

# Secure and secret cooperation in robot swarms

Eduardo Castelló Ferrer<sup>\*1,2,3</sup>, Thomas Hardjono<sup>2</sup>,  
Alex ‘Sandy’ Pentland<sup>1,2</sup>, Marco Dorigo<sup>3</sup>

<sup>1</sup>MIT Media Lab, <sup>2</sup>MIT Connection Science & Engineering  
Massachusetts Institute of Technology, Cambridge, MA 02139, USA

<sup>3</sup>IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

\*E-mail: ecstll@mit.edu.

**The importance of swarm robotics systems in both academic research and real-world applications is steadily increasing. However, to reach widespread adoption, new models that ensure the secure cooperation of large groups of robots need to be developed. This work introduces a novel method to encapsulate cooperative robotic missions in an authenticated data structure known as Merkle tree. With this method, operators can provide the “blueprint” of the swarm’s mission without disclosing its raw data. In other words, data verification can be separated from data itself. We propose a system where robots in a swarm, to cooperate towards mission completion, have to “prove” their integrity to their peers by exchanging cryptographic proofs. We show the implications of this approach for two different swarm robotics missions: foraging and maze formation. In both missions, swarm robots were able to cooperate and carry out sequential operations without having explicit knowledge about the mission’s high-level objectives. The results presented in this work demonstrate the feasibility of using Merkle trees as a cooperation mechanism for swarm robotics systems in both simulation and real-robot experiments, which has implications for future decentralized robotics applications where security plays a crucial role such as environmental monitoring, infrastructure surveillance, and disaster management.**

## Summary

Robot swarms exploiting Merkle trees achieve secret and secure cooperation in sequential missions.

## Introduction

Swarm robotics systems (1) have the potential to revolutionize many industries, from targeted material delivery (2) to precision farming (3, 4). Boosted by technical breakthroughs, such as cloud computing (5, 6), novel hardware design (7–9), and manufacturing techniques (10), swarms of robots are envisioned to play an important role in both industrial (12) and urban (13, 14) activities. The emergence of robot swarms has been acknowledged as one of the ten robotics grand challenges for the next 5-10 years that will have significant socioeconomic impact. Despite having such a promising future, many important aspects which need to be considered in realistic deployments are either underexplored or neglected (15).

One of the main reasons why swarms of robots have not been widely adopted in real-world applications is because there is no consensus on how to design swarm robotics systems that include perception, action, and communication (15). In addition, recent research points out that the lack of security standards in the field is also hindering the adoption of this technology in data-sensitive areas (e.g., military, surveillance, public infrastructure) (16, 17). These research gaps are motivating scientists to focus on new fields of study such as applied swarm security (18, 19) and privacy (20, 21) as well as to revisit already accepted assumptions in the field.

From the origins of swarm robotics research, robot swarms were assumed to be fault-tolerant by design, due to the large number of robot units involved (22–25). However, it has been shown that a small number of partially failed (with defective sensors, broken actuators, noisy communications devices, etc.) (26) or malicious robots (27, 28) can have a significant impact on

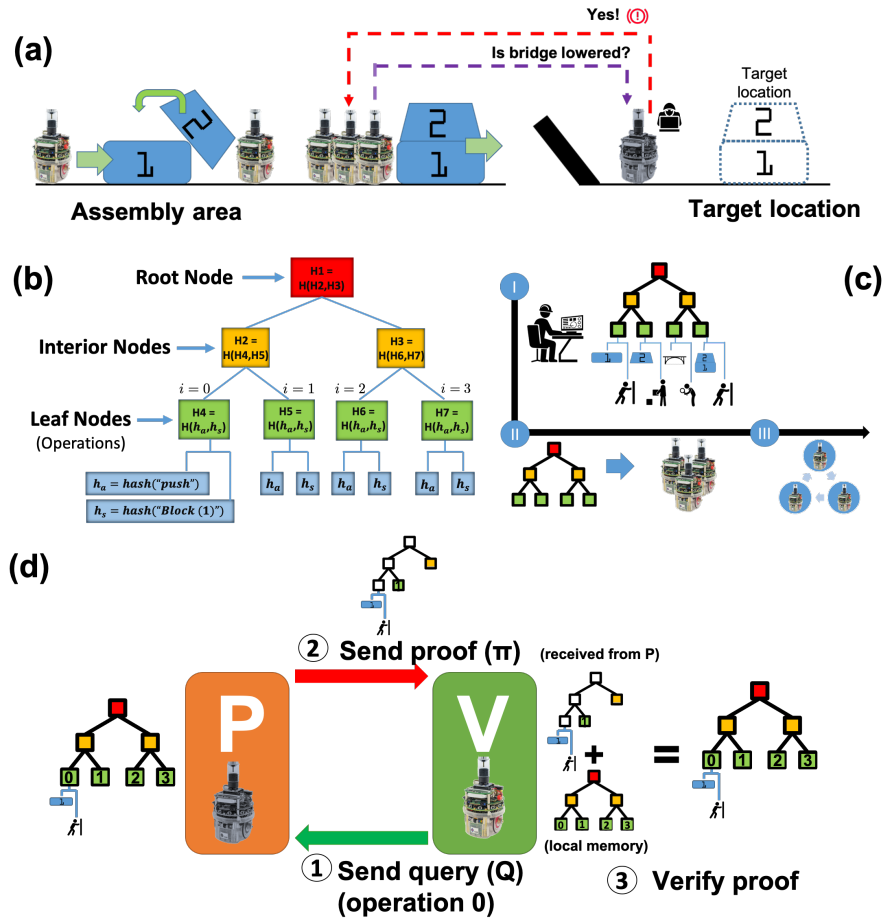


Figure 1: **Towards secure and secret cooperation in swarm robotics missions.** (a) Example of sequential mission (e.g., build a 2-block tower and transport it to a target location) conducted by a robot swarm that can fail if one of the robots involved is tampered with. If a robot spreads erroneous information within the system (e.g., the robot in charge of making the bridge crossable sends a message confirming the bridge is lowered when it is not) the entire mission can fail. (b) Merkle tree (MT) used in this work: a hash-based tree structure where each leaf node stores the hash of an associated operation within the swarm’s mission (i.e., the combination of a robot action hash ( $h_a$ ) and a sensor input hash ( $h_s$ )), while each interior node contains the hash of the combination of its two children. (c) Mission initialization workflow: (I) The swarm’s operator generates a valid MT where all the operations that the swarm needs to perform to complete its mission are included in the correct order. (II) The MT is broadcast to all the robots in the swarm. (III) The mission starts. (d) Workflow for *prover* (P) and *verifier* (V) robots when they exchange queries (Q) and MT “proofs” ( $\pi$ ) to synchronize and complete their corresponding MTs copies.

the overall system reliability and performance. The first surveys of swarm robotics security were presented in (29, 30). These works identified physical capture and tampering with members as significant threats to robot swarms. Physical capture of a robot might not only lead to loss of availability but also to the capture of security credentials or critical details about the swarm operation (31). For instance, if a robot is tampered with and reintroduced into the swarm, an attacker might influence the behaviour of the whole system (25, 27) and eventually hinder the entire mission. Fig. 1(a) depicts an example of such vulnerability, where four operations must be performed in a sequential order by a group of robots: first, block (1) needs to be pushed to an assembly area; second, block (2) has to be placed on top of (1) to assemble a tower-shaped structure; third, the state of a bridge that connects the assembly area and the target location needs to be determined as crossable (i.e., lowered down); fourth, the tower-shaped structure must be pushed to the target location at the other side of the bridge. If a robot is tampered with (e.g., the robot in charge of determining the status of the bridge), it can spread erroneous information within the system (e.g., tell the other robots that the bridge is crossable when it is not), therefore hindering the entire mission. This type of attack would be unique to swarm and multi-robot systems and is particularly critical in situations where robots must share data with other robots in the swarm or with human operators.

In previous swarm robotics work, researchers hard-coded the complete set of rules that trigger the transitions from operation to operation (24, 32, 33) in all robots; that is, each robot in the swarm had a full copy of the information necessary to accomplish its mission. Although this distributed approach is more robust and fault-tolerant than centralized methods, it significantly increases the attack surface (i.e., the total sum of vulnerabilities) for an attacker to figure out the swarm's high-level goals and disrupt the system's behavior (31). Due to these concerns, in this work we aim to find an answer to the following questions: Is there a way to provide the "blueprint" of a robotic mission without disclosing the raw information describing the mission



itself? In other words, is it possible for robots in a swarm to fulfill sequential missions without exposing knowledge about the mission’s high-level objectives? To answer these questions, we present for the first time a cooperation model based on the idea of encapsulating robotic missions into Merkle trees. Our approach allows robots in a swarm to cooperate while minimizing security risks due to issues like physical capture or tampered members. We achieve this by creating a secure data object that is shared by the robots (i.e., a Merkle tree) instead of protecting the communication channel among them (e.g., by encrypting the network connections).

In this paper, we argue that by using Merkle trees (see Fig. 1(b), Materials and Methods, and Supplementary Materials), swarm operators can provide the “blueprint” of the swarm’s objectives without disclosing raw or unprotected data about the mission itself (Fig. 1(c)). More specifically, we introduce for the first time a framework where data verification is separated from data itself. By exchanging cryptographic proofs, robots in the swarm are able to prove to their peers that they know specific pieces of information included in the swarm’s mission and therefore that they are cooperating towards its completion (Fig. 1(d)). This approach was analyzed in simulation and real-robot experiments for two different sequential missions: foraging and maze formation. In both missions, robots were able to cooperate and carry out sequential operations without having explicit knowledge about the mission’s high-level objectives. Our findings show that larger swarms tend to both increase the performance of the system and to diversify the amount of information within the swarm. However, larger numbers of robots as well as longer missions determine a linear increase in the communication requirements of the system. Nevertheless, an analysis of storage usage, communication costs, and computational time for larger-scale missions where the number of operations takes relatively large values reveals that the use of Merkle trees is within reach of current robotic technology.

# Results

## Foraging mission experiments




Figure 2: **Foraging scenario.** (a) Simulated arena used for the foraging mission. The scenario consists of a rectangular area of  $2.5 \times 2.5 \text{ m}^2$  within which robots, tasks (represented as colored cells), and a target area ( $0.5 \times 0.5 \text{ m}^2$ ) located at the center are placed (see Supplementary Materials for a detailed description of the simulation experimental setup). A typical simulated run is displayed in movie M1. (b) Diagram of the e-puck robot together with its sensors layout. E-puck’s size (13 cm tall and 7 cm in diameter) and features make it an ideal platform for simulated and real-world swarm robotics experiments. (c) Real-world equivalent of the scenario shown in (a). In this scenario, colored cells are projected on the arena by using a projector (see Supplementary Materials for a detailed description of the real-robot experimental setup). A typical real-world run is displayed in movie M2.

