



Information transport in communication limited swarms

James Wilson^{1,2} · Sabine Hauert^{1,2}

Received: 8 April 2022 / Accepted: 18 May 2022 / Published online: 20 June 2022
© The Author(s) 2022

Abstract

Users and operators of swarms will, in the future, need to monitor the operations of swarms in a distributed way, without explicitly tracking every agent, and without the need for significant infrastructure or set up. Here we present a method for swarm self-monitoring that enables the aggregate display of information about swarm location by making use of physical transport of information and local communication. This method uses movement already exhibited by many swarms to collect self-reflective information in a fully distributed manner. We find that added swarm mobility can compensate for limited communication and that our self-monitoring swarm system scales well, with performance increasing with the size of the swarm in some cases. When developing systems such as this for real-world applications, individual agent memory will need to be taken into consideration, finding an effective means to spread swarm knowledge among robots while keeping information accessible to users.

Keywords Swarm robotics · Swarm monitoring · Information extraction · Distributed robotics

1 Introduction

Robot swarms create opportunities within intralogistics and could possibly provide storage and organisation solutions to small businesses, such as museums, small retail, or charity stores; businesses that do not have dedicated, organised, warehouse environments [11]. These swarm solutions could be ready out of the box, with minimal infrastructure and setup [6]. Requiring as little as creating a space for charging and marking out a space to store objects. This would allow for easy adoption for these small businesses and move

swarm technology from existing research-focused use cases to more practical uses.

A recent survey conducted by Carrillo-Zapata et al. [1] has shown that there are many consumers within the bracket of small/medium enterprises that are interested in adopting swarm robotic technologies. However, one of the crucial caveats of swarm robotic adoption stems from a need for a business to understand the behaviour swarms exhibit. As a result, creating a distributed method for extracting data to allow a swarm to self monitor may help resolve issues in trust and accountability; allowing a user to monitor current or historic swarm behaviour to provide reassurance.

A distributed user interface such as this would see swarms share and collate self-monitoring information amongst agents, building up a picture of the overall swarm status. This information collection could take the form of human individuals using tablets or mobile devices to interact locally with a swarm, or a central computer node collecting and storing swarm knowledge to provide users with information at a later time. Having access to information on swarm state, location and performance can provide crucial insight to swarm activity. Helping to enable effective Human Swarm Interaction [8] or establish appropriate levels of trust in robotic systems [7].

At present there are many examples of information collection within swarms. However, most of these examples

This work was supported by the UKRI's Trustworthy Autonomous Systems Node in Functionality under Grant number EP/V026518/1. This work was presented in part at the joint symposium of the 27th International Symposium on Artificial Life and Robotics, the 7th International Symposium on BioComplexity, and the 5th International Symposium on Swarm Behavior and Bio-Inspired Robotics (Online, January 25–27, 2022).

✉ James Wilson
j.wilson@bristol.ac.uk

Sabine Hauert
sabine.hauert@bristol.ac.uk

¹ Bristol Robotics Laboratory, Bristol, UK

² Department of Engineering Mathematics, University of Bristol, Bristol, UK

lack a distributed element and are typically performed within a laboratory setting, using cameras, tags or permanent long range networks [9]. Putting together infrastructure such as this is something that may not be viable to pursue when implementing a swarm as a simple, ‘out-of-the-box’, solution designed to meet consumer needs. Existing solutions for real world swarm monitoring rely upon either expansive communication networks or dedicate many agents within the swarm to creating network nodes [2, 10, 14]; methodologies which can be expensive to implement or potentially reduce swarm performance.

When deploying in realistic environments outside of the lab, inter-robotic connectivity cannot always be assumed. The requirement for anticipating and designing a swarm for low communications availability has been highlighted previously by Tarapore et al. [12]. In this paper we address this, presenting a new methodology for information extraction from robot swarms. Instead of a highly connected swarm, we look to have agents physically transport crucial swarm information when outside of communication range, sharing the most up to date information available when locally encountering other agents. This method distributes information throughout the swarm, ready to be collected at any given time by an observer, allowing them to gain insight on swarm state, performance, or diagnostics.

