ADVANCED INTELLIGENT SYSTEMS
Open Access
www.advintellsyst.com

Distributed Situational Awareness in Robot Swarms

*Simon Jones, Emma Milner, Mahesh Sooriyabandara, and Sabine Hauert**

Many-robot systems are becoming a reality for large companies that can invest in bespoke solutions. These systems often require carefully engineered infrastructure and a central planner to coordinate the robots. Outside these controlled environments, robots typically generate a shared situational awareness of the world and state of their task. This requires sophisticated mapping, perception, and control, with changes to the environment or tasks causing challenges to robot deployment. The assumption that centralized situational awareness is needed to deal with real-world complexity may be holding back the field from deploying many-robot systems. Yet potential applications are wide-ranging, including environmental monitoring, construction, agriculture, and logistics. Mainstream adoption requires usability out-of-the-box, in unstructured environments, at a reasonable cost. Distributed situational awareness is proposed as a method to design many-robot systems differently. Distributed situational awareness allows swarms of low-cost robots to rapidly and accurately capture the state of an environment and act accordingly, with no central data storage, modeling, or control. Its distributed nature enhances resilience and redundancy while reducing reliance on infrastructure and central planners. Deploying distributed situational awareness however requires new tools to design hardware and algorithms, demonstrate that it works and is safe, and is intuitive for users of the swarm.

1. Distributed Situational Awareness

Situational awareness is a model for the process by which an operator accepts sensory inputs and knowledge and uses these to synthesize an integrated model of the environment within which they must make decisions and act. The field has expanded to many areas where humans must make decisions to control complex and dynamic systems. The controlling human must not just understand the readings of individual sensors but infer a broader systemic meaning from them within a goal framework to make valid decisions. Endsley divides this process into "perception" of the environment, "comprehension" of the situation in relation to goals, and "projection" into the future.[1]

This article proposes a different concept, distributed situational awareness, which allows a swarm of robots to rapidly and accurately capture the state of an environment and act accordingly, without the need for any heavy infrastructure, central data storage and processing, or control. This is done by having every robot generate its own local situational awareness, which drives its immediate actions. The focus on local information allows robots to rely on limited-range sensing and communication capabilities such as cameras, distance sensors, or Bluetooth, which are widely available at a low cost. This richness of local information combines implicitly to form an emergent overall situational awareness, which drives the behavior of the swarm. Swarms with distributed situational awareness have the potential to be useable out of the box, in a scalable manner, across many applications, including environmental monitoring, construction, agriculture, and logistics.[2]

Distributed situational awareness builds on important pieces of work in swarm robotics, especially on the use of local perception and action to drive the emergent functionality of the swarm.[3,4] Designing these local rules toward desired swarm behaviors is the central challenge of swarm engineering, with solutions found in bioinspiration or automatic discovery using artificial evolution and machine learning.[5] Focus now is on the transition from swarms in the lab to real-world applications.[2] Yet the use of swarm terminology has resulted in unnecessary barriers to their mainstream adoption, mostly due to the perception that individual robots in swarms are too simple or minimal to be used outside the laboratory and that real-world applications require easy access to centralized information about the state of the system to be useful, reliable, and easy to

Dr. S. Jones, E. Milner, Prof. S. Hauert
Bristol Robotics Laboratory
Bristol BS34 8QZ, UK
E-mail: sabine.hauert@bristol.ac.UK

Dr. S. Jones, Prof. S. Hauert
Department of Engineering Mathematics
University of Bristol
Bristol BS81UB, UK

E. Milner
Department of Aerospace Engineering
University of Bristol
Bristol BS81UB, UK

Prof. M. Sooriyabandara
Toshiba Research Europe Limited
Bristol BS1 4ND, UK

The ORCID identification number(s) for the author(s) of this article can be found under https://doi.org/10.1002/aisy.202000110.

DOI: 10.1002/aisy.202000110

interact with, as is the case with classically engineered systems. Distributed situational awareness alleviates some of these concerns by focusing on the information that can be gathered at a local level by individual robots, including more sophisticated robots, how this information can be used to drive human-understandable tasks by the swarm, and how the system can be queried. The human is seen as an embedded agent in this system, also capable of distributed situational awareness.