A set of 100 simulation experiments were carried out to analyze our approach in the foraging scenario shown in Fig. 2(a). The environment consists of a rectangular area segmented in a  $5 \times 5$  grid with a central square representing the target location where e-puck robots (41) (Fig. 2(b)) need to deliver discovered tasks. In Fig. 2(a) tasks are represented by colored cells. Tasks have to be discovered by robots and then delivered to the target location following a pre-

defined sequence of colors. The foraging mission is finished when all tasks in the sequence are delivered in the right order at the target location. In other words, when a robot in the swarm delivers the last task in the sequence. This sequence of colors is encoded by the swarm’s operator as an MT (Fig. 1(b)) with  $n$  operations (i.e., MT length) before the experiment starts (Fig. 1(c)). Therefore, during the mission, robots do not have explicit information about what is the correct sequence of colors to be delivered since they only have the encrypted information included in their MT. In addition, 10 real-robot experiments for a subset of the experimental conditions considered in the simulation experiments were carried out to validate the approach. Fig. 2(c) shows the real-robot foraging environment where tasks are represented by colored cells projected on the arena by using a projector (see Supplementary Materials for the experimental setup of the real-robot experiments).

In the foraging mission, robots search the environment looking for tasks. Once a task is found, robots identify its color and generate  $h_s$ . Then, they generate the value  $h_a$  for all possible actions.<sup>1</sup>  $H_i$  is calculated by hashing  $h_s$  and  $h_a$ :  $H_i = H((h_s, h_a))$  and verified against the current working operation ( $i$ ) of the robot;  $i$  is initialized to 0 at the beginning of the mission and is an index that points to the operation that the robot is currently trying to fulfill in its local MT. In case it is possible to generate a valid proof (see Supplementary Materials for details about MT proofs) with the task’s information ( $\exists \pi_i, H_i$ ), the robot “grabs” the task color and delivers it in the target location. If not, the robot keeps wandering. When an operation is completed successfully, the robot changes the status of  $i$  as completed in its local MT. Then, the robot increases  $i$ :  $i = i + 1$  for  $i \in (0 \leq i \leq n - 1)$ . It is important to note that during the experiments, robots can exchange information with their peers (e.g.,  $i, \pi$ ). If two robots are within communication distance, they can exchange their  $i$  values. If there is a disparity in the values (i.e., one  $i$  is lower than the other) the robot with the lower  $i$  will become a *verifier* ( $V$ ),

---

<sup>1</sup>In the considered foraging mission,  $h_a$  is restricted to the “carry to target” action.

and will send queries ( $Q$ ) to the robot with greater  $i$ , the *prover* ( $P$ ), asking for proofs ( $\pi$ ) about the missing operations as depicted in Fig. 1(d). This process is conducted until the  $i$  value is the same for both robots. By using this method, robots can synchronize their own MT copies and therefore cooperate towards the fulfillment of the swarm’s mission (see Supplementary Materials for robot’s interaction and finite state machine diagrams).

Next, we study how the performance of our approach (i.e., how fast and reliably a particular mission is carried out) changes when varying the MT length and the swarm sizes. We present results obtained in simulation and real-robot experiments where the MT length  $n$  is varied in the set  $[2, 4, 5, 6, 7, 8]$  (simulation) and  $[2, 4, 8]$  (real) and the swarm sizes  $R_n$  in the interval  $1 \leq R_n \leq 16$ . In addition, the amount of communication (i.e., the total amount of data robots exchanged during the mission), and information diversity (i.e., how widely spread information is within the swarm) metrics are calculated and presented. Details about these analysis metrics and statistical methods used can be found in the Materials and Methods section.

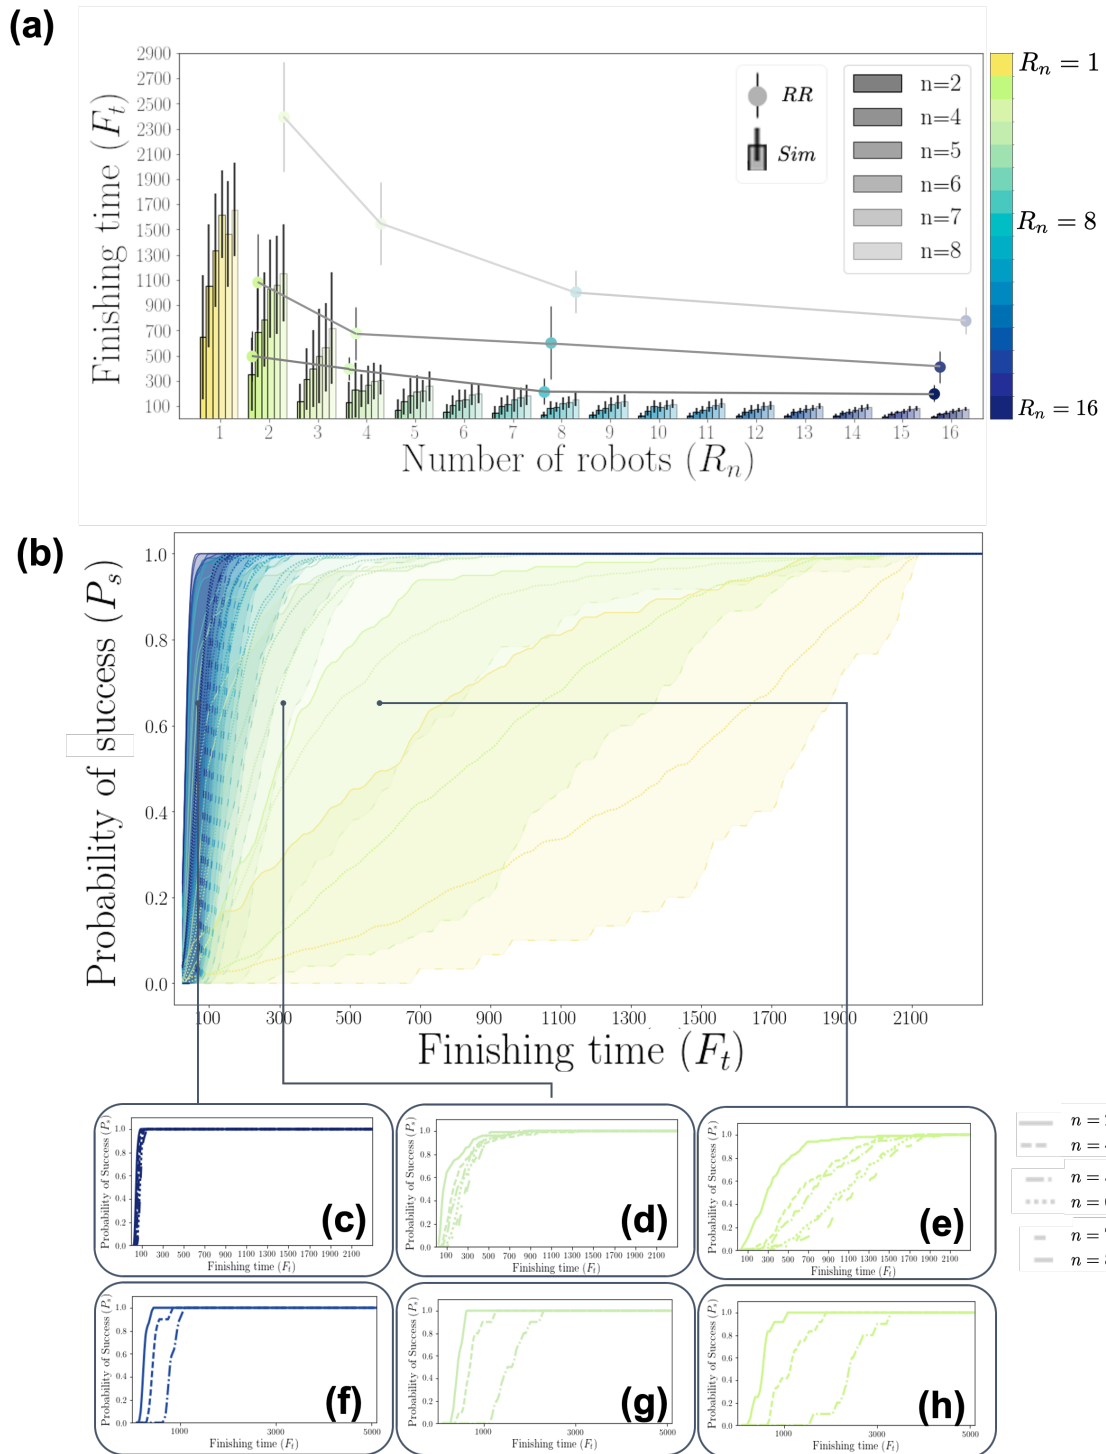


Figure 3: **Finishing times and probability of success for the foraging mission.** (a) Average  $F_t$  (in seconds) and its standard deviations for different MT lengths ( $n$ ) and robot swarm sizes ( $R_n$ ) for both simulation (bars) and real-robot experiments (circles). (b) Probability of success ( $P_s$ ) for all simulation configurations shown in (a) for the execution of the foraging mission with  $R_n$  in the interval  $1 \leq R_n \leq 16$ . For each of the  $R_n$  values a lower-bound (solid line representing  $n = 2$ ), mean of the evaluated interval (dot-dashed line), and the upper-bound (loosely dashed line representing  $n = 8$ ) were included. (c-e) Detailed plots for the simulation experiments depicting the whole  $n$  set for (c)  $R_n = 16$ , (d)  $R_n = 4$ , (e)  $R_n = 2$ . (f-g) Detailed plots for the real-robot experiments depicting the  $n = [2, 4, 8]$  subset for (f)  $R_n = 16$ , (g)  $R_n = 4$ , (h)  $R_n = 2$ .