Our method of swarm communication takes inspiration from Ducatelle et al. [4] in which spatial information is shared throughout a swarm, allowing robots to physically transport information when not connected wirelessly. This system prioritised the most up to date information on robot-to-robot distance to provide navigational assistance to swarm agents. Our system similarly removes focus on creating permanent and stable networks in which robots constantly transmit information to one another. We build upon this existing technology using ad hoc information transmission to spread and update swarm data, ready for user collection at any time.

A recent study has called to attention that explicit representation of swarms can in some cases be detrimental to human-swarm interaction due to the high cognitive complexity required to monitor large volumes of agents [3]. Our method looks to provide a simple and effective means for human-swarm interaction by providing easy to consume summary information of the swarm present in areas the user has identified as key. A visualisation of this, and how it may be expanded to additional information types, can be seen in Fig. 1.

Our distributed approach capitalises upon the repetitive movement many swarms already exhibit when performing logistical or monitoring tasks. These repetitive movements create areas in which robots frequently meet one another, providing them with the opportunity to share information without large communication ranges or networks. By taking advantage of the movement already exhibited in the




Fig. 1 Illustration showing how collected information could be consumed by a user

swarm’s standard behaviour, information can be propagated throughout the swarm over time without the need to explicitly coordinate with one another. Thus, the swarm does not have to sacrifice performance in its underlying task through the need to prioritise communication related movement over task behaviours.

In the following sections we present our swarm monitoring methodology and explore the effects of modifying communication range, number of robots in the swarm, maximum robot speed and robot memory on the accuracy of information extracted from the swarm. We then discuss the considerations that should be made when producing a distributed self-monitoring swarm and the potential difficulties that may be encountered in physical implementation.

2 Distributed swarm information extraction

In the experiments presented in this paper we use a task inspired from logistics as a swarm case study. The chosen task represents an easy to deploy swarm use, aligning with our goal of an out-of-the-box swarm solution. The materials we represent in simulation for this task, in reality, would only require a user held tablet device and paper tags such as QR codes or ArUco tags in addition to the swarm of robots. This simple set up should still provide good insight to the whereabouts and status of the swarm.

We propose a methodology for monitoring a swarm, allowing a user to quickly request swarm information in an ad hoc, distributed manner. In our experiments we mimic a swarm warehouse scenario, introducing a data collection point representing a person with a tablet device entering a swarm-organised storage area. Swarm agents are not aware of the location of the data collection point and only share information with this point when they happen to pass within communications range, travelling with the information indefinitely unless overwritten by newer data.

The task we have selected for the swarms to execute in our experiments mimics the movement that might be seen in

a swarm performing duties they are often designed for, i.e., exploration, area monitoring or multi-object transport. In the chosen task, swarm agents move between a designated central point and one of four additional points of interest. The central point represents a location robots would frequently visit such as a charging station or a warehouse drop off point. Once arriving at any of these four points of interest, the swarm agents spend a random amount of time remaining within a short distance of the point. The agent will then return to the central point or, if already there, randomly select a new point of interest to travel to.

While executing this task, swarm agents will use local interactions with other agents to update a record of swarm information. For the purposes of these experiments, the swarm will be recording which point of interest agents are nearest too at any given time. These points of interest can be thought of as physical, user placed tags positioned in doorways, corridors, or other areas the user has decided are important for gauging swarm distribution. The accuracy of the swarm monitoring information arriving at the collection point will form the performance metric for our system.

3 Methods

The following section details the methodology we have developed to allow a swarm to collect and share self-monitoring information. As well as the methods through which we have scored and simulated the swarm.

3.1 Scenario

Swarm agents move between points of interest gathering and sharing information on robots within proximity of point of interest tags via swarm tables. We have programmed the agents to drive directly to their assigned tag, this movement is illustrated in Fig. 2. For the sake of our experiments, tag name and time of recording are the only data agents currently record. Each agent has their own internal table indexed by numerical IDs. Each of these IDs is uniquely associated with a specific swarm agent, allowing agents to overwrite information relating to a specific robot when newer robot-specific information is received. In real world implementations these IDs could be randomly generated to allow for scalability.