2. Use Case

A small company that organizes bespoke events ranging from 50 to 10 000 attendees is looking for a solution to automate their cloakroom. Venues change location and size. Typically, deploying such a multirobot system would either require bespoke infrastructure (e.g., rails or lines on the ground) or robots that can assess their location, a map of the environment, and can navigate the room using this information. When a belonging is received, it is processed by a central control unit that plans where the belonging goes, a robot is then assigned to the task and deposits the belonging at the correct location. When the user returns to retrieve the belonging, the central control unit remembers where it was stored and directs a robot to the location to retrieve it. This approach requires a globally networked communication system linked to a central control station, sophisticated sensing on board the robot to understand the environment for navigation, infrastructure to allow robots to localize, and this may fail if any of the sensing is off or the belongings have been moved, for example, to free a passage.

Instead, using distributed situational awareness, their robots operate without any central control, global knowledge of their environment, or heavy infrastructure (**Figure 1**). **Table 1** compares this strategy to centralized situational awareness. Local situational awareness is informed by cheap and widely available hardware including a camera, infrared (IR) distance sensors, RFID sensor, and Bluetooth, which determine their interaction range. This distributed situational awareness, combined with the right actions, enables them to complete the desired collective behavior of storing and retrieving items. Installing the robots is as easy as unpacking them, delimiting an operation area either physically (e.g., a walled room) or by IR barriers and defining a deposit and retrieval area for the belongings by posting a marker on the wall.

To use the service, users download the cloakroom app. Belongings are deposited in a small box tagged with an RFID and left in a marked deposit area. Robots waiting nearby pick up the box and move through the storage area by performing obstacle avoidance. After a random amount of time sampled within a predefined range (e.g., 10–360 s), and based on space availability, they deposit the box. Robots are distributed through the storage area and so can quickly identify and navigate to the belongings based on the RFID tag of the box when a request for the belongings is broadcast through the swarm by a user with an app on their phone. The closest robot arrives at the belongings first, picks it up, and returns to the retrieval area using obstacle avoidance and attraction to a large visual marker taped above the deposit area or by following a multihop communication trail back to the user. Manually moving boxes has no impact on the retrieval of the belongings, and as a result humans and robots can work together in a common environment. The robots do not require positioning, dedicated infrastructure, external WiFi, or a central control unit. Instead they have an overall awareness of the storage area, distributed over the swarm. In this example, distributed situational awareness works much closer to the way humans would operate a cloakroom service, as they would typically tag a belonging, deposit it in the storage area, and then would search for the belonging based on its tag when requested by the user. Adding more humans would lead to more organizational capacity, and the front desk operator is not expected to have an overall view of exactly where each belonging is.

3. Scenarios That Benefit from Distributed Situational Awareness

Centralized systems are often the answer when investment can be made upfront in the design and maintenance of a custom solution that is tailored to a specific environment. In these situations, centralization has the advantage of resulting in fast and effective solutions, albeit at the cost of upfront investment, a slower installation, and potential bottlenecks to scale up such as limited communication bandwidth. Instead decentralized situational awareness makes most sense in scenarios that require fast, out-of-the-box setups without heavy infrastructure, robustness

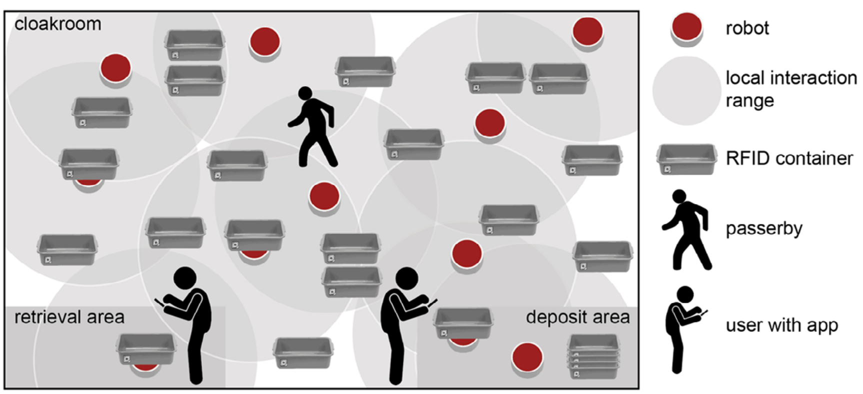

Figure 1. Pop-up cloakroom powered by a swarm of robots using distributed situational awareness.

