every group size begins to deviate from the ideal rates, as overlaps with already visited regions increase and the probability of visiting unvisited cells decreases proportionally to the overall coverage.

4.2 Effect of Velocity Profiles

We investigate the effects of velocity profiles (see Fig. 2a). We first examine the initial performance, reporting coverage percentages at the 100th simulation step. As shown in Fig. 3b, the coverage trends match those of the formal analysis in Table 1. In particular:

Constant: Coverage increases monotonically with the number of splits. For the group sizes considered, 1000 agents achieve the highest coverage, indicating that the optimal group size tends toward ∞ agents, consistent with Table 1. In practice, this suggests distributing resources among as many agents as possible.

Linear: Of the group sizes tested, $n = 500$ agents achieve the highest coverage, implying that the optimal group size lies between the two extremes (1 and ∞). To understand this behavior, consider the relationship between coverage rate and $\sqrt{n}$ from Table 1, which takes the form $y = ax - bx^3$, where y is the coverage rate, $x = \sqrt{n}$, $a, b > 0$ and $n \geq 1$. This cubic equation attains its local maximum at $x = \sqrt{\frac{a}{3b}}$, or equivalently, at $n = \frac{a}{3b}$. For the parameters in our setup, this corresponds to $n = 417$ (and $n = 500$ is the closest from the set of tested agent numbers).

Radius: Coverage remains nearly constant across all group sizes. This indicates that an increase in benefit from distributiveness compensates the reduced mobility of an individual agent at larger group sizes.

Area: Coverage decreases as the group size (n) increases, showing that it is optimal to keep all resources within a single agent.

These trends persist beyond the initial steps, as seen in the top row of Fig. 4. In the constant profile, coverage increases from the darkest to the lightest shade of red (group sizes $1 \rightarrow 1000$). For the linear profile, the performance order is as follows: the second-lightest green (500 agents) achieve the highest coverage, followed by 100, 1000, 10, 2, and 1 agents that perform in decreasing order thereafter. In the radius profile, all shades remain nearly identical, reflecting the uniform coverage across group sizes. In the area profile, coverage decreases from the darkest to the lightest purple ($1 \rightarrow 1000$).




Fig. 4. Coverage trends across velocity profiles. The top row shows collision-free results for 1, 2, 10, 100, 500, and 1000 agents; the bottom row shows results with collisions for 2, 100, and 1000 agents. In each profile, colour intensity encodes group sizes, with darker shades indicating fewer agents.

4.3 Effect of Collisions

While the previous sections discarded the effect of collisions, this section quantifies their impact. In multi-robot systems, when robots are about to collide, they typically slow down and rotate away from each other until the situation is resolved, after which they resume their task at the original velocity [8]. These robot-robot interferences commonly accumulate to long-term [15] and global effects [15,11]. We approximate this behavior by decreasing an agent’s velocity proportionally to the overlap between its footprint and that of any colliding neighboring agents, while ignoring heading corrections. To prevent deadlocks, a residual velocity is maintained even in the case of full overlap.

We evaluate the effect of collisions on all velocity profiles. As shown in the bottom row of Fig. 4, there is no significant decrease in coverage for experiments where the total footprint split over n agents is $A = \pi \times 0.1^2$ sq. units, and the coverage trends closely resemble those observed in the top row of Fig. 4. For larger footprints, deviations may increase due to the higher impact of an individual overlap with other agents.

Notably, unlike classical scalability studies [10,16], where a fixed footprint per agent leads to congestion-induced performance degradation as group size grows, agent density remains constant here. Nevertheless, except for the constant velocity profile, increasing n still reduces agent velocities, emulating congestion like effects.

4.4 Effect of Failures

Typically, in swarm robotics, increasing the number of robots introduces miniaturization-related risks [17,12] that may cause agents to fail over the course of the experiment [4,5]. In this section, we model such effects through a failure rate $k(n)$ that depends only on n . At each simulation step, we sample a uniform random number $r \sim \mathcal{U}(0,1)$ for every agent, and if $r < k(n)$, the agent ceases motion and does not contribute to coverage thereafter. We define $k(n) = \beta(1 - \frac{1}{n^\alpha})$, which is a non-linear, monotonically increasing function with a horizontal asymptote. For any $\alpha > 0$ and $\beta \geq 0$, the failure




Fig. 5. Coverage trends for different failure rates (left: $\alpha = 0.01$ and failure rate $k < 0.007$; right: $\alpha = 0.1$ and failure rate $k \leq 0.05$). α controls the rate at which failure rate approaches its maximum failure rate $\beta = 0.1$, where $k(n) = \beta(1 - \frac{1}{n\alpha})$.

rate for $n = 1$ is 0. Hence, a sole robot would always be assumed fault-free. $\beta > 0$, determines the maximum failure rate for $n \rightarrow \infty$. α controls how rapidly the failure rate approaches β as n increases (except for $n = 1$).

The experiment compares the performance of a single agent with zero failure probability to that of multi-agent groups with failure rate $k(n)$. We use $\beta = 0.1$ and $\alpha \in \{0.01, 0.1\}$ in this study. All agents have constant velocity $V_0 = 0.005$ units/step.