To test the proposed methodology, a custom 2D simulation was designed using python. In these experiments, robots represented by black dots in the simulation visualiser operate within a 20×20 meter arena. Agents, measuring 30cm in diameter, had the capacity to sense range and bearing, execute behaviour to avoid robot to robot and robot to object collisions and send information to one another when in communication range. Movement in these experiments was




Fig. 2 Final frame from a simulated experimental trial depicting the various paths of 5 agents take across the experiment length

simulated through agents moving on a bearing to the point of interest they next needed to visit with some random noise applied to velocity at each time step (± 0.5 cm/s). Agents were unable to pass through one another or obstacles in the environment and travelled in the opposite direction of the boundary/objects when within 30cm range.

In our first experiment, we investigate the trade off between the speed of the robots and their communication range, observing accuracy of the collected knowledge from the swarm as these parameters were modified. In this initial parameter exploration, the positions of the five tag sites and monitoring site were randomly generated. The locations of each of these sites were generated randomly every trial, were constrained to the arena boundaries, and could not be placed within 4 meters of one another. An example visualisation of an environment layout can be seen in Fig. 3.

3.2 Algorithm

The swarm agents append to and update their swarm tables in the following steps:

- (1) The swarm agent starts by checking if there is a tag associated with a point of interest within communication range. The agents then update their personal entry in the swarm table. Alongside this entry, the agent adds a timestamp to indicate the age of the recorded data.
- (2) If there are other robots within the communications range of the swarm agent, the agent iterates through the information stored in the tables of every currently con-




Fig. 3 Randomly generated virtual experiment environment. Displaying multiple tags and robots

nected robot, adding new entries to its personal table for data relating to unrecognised robots and updating entries that indicate having older information than available.

- (3) If the swarm agent is within perception range (5m) of an information collection point, all the information currently gathered is sent to the collecting dataframe and stored at the collection point. Data is then updated and appended as though the collection point is an agent itself. However, upon passing information to the collecting dataframe, the robot clears the swarm table it is currently storing to conserve memory and reduce the computational load of updating other robots with information the collection point already has a record of.

3.3 Performance metric

Throughout these experiments we will be monitoring the ongoing distribution of the swarm, identifying how many agents are within visible proximity of points of interest.

The score used in these experiments will be the total proportion of robot tag associations recorded correctly in the collecting dataframe at the reported time. Scores will range between 0 and 1, with 1 being a perfect categorisation of swarm distribution and 0 being a completely incorrect dataframe. For these experiments, specific IDs are not important due to the homogeneous nature of the swarm in question. Instead, we focus on the system's ability to accurately represent the swarm's distribution at each tag.

Each of our experiments simulated 2000 real-world seconds. This represents approximately 30 min, which was seen as a reasonable monitoring window for an operator. Experiments were repeated 20 times for each parameter subset.

The scores shown in our results are the averaged accuracy in distribution prediction from 1000 to 2000 s across

all repetitions, sampling accuracy once every 10 s. Scores begin from 1000 simulated seconds into the experiment to allow robots to disperse and share information among one another, as they would in the physical implementations previously discussed.

4 Results

4.1 Experiment 1: mobility vs communication distance

For this experiment, parameters ranged between 0 and 1400 cm for communication range and 0–100 cm/s for maximum robot speed. Parameter combinations were also tested across 4 different swarm sizes 5, 10, 20 and 40, to observe the scalability of this system. The results are shown in Fig. 3.

The results show that greater communication range improves the swarm's ability to correctly report on current distribution. With larger values resulting in highly connected networks, the largest of which being able to write to the collecting dataframe almost continuously.

However, as robot speed increases, we can see that across all swarm sizes scores begin to improve at lower communication ranges. In some cases in the 20 and 40 robot example we see very large increases in performance scores between speeds of 10 and 100 cm/s, in one example rising from a score of 0.19–0.79 at a communication range of 400 cm. This provides evidence to the argument that increased swarm mobility can compensate for limited communications by physically transporting information. This is possible even though faster moving robots will be encountering different tags more frequently, increasing the chances that information making its way to the central collection point will be invalid by the time it gets there, thus posing a more challenging task to keep track of. The physical transport of swarm information also appears to improve as swarm size increases, even though there are more robots to keep track of. An aspect of interest seen in these heat-maps is the sudden increase in score as robots drop in speed from 10 to 0 cm/s. This is likely taking place at speeds of 0 cm/s as robots that happen to be placed near the information collection point at the start of the experiment can broadcast their unchanging location along with the location of neighbouring robots.