Fig. 3(a) shows the finishing time ( $F_t$ ) and its standard deviation for several MT length configurations ( $n$ ) and robot swarm sizes ( $R_n$ ) for both simulation and real-robot experiments. According to Fig. 3(a),  $F_t$  decreases when robots are added to the swarm, regardless of the length of the MT. These results also suggest that once a certain number of robots is present ( $R_n \geq 8$ ), the length of the MT ( $n$ ) has little impact on the  $F_t$  of the system. Fig. 3(b) shows how the probability of success ( $P_s$ ) curves change when changing the number  $n$  of tasks for all simulation configurations presented in Fig. 3(a). One can also see that adding robots increases  $P_s$  since the curves become steeper and converge to higher values sooner. However, these results also suggest that as we increase  $n$  (the mission becomes longer),  $P_s$  reaches higher values later (especially for  $R_n \leq 3$ ). This effect can be seen in both simulation Fig. 3(c-e) and real-robot experiments Fig. 3(f-h). It is important to note that  $P_s$  does reach 1 (the maximum value) in all configurations as experiments were run up to a Time Cap (TC), which was never reached (see Supplementary Materials for a detailed description of the experimental setup).

Fig. 4 shows average results and standard deviations for the amount of communication ( $AC$ ) in KB and information diversity measured by Shannon’s equitability index ( $I_e$ ) for different simulation and real-robot experiments. In particular, Fig. 4(a) reports simulation results for the  $R_n$  ( $1 \leq R_n \leq 16$ ) and  $n \in \{2, 4, 6, 8\}$  configurations, while Fig. 4(b) reports results of real robot experiments for  $R_n = [2, 4, 8, 16]$  and  $n \in \{2, 4, 8\}$  configurations. In both figures,  $I_e \in [0, 1]$  gives information on the variability in the number of operations performed by each robot in the swarm during the mission. Lower values indicate more uneven distributions (e.g., one single robot completing all operations) while higher values indicate more uniform distributions (e.g., all robots completing the same number of operations). Figs. 4(a-b) also show that  $AC$  increases linearly with  $R_n$  since there are more robots exchanging MT proofs. Moreover, MTs with larger  $n$  values make the proofs robots exchange “heavier” (i.e.,  $\pi$  is longer), which also contributes to increase  $AC$ . However, larger  $R_n$  values tend to increase information evenness

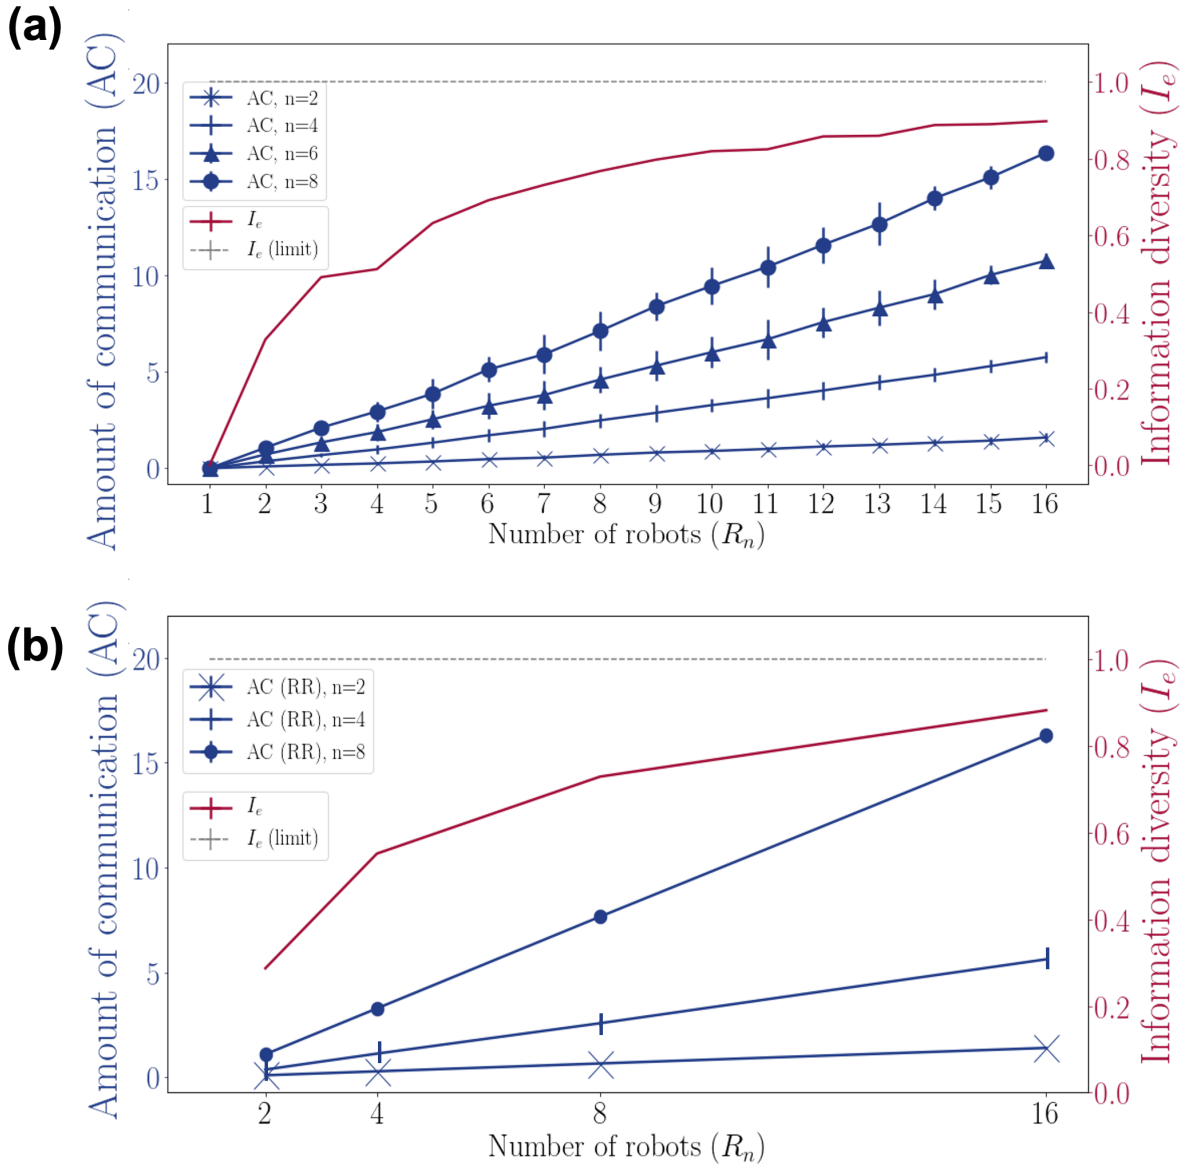


Figure 4: **Amount of communication (AC) and information diversity ( $I_e$ ) metrics for the foraging mission.** The AC and  $I_e$  metrics are depicted in blue and red colors, respectively, for different  $R_n$  and  $n$  values used in simulation (a) and real-robot (b) experiments. The upper-bound limit for the  $I_e$  metric is depicted in both figures with a dashed grey line. Averaged results and standard deviations in both scenarios suggest a direct relationship between the two metrics in the foraging mission.

( $I_e$ ) in the swarm. In other words, sequential operations are more evenly distributed. Results for both simulation and real-robot experiments suggests that, as we increase  $R_n$ , the number of operations fulfilled by the robots tends to become more uniform. Fig. 4 also suggests that information diversity might need a really large  $R_n$  value to converge to 1 (i.e., to complete “evenness”).

## Maze formation mission experiments

In the second experiment, we gave the robot swarm a maze formation mission (Fig. 5(a)-(c)) where we conducted 100 simulations and 15 real-robot experiments for each considered configuration. Fig. 5(a) represents the “blueprint” of a  $5 \times 5$  maze where 0 represents an empty space, 1 a wall, and \* and @ the entrance and the exit of the maze, respectively. As in the foraging mission, robots first wander around the arena. However, instead of looking for colored cells, in this mission robots search for (x,y) maze coordinates. By knowing the cell dimensions ( $0.5 \times 0.5 \text{ m}^2$ ) robots can calculate the (x,y) coordinates of cells they are located at. Every time a robot discovers a new (x,y) coordinate pair (e.g., (1,1)), it uses it to generate  $h_s$ . Then, it generates the value  $h_a$  for all possible actions.<sup>2</sup> Like in the foraging mission,  $H_i$  is calculated by hashing  $h_s$  and  $h_a$ :  $H_i = H((h_s, h_a))$  and verified against the current  $i$  of the robot. In case it is possible to generate a valid proof  $\pi$  with this information ( $\exists \pi_{(i, H_i)}$ ) the robot finds the center of the cell and stops there (Fig. 5(b-c)). If not, the robot keeps wandering. In the same way as in the foraging mission, robots avoid already completed operations (in this case, they avoid stopping in already occupied cells) by receiving the proof that  $i$  was already completed. Additional information about robot’s behavior in the maze formation mission can be found in the Supplementary Materials.

As with the foraging experiment, the maze formation mission is finished once all operations

---

<sup>2</sup>In the considered maze formation mission,  $h_a$  is restricted to the “stop” action.






Figure 5: **Maze formation scenario.** (a)  $5 \times 5$  matrix used to represent a maze. Four different elements are included in the array: 0 (black) represents an empty space, 1 (red) represents a wall, @ and \* (blue) represent the entrance and the exit of the maze, respectively. (b) The maze depicted in (a) built by a swarm of simulated robots (a typical simulation run is displayed in movie M3) where the path between entrance and exist is depicted with a dashed line. (c) The maze depicted in (b) built by a swarm of real e-puck robots (a typical real-robot run is displayed in movie M4). (d) Average  $F_t$  (in seconds) and its standard deviation for the maze formation mission with different robot swarm sizes for both simulation (blue) and real-robot (orange) experiments. All results were obtained with a 16 operations ( $n = 16$ ) maze as shown in (a-c). (e) Amount of communication ( $AC$ ) and information diversity ( $I_e$ ) metrics in blue and red colors, respectively, for both simulation (circles) and real robots (triangles and squares) and for different  $R_n$  values for the maze formation mission. Upper-bound limit values for both  $AC$  (star) and  $I_e$  (line) are depicted with dashed black lines respectively. Average results suggest no relationship between the two metrics.

have been completed. In this maze formation mission,  $n$  was fixed to 16 in order to match the number of cells where the value 1 is present in Fig. 5(a). In addition, we explored a wider  $R_n$  range ( $16 \leq R_n \leq 28$ ). Fig. 5(d) shows average  $F_t$  and standard deviations for the maze formation mission for both simulation and real-robot experiments: as it was the case in the corresponding foraging experiment (Fig. 3(a)), larger  $R_n$  values reduce  $F_t$ . However, beyond a certain  $R_n$  value ( $R_n \geq 24$ ), the impact on  $F_t$  becomes very small. Fig. 5(e) shows the  $AC$  and  $I_e$  metrics for the maze formation mission. Fig. 5(e) follows the same pattern as Fig. 4:  $AC$  increases linearly with  $R_n$ . Fig. 5(e) includes the  $AC$  upper limit following equation 1 (see Statistical methods). According to Fig. 5(e), both simulation and real-robot  $AC$  fall under the limit and suggest that real-robot experiments require a lower amount of communication than their simulation counterparts. Finally, Fig. 5 (e) shows that in this scenario complete “evenness” of information (i.e.,  $I_e = 1$ ) is achieved.