As shown in Fig. 5, for small $\alpha = 0.01$, large group sizes (e.g., 1000 agents) remain robust to failures and consistently outperform a single agent, whereas small groups eventually fall behind due to a larger proportion of agents failing. For the large $\alpha = 0.1$ and for the same n , the failure rate approaches β more rapidly, causing agents to fail earlier in the simulation. Under this condition, a single agent overtakes every group size n even before covering a significant proportion of the environment.

5 Conclusions

In this work, we formulated the split over n resource sharing problem. It explores a fundamental question in swarm robotics: What level of distributiveness is optimal? We presented a case study built around a simple coverage task. We formally derived the initial coverage rates for a range of velocity profiles. Through a series of computer simulations, we showed that technological constraints due to miniaturization, such as reduced mobility and increased failure rates, shift the optimal level of distributiveness. Depending on the type of robots available, our findings could inform system designers in determining whether a single highly capable robot, several moderately capable robots, or a large group of smaller, less capable robots would be most effective for a coverage task. Future work will investigate the performance of heterogeneous robot swarms, combining highly capable and less capable robots. We will also study how these findings translate to physics-based simulations, assessing whether the observed trends hold.

Acknowledgments. This research was supported in part by the Mitacs Globalink Research Award, BMBF (Robotics Institute Germany; grant no. 16ME1001) and EU Horizon Europe Framework Programme (“OpenSwarm”; grant no. 101093046).

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Amsters, R., Slaets, P.: Turtlebot 3 as a Robotics Education Platform, pp. 170–181 (01 2020). https://doi.org/10.1007/978-3-030-26945-6_16
2. Bradski, G.: The opencv library. In: Dr. Dobb’s Journal of Software Tools (2000)
3. Breitenmoser, A., Martinoli, A.: On combining multi-robot coverage and reciprocal collision avoidance. Springer Tracts in Advanced Robotics, vol. 112, p. 49–64. Springer Japan, Tokyo (2016). https://doi.org/10.1007/978-4-431-55879-8_4, <https://infoscience.epfl.ch/handle/20.500.14299/127959>
4. Caprari, G., Estier, T., Siegwart, R.: Fascination of down scaling-Alice the sugar cube robot. In: Proceedings of The IEEE International Conference on Robotics and Automation (ICRA), Workshop on Mobile Micro-Robots (2000)
5. Chin, K.Y., Khaluf, Y., Pincirol, C.: Minimalistic collective perception with imperfect sensors. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 8862–8868. IEEE (2023)
6. Choset, H.: Coverage for robotics – a survey of recent results. Annals of Mathematics and Artificial Intelligence **31**(1), 113–126 (2001)
7. Dimidov, C., Oriolo, G., Trianni, V.: Random walks in swarm robotics: An experiment with kilobots. pp. 185–196 (09 2016). https://doi.org/10.1007/978-3-319-44427-7_16
8. Goldberg, D., Matarić, M.J.: Interference as a tool for designing and evaluating multi-robot controllers. In: Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence. pp. 637–642 (1997)
9. Gonçalves, P., Torres, P., Alves, C., Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klapotcz, A., Magnenat, S., Zufferey, J.C., Floreano, D., Martinoli, A.: The e-puck, a robot designed for education in engineering. Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions **1** (01 2009)
10. Hamann, H.: Superlinear scalability in parallel computing and multi-robot systems: Shared resources, collaboration, and network topology. In: Berekovic, M., Buchty, R., Hamann, H., Koch, D., Pionteck, T. (eds.) Architecture of Computing Systems – ARCS 2018. pp. 31–42. Springer International Publishing, Cham (2018)
11. Hamann, H., Reina, A.: Scalability in computing and robotics. IEEE Transactions on Computers **71**(6), 1453–1465 (2022). <https://doi.org/10.1109/TC.2021.3089044>
12. Pfeiffer, C.: Fundamental efficiency limits for small metallic antennas. IEEE Transactions on Antennas and Propagation **65**(4), 1642–1650 (2017)
13. Ramesh, M., Imeson, F., Fidan, B., Smith, S.L.: Anytime replanning of robot coverage paths for partially unknown environments. IEEE Transactions on Robotics **40**, 4190–4206 (2024). <https://doi.org/10.1109/TR0.2024.3454417>
14. Rubenstein, M., Ahler, C., Nagpal, R.: Kilobot: A low cost scalable robot system for collective behaviors. In: 2012 IEEE International Conference on Robotics and Automation. pp. 3293–3298 (2012). <https://doi.org/10.1109/ICRA.2012.6224638>
15. Schroeder, A., Trease, B., Arsie, A.: Balancing robot swarm cost and interference effects by varying robot quantity and size. Swarm Intelligence **13**(1), 1–19 (2019)
16. Soma, K., Vardharajan, V.S., Hamann, H., Beltrame, G.: Congestion and scalability in robot swarms: A study on collective decision making. In: 2023 International Symposium on Multi-Robot and Multi-Agent Systems (MRS). pp. 199–206 (2023). <https://doi.org/10.1109/MRS60187.2023.10416793>

17. Yang, C., Van Der Drift, E., French, P.: Review of scaling effects on physical properties and practicalities of cantilever sensors. *Journal of Micromechanics and Microengineering* **32**(10), 103002 (2022)