4.2 Experiment 2—tag distance vs accuracy

In the second experiment we look to see how accuracy changes with the distance between tag sites and collection point. In this test we once again perform the same parameter sweep, but in this experiment we keep the tag sites in fixed positions. Figure 4. shows representative examples of these experiments, displaying the scores for a swarm of 20 robots




Fig. 4 Heat maps showing scores for distribution-reporting accuracy across swarms of size 5, 10, 20 and 40. These heat maps show the scores of 20 averaged trials for each subset of communications range and robot speed

(this swarm size was chosen as it reflects the number of physical robots we have available for this task for real-world trials) for the closest (5.66 m) and furthest (25.46 m) tag sites. The scores in this figure are now broken down by tag, showing the ratio of correctly recorded robots at a tag in the collecting dataframe vs the actual number of robots at said tag. We can see a similar pattern to those seen in experiment 1 forming across the spread of parameters.

This again shows that limited communications range can be some-what compensated for through increased mobility. We also see that distance between the information collection site and the tag has some effect on the accuracy of the data. Intuitively, information takes longer to arrive from distant tags and thus scores decrease as distance increases. In addition to this, robots will encounter new information from nearer tags more frequently, possibly displacing information on further tags as agents move. Despite this, most of the score change takes place in the extremes of the parameters explored here. This suggests that some parameter combinations could still form reasonably robust parameter sets for communication limited swarm monitoring, e.g., swarms of 20 and 40 robots, with max speeds over 60 cm/s, can

consistently achieve accuracy over 70% with a communication range of only 6 m.

4.3 Experiment 3—reduced memory

The results so far have pointed towards a scalable system. However, we currently assume that each swarm agent will have the capability to store a swarm-worth of information. This may not always be the case in situations with very large swarms or high volumes of robot data.

To test the resilience of our information extraction methodology we repeated Experiment 2, this time reducing the number of individual robot information each robot can store from 20 (the whole swarm size) to 5. In this experiment we still separate results by tag distance due to the large difference in success as tags move away from the collection point, an effect that cannot be clearly seen across the averaged tags data. These results are shown in Fig. 5.

Here we can see that there is a considerable drop in scores when looking at the average performance. Investigating further we can see that the accuracies of the tag at the closest point go almost unchanged; however, there is a considerable




Fig. 5 Heat maps for scores recorded for the closest and furthest tags in the virtual environment

decrease in score at the furthest point. This suggests that there will be issues with memory use when it comes to particularly large swarms or when surplus data is required from the swarm. With this being so, the total memory available to the swarm will increase with swarm size. If knowledge can be effectively distributed across the swarm while keeping information accessible, it may be possible to capitalise on a swarm's collective memory, making this method of swarm information extraction truly scalable.

5 Discussion

This methodology for swarm monitoring has shown that information accuracy can be increased in systems with limited communication by increasing the mobility of the swarm. One of the overarching decisions to be made when designing these systems will be around the acceptable level

of information loss. Depending on context and how critical a scenario might be, an accuracy as low as 0.6 could be reasonable to get an approximate image of swarm behaviour creating an interface resembling that shown in Fig. 1. However, in more critical scenarios with an emphasis on accuracy and up to date information, much higher scores may be required. Levels of acceptable accuracy within the system we have proposed should be further investigated in the future.

A user could use this information to travel closer to areas of interest, utilizing the increased accuracy that comes with reduced distance to gain additional insight on the swarm activities. The transparency a system like this would provide addresses the importance of understanding between user and swarm; one of the fundamental requirements needed for the practical adoption of swarm technology. The experiments conducted in this paper open up the opportunity for new research in how users can quickly monitor and control distributed systems, contributing towards the need to monitor swarms for safety purposes [5] and increasing user trust in swarm systems through better explained behaviour (Fig. 6).

This system could also be used to accompany a swarm focused Ethical Black Box. Winfield and Jirotko [13] have pointed towards a requirement for the robotic equivalent of a Flight Data Recorder when creating robotic systems for consumer adoption. A suggestion which is viable when implementing on individual robots but, with swarm robot solutions focusing on distributed, inexpensive, and mass-produced hardware, implementing a black box on every individual swarm member may have its difficulties. Having an additional method that records swarm level information could provide context to individual robot's data and could prove to be a valuable tool.