## Experiments with larger-dimension missions




Figure 6: **LEGO<sup>®</sup> models with their corresponding piece count.** LEGO<sup>®</sup> models are a good example of complex sequential missions that could be encapsulated in Merkle trees.

Encouraged by these results, we found appropriate to analyze the feasibility of our approach in complex missions where the number of operations takes relatively large values. Fig. 6 shows

Table 1: Memory, maximum amount of communication per robot ( $AC/R_n$ ), and computation times for the models depicted in Fig. 6. Time measures (in seconds) are averages (standard deviations in brackets) over 100 runs using the Gumstix Overo board, the on-board computer mounted on the e-puck robots.

Model in Fig. 6	Memory (KB)	$AC/R_n$ (MB)	Average computation times* (s)
(a) Millennium Falcon (7541 pieces)	235 KB	3.39 MB	<b>G:</b> 0.35 (0.001) <b>P:</b> 0.0017 (0.0002) <b>V:</b> 0.002 ( $7.96 \cdot 10^{-5}$ )
(b) Taj Mahal (5923 pieces)	185 KB	2.54 MB	<b>G:</b> 0.28 (0.004) <b>P:</b> 0.0016 ( $7.40 \cdot 10^{-5}$ ) <b>V:</b> 0.002 (0.0003)
(c) Bugatti Chiron (3599 pieces)	112 KB	1.46 MB	<b>G:</b> 0.16 (0.001) <b>P:</b> 0.0015 ( $3.61 \cdot 10^{-5}$ ) <b>V:</b> 0.0018 ( $9.37 \cdot 10^{-5}$ )

\* **G:** time to generate the complete Merkle tree; **P:** time for the generation of a proof; **V:** time for the verification of a proof.

different LEGO<sup>®</sup> models where a sequential set of operations is required to achieve the final outcome (i.e., build the replica). These models<sup>3</sup> are good projections of the missions presented previously, especially since  $n$  takes a relatively large value. Due to the possibility of accurately calculating the upper limit of  $AC$  required to make all robots complete their MTs (see Statistical methods, Eq. 1) as well as the overall size of the MT stored by robots, we can compute Fig. 6's corresponding MTs and measure their memory, maximum amount of communication per robot ( $AC/R_n$ ) requirements. We also computed the average computation time to generate the complete Merkle tree (G), the time for the generation of a proof (P), and for the verification of a proof (V). Results for the aforementioned models are given in Table 1.

<sup>3</sup>Three of the replicas with the largest piece count according to the current LEGO<sup>®</sup> catalog: <https://shop.lego.com/en-US/>

## Discussion

We showed that two of the main MT properties (i.e., correctness and security) open a new path towards secure and secret cooperation in robot swarms. Regarding the security aspect, by using this approach, in order to cooperate, the robots in a swarm are required to prove to their peers that they fulfilled certain actions or that they know or own some particular information (i.e., proof-of-ownership (34)), rather than merely rely on information received from other robots (sensor data, votes, etc.). This approach makes robots resistant against potential threats such as tampering attacks since any change in the operation's data (i.e.,  $s$ ,  $a$ ) will necessarily change the proof's outcome. Regarding the secrecy component, with the use of MTs robot swarms are now able to separate the mission data from its verification. This allows robots to verify that an operation was carried out by a member of the swarm without knowing what this operation entailed or which robot took part in its completion. This makes physical capture attacks inefficient since individual robots might not have enough raw or unprotected information to describe the high-level swarm's missions and goals, especially in large systems. However, this does not prevent robots in a swarm from cooperating to fulfill complex missions since robots can still prove to their peers that certain operations were identified and completed.

This approach was tested in two different scenarios: a foraging and a maze formation mission. In the foraging case, results suggest that increasing the swarm size has a positive impact on the performance of the system: larger swarm size have lower finishing time ( $F_t$ ) and a higher probability of success ( $P_s$ ). Results also show that the amount of communication ( $AC$ ) grows linearly with the swarm size ( $R_n$ ), which in extreme situations (e.g., very large swarms) could be detrimental to the system since individual robots might not be able to cope with the bandwidth requirements. In contrast, increasing  $R_n$  has the positive effect of increasing the information diversity ( $I_e$ ) since we are increasing the probability of reaching more uniform distributions of

completed operations within the swarm.

In the maze formation mission, where  $R_n$  and  $n$  take larger values, results also suggest that  $R_n$  maintains an inverse relationship with  $F_t$ . In the maze formation mission,  $I_e$  is maximized (i.e.,  $I_e = 1$ ), which is possible since in this mission, when robots find a cell where they can generate a valid  $\pi$  proof, they stop at its center, in this way making robots capable of fulfilling only one operation per mission, in contrast to the foraging scenario, where one robot might be able to complete several operations. Even though  $AC$  grows linearly with  $R_n$ , this still does not represent a challenging situation for the swarm (e.g., 72 and 70 KB are exchanged in a 28 robots system for both simulation and real-robot experiments). Counterintuitively, real-robot experiments require a lower amount of  $AC$  to complete the maze-formation mission than their simulation counterparts. This effect is due to the incapability of the simulation engine to take into account network lag: delays in the network produce situations where robots are able to find and complete the mission, while other robots in the swarm are still trying to synchronize on previous completed operations. Finally, it is interesting to emphasize that, due to these properties, robot swarms can complete complex missions such as the maze formation one without the means to infer high-level details such as where the entrance or the exit might be located.

Finally, Table 1 shows that neither the memory, nor the communication requirements per robot, nor the average computation times of the corresponding MTs are out of reach of current commodity hardware (e.g., Overo Gumstix) and therefore our approach is feasible for current robot platforms. It is important to note that more than 99% of the computation time is taken by the generation (G) of the MT that only takes place at the beginning of the mission, while the proof assembly (P) and validation (V) take a nearly insignificant amount of time.

## Limitations and future work

### The complete set of operations (and its order) must be known at design time

It is important to note that, as described previously, the current approach needs a swarm operator to design and encode the complete sequence of operations included in the MT before the mission starts. Therefore, any deviation from the original sequence included in the MT might be difficult to adapt to. **Possible improvement:** it is possible to make incremental updates to the mission and encapsulate the outdated MT within a newer version that takes into account all necessary changes. In that situation, a new MT root might be necessary to reflect the updated mission sequence. However, the most critical point for this solution is to trust the source that distributes the new root node hash. Fortunately, new methods such as blockchain-based communication are showing a promising way to make robots rely on trusted sources (27, 28).

### Generation of a valid MT proof does not imply execution

Another limitation of the system is that even if a robot discovers a valid combination of  $h_a$  and  $h_s$  values, thus being able to generate a valid MT proof, there is no guarantee the robot fulfills the corresponding operation in the physical world (e.g., deliver discovered colors to a target location or stop at the right coordinates). **Possible improvement:** it is possible to add additional variables to the operation encoding process, for instance, forcing the robot to do the right action  $a$ , with the right sensor input  $s$ , at the right location  $l$  such as  $H_i = H((h_s, h_a, h_l))$ . Another possible way to tackle this problem is the addition of robot “validators”, which verify that every operation claimed by the robots is completed.

Perhaps the most promising direction for future research is the possibility to implement missions where heterogeneous swarms with different sensing, computation, and actuation capabilities can cooperate and collaborate. Along these lines, the increasing amount of attention that swarm robotics is gaining envisages a future where different swarms run by different types of institutions (e.g., private, public) can co-exist in a same location (e.g., urban and disaster

areas, battlefield). The method described in this research offers a path to make different swarms cooperate without exposing sensitive data about their internal processes, goals, and organizations. Finally, a new type of service where the robot commands are encapsulated within an MT and whose output is the cryptographic proof that the job has been completed is now a possibility. In the last section of the supplementary materials of this work is described an early-stage example in this direction where a web-based marketplace allows a robot swarm to form shapes on-demand. In the example, potential customers upload an MT describing the desired shape and send it to a smart-contract in the Ethereum blockchain, which returns a proof when the work is being completed by the robots.

## Materials and Methods

### Merkle trees for robot swarms

A Merkle tree (MT) (37) is a hash-based tree structure where data is not stored in the inner nodes but in the leaves (Fig. 1(b)). Every leaf node (i.e., operation) encapsulates the combined hash of two hashes:  $h_s$  (hash of the robot sensor's input) and  $h_a$  (hash of the robot's action). These two hashes describe an atomic step within the swarm's high-level mission. For instance, the hash of the action "push" ( $h_a$ ) and the hash of the sensor input "block (1)" ( $h_s$ ) are included in the first operation by using the hash function  $H: H(h_a, h_s)$ . In this work, we assume that an external entity (e.g., the swarm's operator) encodes a valid MT before the swarm robotics mission can take place. This MT contains all needed information about the operations that the swarm has to execute to fulfill the mission in the correct order (Fig. 1(c) - I). The resultant MT is broadcast to all the robots (Fig. 1(c) - II) before the mission starts (Fig. 1(c) - III). Since MTs are an encrypted data structure, robots do not have direct access to the raw information concerning the mission, but only to the hash values comprising the MT are known.