6 Conclusions and further work

We have demonstrated a method for monitoring swarm information in a decentralised manner by capitalising upon existing movement in the swarm. This method allows for breaks in communications networks, having agents physically transport information when they are unable to broadcast to neighbours.

We have shown that, when attempting to collect swarm data, increasing mobility of the swarm can compensate for limited communication. We also found that performance improved as swarm size increased. However, when considering that agents may be memory limited in certain circumstances, the scalability of this system will depend upon effective use of swarm memory. Thus, our future work may include attempting to distribute information storage amongst the swarm, using the swarm's collective memory to create a truly scalable solution. We may also explore methods in




Fig. 6 Heat maps for scores in simulation with memory limited agents. Displaying the average of all tags, the scores for the closest tag and the furthest

predicting the accuracy of received information to allow users to more effectively gauge trust in swarm data and act accordingly. This system could improve human-swarm interaction by suggesting the most effective areas for a human to monitor a swarm.

The end goal of these developments will be to use the information gained from our proposed monitoring methodology to create an interface allowing workers within small/medium businesses to easily understand, interact with, and trust a swarm. Opening up new opportunities for real world swarm applications and creating a swarm tool that can provide a competitive means for the intralogistic management of goods for businesses. Businesses that would not traditionally have access to automated solutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Carrillo-Zapata D, Milner E, Hird J, Tzoumas G, Vardanega PJ, Sooriyabandara M, Giuliani M, Winfield AF, Hauert S (2020) Mutual shaping in swarm robotics: user studies in fire and rescue, storage organization, and bridge inspection. *Front Robot AI* 7:53
2. Christensen AL, Oliveira SM, Postolache O, De Oliveira MJ, Sargento S, Santana P, Nunes L, Velez F, Sebastião P, Costa V, et al (2015) Design of communication and control for swarms of aquatic surface drones. In: ICAART 2015-7th international conference on agents and artificial intelligence, proceedings, pp 548–555
3. Divband Soorati M, Clark J, Ghofrani J, Tarapore D, Ramchurn SD (2021) Designing a user-centered interaction interface for human-swarm teaming. *Drones* 5(4):131
4. Ducatelle F, Di Caro GA, Förster A, Bonani M, Dorigo M, Magnenat S, Mondada F, O'Grady R, Pinciroli C, Rétornaz P et al (2014) Cooperative navigation in robotic swarms. *Swarm Intell* 8(1):1–33
5. Hunt ER, Hauert S (2020) A checklist for safe robot swarms. *Nat Mach Intell* 2(8):420–422
6. Jones S, Milner E, Sooriyabandara M, Hauert S (2020) Distributed situational awareness in robot swarms. *Adv Intell Syst* 2(11):2000,110
7. Kok BC, Soh H (2020) Trust in robots: challenges and opportunities. *Curr Robot Rep* 1–13
8. Kolling A, Walker P, Chakraborty N, Sycara K, Lewis M (2015) Human interaction with robot swarms: a survey. *IEEE Trans Hum Mach Syst* 46(1):9–26
9. Millard AG, Redpath R, Jewers AM, Arndt C, Joyce R, Hilder JA, McDaid LJ, Halliday DM (2018) Ardebug: an augmented reality

- tool for analysing and debugging swarm robotic systems. *Front Robot AI* 5:87
10. Stirling T, Roberts J, Zufferey JC, Floreano D (2012) Indoor navigation with a swarm of flying robots. In: 2012 IEEE international conference on robotics and automation. IEEE, pp 4641–4647
 11. Stormont DP (2008) Analyzing human trust of autonomous systems in hazardous environments. In: Proc. of the human implications of human-robot interaction workshop at AAAI, pp 27–32
 12. Tarapore D, Groß R, Zauner KP (2020) Sparse robot swarms: moving swarms to real-world applications. *Front Robot AI* 7:83
 13. Winfield AF, Jirotko M (2017) The case for an ethical black box. In: Annual conference towards autonomous robotic systems. Springer, pp 262–273
 14. Witkowski U, El Habbal MAM, Herbrechtsmeier S, Tanoto A, Penders J, Alboul L, Gazi V (2008) Ad-hoc network communication infrastructure for multi-robot systems in disaster scenarios. In: Proceedings of IARP/EURON workshop on robotics for risky interventions and environmental surveillance (RISE 2008), Benicassim, Spain

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.