At the beginning of the mission, all robots aim to fulfill the first operation. However, only

one will succeed in doing so. The successful robot, first, wanders around the environment; second, when it perceives a new sensor input (e.g., block (1)) it generates a hash ( $h_s$ ) from it; third, it generates another hash ( $h_a$ ) for an action it can perform (e.g., “push”); fourth, it combines  $h_s$  and  $h_a$  and rehashes them to obtain the hash string  $H_{(h_s, h_a)}$ ; fifth, it matches  $H_{(h_s, h_a)}$  against the hash value that corresponds to the first operation of the MT. At this point, the robot executes the resulting command (e.g., “push block (1)”). The robot that fulfills this operation, can regard the operation as completed. Once the operation is completed, the (*prover*) robot has the ability to generate cryptographic “proofs” ( $\pi$ ) that demonstrate to other (*verifier*) robots that it actually carried it out that operation because it has information corresponding to that operation that it could not have otherwise (i.e.,  $h_s$  and  $h_a$ ). In this case, the proof  $\pi$  is a list of hashes that *verifier* robots can use to calculate the MT root node hash (shared by all robots in the swarm) by using  $h_a$  and  $h_s$  as inputs. These proofs can be sent and received by robots in the swarm, allowing them to synchronize, check, update, and complete their corresponding MT until the whole mission is completed. It is important to note that these proofs can be exchanged without disclosing any raw or unprotected data about the mission itself. The typical workflow between *prover* (P) and *verifier* (V) follows the pattern depicted in Fig. 1(d): (1) V sends a query Q to P regarding a particular operation in the MT (e.g., operation 0). (2) P sends a proof  $\pi$  (list of hashes) that demonstrates knowledge about the requested operation. (3) V verifies the proof by computing in a bottom-up fashion the received information. The proof is regarded as valid if V can generate the root node hash it keeps in its memory (shared by P and V) by using  $\pi$ . More details about MT structure, proofs, and data encapsulation are provided in the Supplementary Materials.



## Statistical methods

To evaluate and analyze our approach we rely on three metrics:

**Performance.** These measures show the speed and reliability with which a mission is carried out. In this paper, we use the mission’s finishing time  $F_t$  to measure the amount of time required to fulfill the swarm’s mission. In addition, we use an estimate  $P_s$  of the probability that the system attains its target objective in an amount of time  $\tau$  (38). Formally, let  $j \in \{1, \dots, k\}$  be the index of an experiment,  $r_j$  be the run time of experiment  $j$ , and experiments where  $r_j < TC$  are considered successful experiments. The estimate  $P_s$  of the probability of success of the system over time (up to TC) is defined as  $P_s(\tau \leq t) = \{j | r_j \leq t\} / k$ .

**Amount of communication.** The amount of communication ( $AC$ ) is computed by multiplying the number of times the *prover-verifier* workflow (Fig. 1(d)) takes place during the mission by the size (in bytes) of the proof ( $\pi$ ) the robots exchanged. When the MT is perfectly balanced (i.e., the number  $m$  of leaves is a power of 2), the size of the proof  $\pi$  is  $\log_2(m) + 2$ : the number of hashes to reach the root node plus the  $h_s, h_a$  hashes. In missions with  $n$  operations where all robots need to complete their MTs, an upper limit  $AC_{ul}$  of the  $AC$  metric can be calculated with the following equation:

$$AC_{ul} = P_n \cdot P_l \cdot |H| \quad (1)$$

where  $P_n = ((R_n - 1) \cdot n)$  is the total number of proofs exchanged,  $P_l = \log_2(\hat{n}) + 2$  is the length of the proof,  $\hat{n}$  is the smallest power of 2 that is greater or equal to  $n$ , and  $|H|$  is the size (in bytes) of the hash function used. The hash function used in this work (SHA256) has a hash size ( $|H|$ ) of 32 bytes. It is important to note that the hash size can be increased (e.g., SHA3-512, SHA3-1024) in order to achieve improved security.

**Information diversity.** In this research, robots are only in contact with the raw sensor

and action information from the operations they carry out themselves (i.e., input data for the  $h_a$  and  $h_s$  hashes). However, in swarm robotics applications, it is difficult to fully ensure a uniform distribution of operations or tasks among the robots. In this scenario, certain robots might be able to accumulate raw or unprotected information that could be stolen if they are subject to attacks (e.g., physical capture). To analyze this phenomenon, we introduce Shannon’s equitability ( $I_e$ ) to measure “evenness”, that is, to measure how widely spread raw information is within the swarm.

Shannon’s equitability ( $I_e$ ) can be calculated by dividing Shannon’s index  $I$  by  $I_{max}$ :

$$I_e = \frac{I}{I_{max}} \quad (2)$$

where Shannon’s index  $I = -\sum_{i=1}^S p_i \ln p_i$  (i.e., Shannon’s entropy (39)) is a mathematical measurement used to characterize diversity in a community. In this case,  $S$  represents the total number of robots that took part in the mission ( $R_n$ ),  $p_i$  is the percentage of operations robot  $i$  conducted compared to the mission’s length ( $n$ ), and  $I_{max} = \ln S$ .  $I_e$  assumes a value between 0 and 1. Lower values indicate more uneven distributions while higher values indicate more uniform distributions. In this case, 1 represents complete “evenness”: all robots carried out the same number of operations and therefore were exposed to the same amount of raw information.

## References

1. M. Dorigo, M. Birattari, M. Brambilla, Swarm Robotics, *Scholarpedia*, 1463, (2014).
2. R. D’Andrea, Guest Editorial: A Revolution in the Warehouse: A Retrospective on Kiva Systems and the Grand Challenges Ahead, *IEEE Transactions on Automation Science and Engineering*, 638 (2012).

3. T. Blender, T. Buchner, B. Fernandez, B. Pichlmaier, C. Schlegel, Managing a Mobile Agricultural Robot Swarm for a seeding task, *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, 6879–6886 (2016).
4. D. Albani, J. IJsselmuiden, R. Haken, V. Trianni, Monitoring and mapping with robot swarms for agricultural applications, *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 1–6 (2017).
5. Y. Chen, Z. Du, M. García-Acosta, Robot As a Service in Cloud Computing, *Proceedings of the 2010 Fifth IEEE International Symposium on Service Oriented System Engineering, SOSE '10* (IEEE Computer Society, Washington, DC, USA, 2010), 151–158 (2010).
6. H. He, S. Kamburugamuve, G. Fox, W. Zhao, Cloud based Real-time Multi-robot Collision Avoidance for Swarm Robotics, *International Journal of Grid and Distributed Computing* 9, 339-358 (2016).
7. M. H. A. Majid, M. R. Arshad, Design of an autonomous surface vehicle (ASV) for swarming application, *2016 IEEE/OES Autonomous Underwater Vehicles (AUV)*, 230–235 (2016).
8. Y. Mulgaonkar, A. Makineni, L. Guerrero-Bonilla, V. Kumar, Robust Aerial Robot Swarms Without Collision Avoidance, *IEEE Robotics and Automation Letters* 3, 596–603 (2018).
9. D. Rus, M. T. Tolley, Design, fabrication and control of origami robots, *Nature Reviews Materials* 3, 101–112 (2018).
10. E. Castello Ferrer, S. Y. Hanay, Demo: A Low-cost, Highly Customizable Robotic Platform for Testing Mobile Sensor Networks, *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc '15* (ACM, New York, NY, USA, 2015), 411–412 (2015).

11. N. Oxman, J. Duro-Royo, S. Keating, B. Peters, E. Tsai, Towards Robotic Swarm Printing, *Archit Design*, 84: 108-115 (2014).
12. D. Chaudhari, S. A. Bhosale, A review on management of logistics using swarm robotics, *2016 International Conference on Communication and Signal Processing (ICCSP)*, 0371–0373 (2016).
13. A. Heinla, A. Martinson, K. Volkov, A. Macks, L. Roberts, I. Mandre, M. Liivik, T. Liivik, I. Liivik, Method and system for autonomous or semi-autonomous delivery, U.S. Patent 0232839A1, April, 9, 2018.
14. A. L. Alfeo, E. Castello Ferrer, Y. Lizarribar Carrillo, A. Grignard, L. Alonso Pastor, D. T. Sleeper, M. G.C.A. Cimino, B. Lepri, G. Vaglini, K. Larson, M. Dorigo, A. Pentland, Urban Swarms: A new approach for autonomous waste management, *2019 International Conference on Robotics and Automation (ICRA)*, 4233–4240 (2019).
15. G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, B. J. Nelson, B. Scassellati, M. Taddeo, R. Taylor, M. Veloso, Z. L. Wang, R. Wood, The grand challenges of Science Robotics, *Science Robotics*, **3** (2018).
16. P. Scharre, Robotics on the Battlefield Part II: The Coming Swarm, *Center for New American Security*, 1–4 (2014).
17. S. Vivek, D. Yanni, P. J. Yunker, J. L. Silverberg, Cyberphysical risks of hacked internet-connected vehicles, *Phys. Rev. E*, **100**, 012316 (2019).
18. R. N. Akram, K. Markantonakis, K. Mayes, O. Habachi, D. Sauveron, A. Steyven, S. Chaumette, Security, privacy and safety evaluation of dynamic and static fleets of drones, *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, 1–12 (2017).

19. S. Gil, S. Kumar, M. Mazumder, D. Katabi, D. Rus, Guaranteeing spoof-resilient multi-robot networks, *Autonomous Robots* **41**, 1383–1400 (2017).
20. A. Prorok, V. Kumar, A macroscopic privacy model for heterogeneous robot swarms, *ANTS 2016. Lecture Notes in Computer Science* **9882**, 15–27 (2016).
21. A. Prorok, V. Kumar, Towards Differentially Private Aggregation of Heterogeneous Robots, *Distributed Autonomous Robotic Systems, Springer Proceedings in Advanced Robotics* **6**, 587–601 (2018).
22. A. L. Christensen, R. O’Grady, M. Birattari, M. Dorigo, Exogenous Fault Detection in a Collective Robotic Task, *Advances in Artificial Life, Proceedings of ECAL 2007* **4648**, 555–564 (2007).
23. A. L. Christensen, R. O’Grady, M. Dorigo, Synchronization and Fault Detection in Autonomous Robots, *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-2008)*, 4139–4140 (2008).
24. S. Nouyan, R. Groß, M. Bonani, F. Mondada, M. Dorigo, Teamwork in Self-Organized Robot Colonies, *IEEE Transactions on Evolutionary Computation* **13**, 695–711 (2009).
25. A. G. Millard, J. Timmis, A. F. T. Winfield, Towards Exogenous Fault Detection in Swarm Robotic Systems, *Towards Autonomous Robotic Systems. TAROS 2013. Lecture Notes in Computer Science*, 429–430 (2013).
26. J. D. Bjercknes, A. F. T. Winfield, On Fault Tolerance and Scalability of Swarm Robotic Systems, *Distributed Autonomous Robotic Systems: The 10th International Symposium* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013), 431–444 (2013).

27. V. Strobel, E. Castello Ferrer, M. Dorigo, Managing Byzantine Robots via Blockchain Technology in a Swarm Robotics Collective Decision Making Scenario, *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 541–549 (2018).
28. V. Strobel, E. Castello Ferrer, M. Dorigo, Blockchain Technology Secures Robot Swarms: A Comparison of Consensus Protocols and Their Resilience to Byzantine Robots, *Frontiers in Robotics and AI* **7**, 54 (2020).
29. A. F. Winfield, J. Nembrini, Safety in numbers: fault-tolerance in robot swarms, *International Journal of Modelling, Identification and Control* **1**, 30–37 (2006).
30. F. Higgins, A. Tomlinson, K. M. Martin, Survey on Security Challenges for Swarm Robotics, *IEEE 2009 Fifth International Conference on Autonomic and Autonomous Systems*, 307–312 (2009).
31. I. Sargeant, A. Tomlinson, Review of Potential Attacks on Robotic Swarms, *Proceedings of 2016 SAI Intelligent Systems Conference (IntelliSys)*, *Lecture Notes in Networks and Systems* **16**, 628–646 (2018).
32. M. Dorigo, D. Floreano, L.M. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Biorattari, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A.L. Christensen, A. Decugniere, G. Di Caro, F. Ducatelle, E. Ferrante, A. Forster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. Montes de Oca, R. O’ Grady, C. Pinciroli, G. Pini, P. Retornaz, J. Roberts, V. Sperati, T. Stirling, T. Stutzle, V. Trianni, E. Tuci, A.E. Turgut, F. Vaussard, Swarmanoid: A Novel Concept for the Study of Heterogeneous Robotic Swarms, *IEEE Robotics Automation Magazine* **20**, 60–71 (2013).

33. J. Werfel, K. Petersen, R. Nagpal, Designing Collective Behavior in a Termite-Inspired Robot Construction Team, *Science* **343**, 754–758 (2014).
34. S. Halevi, D. Harnik, B. Pinkas, A. Shulman-Peleg, Proofs of Ownership in Remote Storage Systems. *Proceedings of the 18th ACM Conference on Computer and Communications Security*, 491–500 (2011).
35. E. C. Ferrer, O. Rudovic, T. Hardjono, A. Pentland, Robochain: A secure data-sharing framework for human-robot interaction. *arXiv preprint arXiv:1802.04480* (2018).
36. V. Strobel, M. Dorigo, Blockchain technology for robot swarms: A shared knowledge and reputation management system for collective estimation, *Swarm Intelligence – Proceedings of ANTS 2018 – Eleventh International Conference* **11172**, 425–426 (2018).
37. R. C. Merkle, A Digital Signature Based on a Conventional Encryption Function, *Advances in Cryptology — CRYPTO '87*, 369–378 (1988).
38. L. Garattoni, M. Birattari, Autonomous task sequencing in a robot swarm, *Science Robotics* **3**, (2018).
39. C. E. Shannon, A Mathematical Theory of Communication, *Bell System Technical Journal* **27**, 379—423 (1948).
40. C. Pinciroli, V. Trianni, R. O’Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L.C. Gambardella, M. Dorigo, ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems, *Swarm Intelligence* **6**, 271–295 (2012).
41. F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J.C. Zufferey, D. Floreano, A. Martinoli, The e-puck, a Robot Designed for Education in Engineer-

- ing, *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions* **1**, 59–65 (2009).
42. J.A. Marvel, R. Bostelman, J. Falco, Multi-Robot Assembly Strategies and Metrics, *ACM Comput. Surv.* **51**, 1-32, (2018).
43. A. Farinelli, N. Boscolo, E. Zanutto, E. Pagello, Advanced approaches for multi-robot coordination in logistic scenarios, *Robotics and Autonomous Systems* **90**, 34–44 (2017).
44. I. Roman-Ballesteros and C. F. Pfeiffer, A Framework for Cooperative Multi-Robot Surveillance Tasks, *Electronics, Robotics and Automotive Mechanics Conference (CERMA'06)* **2**, 163–170 (2006).
45. A. Stranieri, A.E. Turgut, M. Salvaro, G. Francesca, A. Reina, M. Dorigo, and M. Birattari, IRIDIA's Arena Tracking System, (*Tech. rep. TR/IRIDIA/2013-013*, IRIDIA, Université libre de Bruxelles, 2013), 1–7 (2013).
46. M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, B. Wheeler, A. Ng, ROS: an open-source Robot Operating System, *ICRA Workshop on Open Source Software* **3**, (2009).
47. A. Tiderko, F. Hoeller, T. Röhling, The ROS Multimaster Extension for Simplified Deployment of Multi-Robot Systems, *Robot Operating System (ROS): The Complete Reference* **1**, 629–650, (2016).
48. A. Drogoul, E. Amouroux, P. Caillou, B. Gaudou, A. Grignard, N. Marilleau, P. Taillandier, M. Vavasseur, D.A. Vo, J.D. Zucker, GAMA: A Spatially Explicit, Multi-level, Agent-Based Modeling and Simulation Platform, *Advances on Practical Applications of Agents and Multi-Agent Systems* **1**, 271–274 (2013).



49. A. D. Maynard, Navigating the fourth industrial revolution, *Nature nanotechnology* **10**, 1005–1006 (2015).
50. W. E. Forum, The future of jobs: Employment, skills and workforce strategy for the fourth industrial revolution. *Global Challenge Insight Report, World Economic Forum*, 1–167 (2016).
51. K. Schwab, The fourth industrial revolution, *World Economic Forum*, (2017).
52. E. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, Industrial internet of things: Challenges, opportunities, and directions, *IEEE Transactions on Industrial Informatics* **14**, 4724–4734 (2018).
53. S. Park, N. Crespi, H. Park, S.-H. Kim, Iot routing architecture with autonomous systems of things, *2014 IEEE World Forum on Internet of Things (WF-IoT)* **1**, 442–445 (2014).
54. G. A. Bekey, Autonomous robots: from biological inspiration to implementation and control, *MIT press*, (2005).
55. F. Fahimi, Autonomous robots: Modeling, Path Planning, and Control, *Springer*, (2009).
56. J. Lee, B. Bagheri, and H.-A. Kao, A cyber-physical systems architecture for industry 4.0-based manufacturing systems, *Manufacturing Letters* **3**, 18–23 (2015).
57. Y.-Y. Hsieh, The rise of decentralized autonomous organizations: Coordination and growth within cryptocurrencies, *Ph.D. dissertation, The University of Western Ontario*, (2018).
58. M. Swan, Blockchain: Blueprint for a new economy, *O'Reilly Media* (2015).
59. M. Swan, Blockchain thinking: The brain as a Decentralized Autonomous Corporation [Commentary], *IEEE Technology and Society Magazine* **34**, 41–52 (2015).

60. V. Buterin, A next-generation smart contract and decentralized application platform, *White Paper*, 1–36 (2014).
61. I. Lianos, P. Hacker, S. Eich, G. Dimitropoulos, *Regulating Blockchain: Techno-Social and Legal Challenges*, *Oxford University Press*, (2019).
62. P. Veena, S. Panikkar, S. Nair, P. Brody, Empowering the edge: practical insights on a decentralized internet of things, *IBM Institute for Business Value*, (2015).

## Acknowledgment

**Funding:** This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 751615. Marco Dorigo acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Research Director. **Author contributions:** E.C.F. designed and realized the system as well as performed the experiments. E.C.F. and M.D. analyzed the results and wrote the manuscript. T.H., A.P., and M.D. directed the project. **Competing interests:** The authors declare that they have no competing interests.

## Supplementary materials

Section S1. Detailed description of an MT structure, proof, and data encapsulation

Section S2. Detailed description of the real-robot experimental setup

Section S3. Detailed description of the simulation experimental setup

Section S4. Description of the finite state machines controlling the robots

Section S5. IRIDIA Swarm Marketplace

Figure S1. Description of a typical MT structure and proof

Figure S2. Robot interaction and FSM diagrams

Figure S3. Experimental setup for the real-robot experiments

Figure S4. IRIDIA Swarm Marketplace

Movie M1. Typical simulation run for the foraging experiments

Movie M2. Typical real-robot run for the foraging experiments

Movie M3. Typical simulation run for the maze formation experiments

Movie M4. Typical real-robot run for the maze formation experiments

Movie M5. Tangible Swarm (Supporting features)

# Supplementary materials

## S1 - Detailed description of an MT structure, proof, and data encapsulation

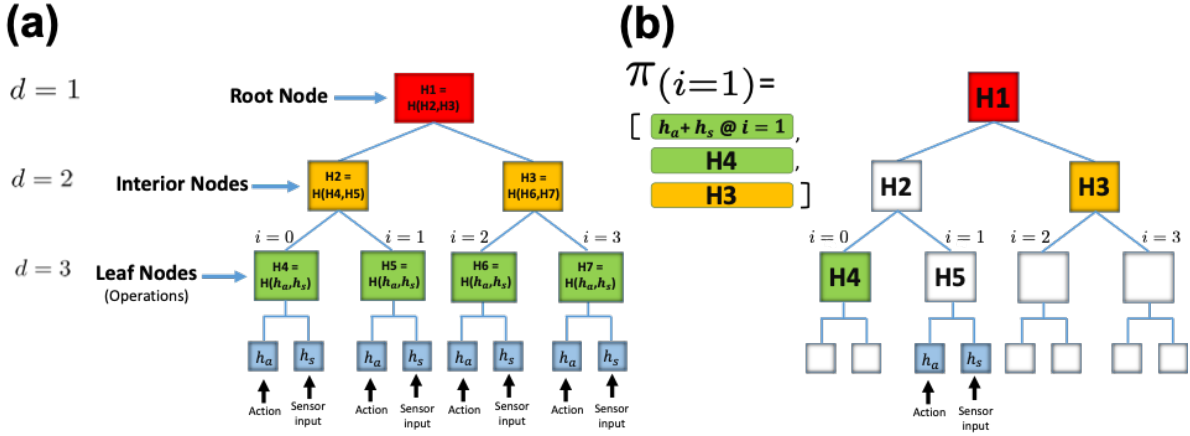


Figure S1: **Description of a typical Merkle tree (MT) structure and proof.** (a) Merkle tree implementation ( $d = 3$ ) with 4 leaves, 2 interior, and 1 root nodes. Each leaf node (green) encapsulates the hash of two hashes: a robot action ( $h_a$ ) and a robot sensor input ( $h_s$ ). Each leaf node represents one atomic operation that robots should complete sequentially ( $i = 0, i = 1, \dots, i = m - 1$ ) in order to fulfill the swarm’s mission. Interior nodes encapsulate the hash of their two children, and the root node encapsulates the hashes of its two interior nodes. (b) Describes the path (in color) in order to get the proof for leaf node  $i = 1$ .

A Merkle tree (MT) of depth  $d$  is a tree with  $m = 2^{d-1}$  leaf nodes, each one with an index  $i \in [0, m - 1]$ . Each leaf node encapsulates a hash string of an associated operation, while each interior node contains the hash of the combination of its two children. A depiction of a complete MT for  $d = 3$  is given in Fig. S1(a). Each operation (green blocks) encapsulates the combined hash of two hashes:  $h_s$  (hash of the sensor’s input  $s$ ) and  $h_a$  (hash of the robot action  $a$ ). These two hashes describe an atomic step within the swarm’s high-level mission. For instance, the hash of the action “push” ( $h_a$ ) and the hash of the sensor input “block (1)” ( $h_s$ ) would be included in one of the operations by using the hashing function  $H: H(h_a, h_s)$ . As outlined in Fig. 1(d), when a *verifier* ( $V$ ) robot queries a *prover* ( $P$ ) robot with updated

information (i.e.,  $P$ 's index  $i$  is greater than  $V$ 's index  $i$ ),  $P$  returns a chain of digests  $\pi$  needed to compute the root node digest (red block).  $V$  keeps (at least) a copy of the root node hash itself (computed before the experiment starts), and checks  $\pi$  by trying to recompute the root node hash in a bottom-up manner.

Fig. S1(b) shows the proof  $\pi$  for a fetch at operation  $i = 1$ . It consists of four elements in sequence, the two hashes  $h_a$  and  $h_s$ , the hash H4, and the hash H3. Then, the verification proceeds bottom up:  $V$  computes the hashes of the two hashes  $h_s$  and  $h_a$ , which is H5, and concatenates it with the hash of H4 provided in  $\pi$ . Next, it concatenates H4 and H5 and computes the hash of the digest for node H2. Then, it concatenates H3 provided in  $\pi$ , with the computed H2, and hashes the result. Finally, it checks whether this computed digest equals H1 (root node hash).

In this research, we are interested in two properties derived from the workflow between  $P$  and  $V$  robots depicted in Fig. 1(d): correctness and security. First, correctness implies that when  $P$  executes a query  $Q$  over its own MT, the generated proof  $\pi$  requires knowledge of  $h_a$  and  $h_s$  and can be easily verified by  $V$  locally. This property opens the path to secret cooperation between robots since encrypted verification and validation of data (e.g., by using cryptographic hashes) can be exchanged within the swarm without disclosing any raw or unprotected information. Second, security implies that a computationally limited, deceiving  $P'$  cannot induce  $V$  to accept a faulty answer. The basis of this property is the use of collision-resistant hashes: if  $P'$  can cause  $V$  to accept an incorrect answer then the proof returned by  $P'$  will yield a collision (i.e., two different inputs produce the same output hash value). This research assumes the hash generation method used (SHA256) is collision-free, that is, that the probability that different inputs produce the same output is negligible. This property leads to secure cooperation between robots.

## S2 - Detailed description of the real-robot experimental setup

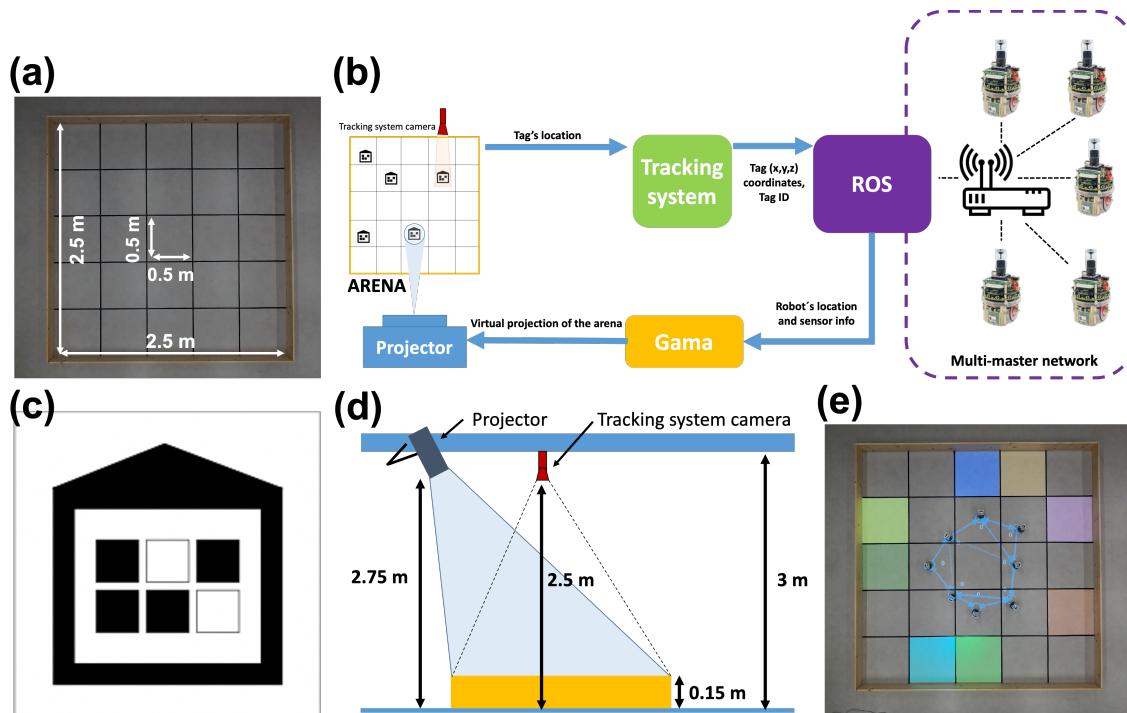


Figure S2: **Experimental setup for the real-robot experiments.** (a) arena (with measures) used during the real-robot experiments. (b) Information workflow for the real-robot experiments. The diagram shows how robot tags' location are captured by a vision-based tracking system, coordinates for each tag are calculated, and used by the robots connected in a network. Then, relevant information such as robot's sensor input or status can be projected on the arena where experiments are being run in real-time. (c) Tag with ID number 34 encoded in binary as a matrix of black and white squares. (d) Diagram with the physical disposition of the sensing (camera) and actuation (projector) elements introduced in (b). (e) Image of a real-robot experiment for the foraging mission using the information workflow described in (b).

Fig. S2(a) shows a bird's eye view of the arena built for the real-robot experiments. The arena is a rectangular area of  $2.5 \times 2.5 \text{ m}^2$  built by 4 wooden planks attached to each other using hot melt glue. Each one of these planks has the following measures: 2.5 m (length)  $\times$  0.15 m (height)  $\times$  0.05 m (depth). The experimental area is divided into a  $5 \times 5$  grid where each cell corresponds to a  $0.5 \times 0.5 \text{ m}^2$  area. The black lines between the cells were created by using black duct tape which also covers the rim of the linoleum tiles underneath.

Fig. S2(b) shows the information workflow that takes place during the real-robot experiments. First, each robot carries a tag (Fig. S2(c)) on top of its omnidirectional camera cylinder; this tag encodes a specific robot ID used in the system. Second, tags are recognized by a vision-based tracking system (45) using a camera (Prosilica GC 1600C) right on top of the arena (Fig. S2(d)). Third,  $(x,y,z)$  coordinates for each tag are calculated with respect to the bottom-left corner of the arena and sent to a ROS (46) multi-master (47) instance connected to the robots through a WiFi router. With this configuration, robots can access their own position and orientation and use them during the mission (e.g., locate the target location, find the center of each cell, etc.). The WiFi network is used as a low-level communication medium for the robots. In order to keep the decentralized aspect of the robot swarm, we designed a higher-level protocol where each robot uses its own ROS multi-master instance to communicate with other robots (e.g., send queries, receive proofs, etc.). This topology avoids that the system relies on a single ROS master for data exchange. Robots communicate in a peer-to-peer fashion with each other since their own multi-master instances can only exchange data if robots are within communication range ( $C_{range}$ ).

### **S2.1 - Supporting features of the system**

In order to better understand the robot swarm behavior during an experiment, we built a system capable of collecting relevant information—such as the robots’ internal states and the interaction patterns within the swarm—from the robots during the mission and of displaying it on the robot swarm working space in real-time. Following the workflow depicted in Fig. S2(b), information about the robots’ location, sensor input and internal state is extracted from the multi-master network at a frequency of 10 Hz and sent to an agent-based simulation software named Gama (48). With this information, Gama is able to create a virtual representation of the current state of the arena which is projected on the arena itself by using a short throw projector

(Optoma EH460ST). Fig. S2(d) shows a diagram of the arrangement of the tracking system camera and of the projector with respect to the arena. Fig. S2(e) shows a real-robot foraging experiment where the virtual representation of the arena (e.g., colored cells, communication patterns between the robots, robot’s internal variables, etc.) are projected on the arena and robots in real-time. Detailed information about the amount of information displayed in both foraging and maze-formation missions is included in the M2 and M4 movies.

The idea of building a close-loop control where a camera (e.g., performing tag tracking, pattern recognition, etc.) acts as the sensing part and a projector is the actuation component can be extended to other uses and scenarios. For that reason, we designed a tool named *Tangible Swarm*<sup>4</sup> that has the aim to display in a real-time fashion relevant information about a swarm robotics system at the same place where the swarm is conducting its mission. Information such as IDs, sensor inputs, robot trajectories, distance between robots in a formation scenario, battery status, communication patterns, etc. can be easily coded, displayed and customized for different robot swarm missions. Movie M5 shows several of the aforementioned capabilities and gives an overview of the tool capabilities.

### **S3 - Detailed description of the simulation experimental setup**

For the results obtained in the simulated foraging mission, the following parameters were used: the color  $k$  of a task is chosen in the set {green, magenta, blue, yellow, red, cyan, lime, orange}. Only one task of each color  $k$  is placed in the arena at the beginning of each run. For example, in the run where the MT has 4 operations ( $n = 4$ ), a total of 4 tasks were randomly scattered in the arena. For the robot communication and interaction space (as depicted in Fig. S3) we used the following parameters: The robot communication range ( $C_{range}$ ) was set to 1 m (maximum range of the range and bearing device of the e-puck robots), the vision sensing distance ( $V_{range}$ )

---

<sup>4</sup><https://github.com/edcafenet/TangibleSwarm>






Figure S3: **Robot interaction diagram.** Every robot controller relies on three main thresholds:  $O_{range}$  represents the range where obstacles are detected by robots in order to avoid collisions.  $V_{range}$  represents the robot's maximum vision range to detect tasks (i.e., see LED markers) in the arena.  $C_{range}$  represents the maximum required distance for robots to communicate with one another. Finally, robots are able to locate themselves in the arena's 2D space either to find the target location (foraging mission) or to obtain the (x,y) coordinates of the grid (maze formation mission).

was set to 0.35 m (maximum range to detect an LED marker from within the same cell). LEDs are used in the foraging mission to help robots visually locate the cell and define its color. The robot obstacle detection range ( $O_{range}$ ) was set to 0.10 m. In addition, robots are able to locate themselves (Fig. S3) in the 2D arena by using the simulator positioning sensor, this information is used to travel to the center of the cell of the target location. Finally, the Time Cap (TC) for each experiment was set to 5,100 seconds (maximum battery duration<sup>5</sup>). All the parameters mentioned previously were used also for the maze formation mission, excluding those concerning the discovery of tasks. In contrast to the foraging mission, in the maze formation mission,

<sup>5</sup>According to manufacturer: [https://www.gctronic.com/doc/index.php/Overo\\_Extension#Consumption](https://www.gctronic.com/doc/index.php/Overo_Extension#Consumption)

robots are able not only to locate themselves in the 2D arena ( $2.5 \times 2.5 \text{ m}^2$ ), but also to calculate the (x,y) coordinates of the grid which they used as input to generate the  $h_s$  hash. All simulation experiments presented in this paper were conducted in ARGoS (38), a modular multi-robot simulation and development environment.

## S4 - Description of the finite state machines controlling the robots

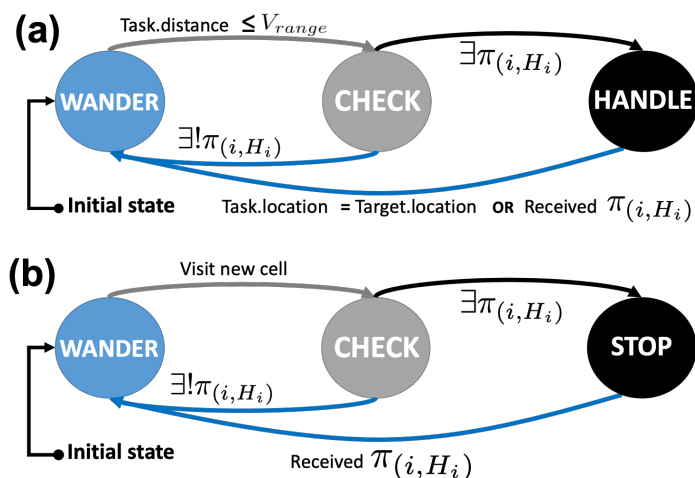


Figure S4: **Robot FSM diagrams.** (a) Finite State Machine (FSM) controller for the foraging mission. Three behaviors —**Wander**, **Check**, and **Handle**—are coded in every robot. (b) FSM controller for the maze formation mission. This FSM differs from the one in (a) in two ways. First, the **Handle** behavior is substituted by the **Stop** behavior, which makes robots stop at the center of a grid cell instead of carrying a token to the target. Second, the return condition to the **Wander** behavior is only executed if robots receive proof that the cell they are trying to occupy has been already filled.

Fig. S4(a) depicts the finite state machine (FSM) that controls the robot for the foraging mission, which relies on three basic behaviors:

**Wander.** The robot performs a random walk searching for tasks in the arena. If the robot detects a task within its vision range ( $\text{Task.distance} \leq V_{range}$ ), it executes the **Check** behavior; otherwise it continues searching. During the execution of this action the robot is able to detect obstacles such as walls or other robots (within distance  $O_{range}$ ) and avoid them.

**Check.** Once a task is within the robot’s vision range, the robot can sense the color of the LED marker associated with it. This information is used by the robots in order to generate the sensor input hash ( $h_s$ ). For the considered foraging mission, the action hash  $h_a$  encodes the “carry to target” action. Robots executing the **Check** behavior combine  $h_s$  and  $h_a$  to generate a hash ( $H_i = H((h_s, h_a))$ ) that is used to generate the proof  $\pi$ . In case a proof  $\pi$  exists for the current working operation  $i$ :  $\exists \pi_{(i, H_i)}$ , the combination of task color and action (i.e., operation) can be verified as part of the MT. Otherwise ( $\exists! \pi_{(i, H_i)}$ ), the robot returns to the **Wander** behavior.

**Handle.** In case the robot generates a valid proof for the visible task, the robot travels to the center of the cell. Then, the robot delivers the color to the target located at center of the arena. Once the robot reaches its destination (`Task.location = Target.location`), the robot releases the task information and changes the status of  $i$  as completed in its local MT. Then, the robot increments the pointer of the current working operation index:  $i = i + 1$  for  $i \in [0, n-1]$ . In case the robot receives a proof that the  $i$  has been already completed by another robot while carrying the task, the robot updates its own MT with the received information and returns to the **Wander** behavior.

Fig. S4(b) depicts the FSM that controls the robot for the maze formation mission. There are three differences with respect to the foraging mission. First, robots use the (x,y) coordinates of visited cells to generate  $h_s$ . Second, the action hash  $h_a$  encodes the “stop” action. Third, in case a valid proof  $\pi$  is generated with  $h_s$  and  $h_a$ , robots execute the **Stop** behavior, which makes them find the center of the visited cell and stop there. After the **Stop** behavior is executed, robots cannot run other commands and they remain moveless until the end of the run.

## **S5 - IRIDIA Swarm Marketplace**

Academic and industry leaders agree that we are currently on the verge of the fourth industrial revolution (49–51), one that may be marked by the omnipresence of networks instead of stand-alone systems. More and more machines are interconnected in these networks, objects such as sensors, smart devices, factories, and so on, are gradually creating a global industrial Internet of Things (IoT) (52). It is expected that this fourth industrial revolution will create new ways for people, machines, and organisations to be interconnected and communicate with each other by exchanging and sharing information. At the core of this new set of technologies there is a fundamental shift from centralized and human-mediated systems to decentralized and autonomous systems (53). Autonomous robots (54, 55) and cyber-physical systems (56) belong to a growing category of devices that can be programmed to perform tasks with little or no human intervention. They can vary significantly in size, functionality, mobility, dexterity, intelligence, and cost, to provide from robotic process automation to data capturing capabilities.

In the meantime, in the blockchain space, new concepts such as the Decentralized Autonomous Organization (DAO) (57–59) are starting to gain popularity. A DAO is an organization that is run through rules encoded as smart contracts (i.e., computer code embedded in a blockchain that directly controls the transfer of digital assets between parties under certain conditions) (60). In these organizations, machines and people can cooperate without the need to be incorporated into traditional business identities (61). Such autonomous organizations can charge users for the services they provide, in order to pay others for the resources they need. Essentially, DAOs are entities that often need minimal or no input at all to be able to operate, and that are used to execute smart contracts and record activity on the blockchain.

In recent years, the DAO field has also explored the combination with robotics creating a type of decentralized entity where part of the control logic resides in a smart contract while the actuation part resides in the physical world (62). However, there is still no example in the liter-

ature of how an institution can start offering autonomous systems services and match the needs of potential customers that might require those services on demand. With the aim to explore this path, we present the IRIDIA<sup>6</sup> Swarm Marketplace: a web-based service marketplace whose logic is coded in a smart contract and uploaded in the Ethereum blockchain. In the proposed approach, staff from IRIDIA advertises the robotics services available (number of robots, duration of service, and price). Then, customers are able to purchase these services and pay the price with their own crypto-wallets. Then they can upload an MT with the list of operations the robots need to complete. Once the service is completed, customers get the cryptographic proof that the robots completed all the operations included in the MT, which allows them to trust the system and understand the service was not faked. Finally, customers get pictures and video footage of the final work the robots conducted. Videos are uploaded by the system to the public IRIDIA Youtube playlist together with the root of the provided MT.

Fig. S5 shows the general framework and main components of the IRIDIA Swarm Marketplace: a web-based marketplace for swarm robotics services built on blockchain technology. Fig. S5(a) depicts the main workflow of the proposed application. Let's consider an example of how this market place could be used. In the planning of research activities there are periods of time in which IRIDIA's robot swarm remains idle. When this happens, the IRIDIA staff members taking care of the robot swarm allocation to the different research activities will broadcast the robot swarm availability through a market-based website as the one depicted in Fig. S5(b). In this website, IRIDIA staff can log into the platform by using lab-controlled Ethereum accounts and the MetaMask interface (a popular Ethereum wallet that allows your browser to connect to the Ethereum blockchain). Then, a new service can be added where the number of robots available, the amount of time they are available for, and the price charged for their use are indicated. Once the service is broadcast, potential customers interested in using the robot

---

<sup>6</sup>IRIDIA is the name of the artificial intelligence lab of the Université Libre de Bruxelles, where the real-robot experiments presented in this paper took place.




Figure S5: **IRIDIA Swarm Marketplace.** (a) Information flow for the proposed application where: (1) IRIDIA staff adds services to the smart-contract by using the web interface depicted in (b). (2) Customers can purchase available services (i.e., a certain number of robots, for a time, for a price) by providing a valid MT as input and paying the price listed in cryptocurrency (i.e., Ethers). (3) Then, the smart-contract sends the MT to the robot swarm and the mission starts. (4) When the job is finished, the resultant hashes that prove that the robots did the job requested are sent to the smart-contract and displayed in the results section of the web interface. (5) The smart-contract displays the results video of the job done by the robots in a public site where the customer can retrieve it. (b) Web interface where services are added by IRIDIA staff and purchased by customers.

swarm can purchase the service and upload their own MT in the system. After the received MT is validated, the smart-contract implementing the Swarm Marketplace sends the information to the robots so that they can start working. Robots complete the set of actions the same way it was described in the real-robot experiments of this work. In our initial demonstration example, we provide customers robots that have the maze-formation set of actions so that, when they upload their MTs, they can choose which kind of pattern to create. Once the mission described in the uploaded MT is completed by the robots, the complete set of proofs with their correspondent  $h_a, h_s$  values is sent to the smart-contract and stored there. Finally, customers can make sure their job has been completed by the robots since the values that were used as inputs for the MT have been discovered by the robots. Additionally, a video of the robots performing the mission is uploaded to the IRIDIA Youtube playlist and available for the customer. Fig. S5(b) shows the web interface where the buyer and seller broadcast and purchase services. A live demo of the IRIDIA Marketplace is accessible at [www.blockchainswarm.eu](http://www.blockchainswarm.eu).