Table 1. Strategies for the deployment of multirobot systems in a cloakroom scenario using centralized situational awareness and distributed situational awareness.

Task	Central situational awareness	Distributed situational awareness
Install cloakroom	WiFi is configured on robots. Map of the environment is generated and uploaded to the central controller. Robot positioning system is installed and tested.	Robots are taken out of their box. The area is delimited (physically or virtually). A sign is posted on a wall near the deposit and retrieval area.
Deposit belongings	The central controller receives the request and delegates one robot to pick up the belongings and designates a location for it to be stored. The central controller computes the path for the robot to follow.	Any robot present in a designated "deposit area" can pick up belongings deposited in a box tagged with an RFID, move it around the storage area, and deposit it after a random amount of time and based on space availability. Robots perform random walks, or other minimal navigation strategies, based on collision avoidance to move through their environment.
Query the environment	The state of the storage area is stored in a central database that can be queried by external users. This provides them with a map of each belonging's location.	Robots can rapidly gather an inventory of RFID tags in their local environment and transmit this information to the user's app by forwarding it through the swarm. Alternatively, robots can change color to signal a specific belonging in their neighborhood.
Retrieve belongings	The central controller receives a request to retrieve belongings. The belonging is located on the central map and a designated robot is guided to its last known location and then back to the retrieval area.	A request to retrieve a belonging is transmitted to the swarm by the user through their app. Robots then sample their local environment for the corresponding RFID tag and pick it up when found. They then proceed to retrieval area. Once the belonging is received, the user deactivates their request.

(i.e., failure of robots doesn't fully impede the system), and adaptability. Adaptability here relates to the ability of the swarm to be effective across different deployment scenarios and within scenarios that are dynamic. Performance and cost are linked to the number of robots deployed, allowing users to purchase just a few robots to start and then scale up when appropriate. The lack of centralized information however means the system may be more difficult to optimize for performance and is more challenging to design. In many situations, the solution may be a hybrid between central and distributed situational awareness. For example, in the cloakroom scenario, a central station may still be installed that collects regular snapshots from the swarm regarding storage organization, can serve as an interface to a user who hasn't installed the app on their phone, and could be used to monitor individual robot health or the health of the swarm.

4. Reimagining How Things Are Done

Distributed situational awareness allows us to entirely redesign space, as only local structure matters in the performance of the system. The resulting emergent storage organization may look atypical for an outside observer. In the cloakroom scenario, for example, this may lead to storage space that looks much more organic than would typically be human-organized. By following simple local rules, belongings may end up organized in clusters, spirals, or uniformly distributed over the environment, all patterns that are atypical for most cloakrooms but are reminiscent of patterns found in natural swarm structures including termite mounds and ant nests. Resulting patterns may be better suited for local situational awareness and action, while still achieving the overall aims of the swarm, here storage and retrieval. In the same vein, distributed situational awareness could allow for crops to be mixed within a single farm field, with local awareness and action leading to the cultivation of individual plants rather than a full field. So-called edible forests could grow more like a natural, resilient ecosystem. Construction swarms would produce buildings that are built following rules, not blue prints, leading to entirely new designs that are emergent. Swarms used for environmental monitoring, rather than relying on a stream of information being aggregated on a central controller, could display hazard levels by making them visible locally through illumination.

5. Misconceptions

5.1. Real-World Applications Require Centralized Information

One of the key challenges in affecting this paradigm shift from centralized to distributed awareness is the preconception that to operate, multirobot systems have to have a centralized high-quality complete model of the world on which to make decisions. As more sophisticated single-robot systems emerge, it's natural to believe that multirobot systems equally have to be increasingly sophisticated as a collective. Instead, sensing and modelling of the world can be distributed, with actions mostly reactive to these local inputs. If this paradigm shift sounds familiar, it's because it is similar to the shift to behavior-based robotics, with emphasis given on the ability of robots to react to their local environment following a simple set of rules rather than to a sequence of modelling/planning/acting steps.[6] The simplicity of behavior-based robotics powered one of the first successful commercial domestic robots with the Roomba. Many-robot systems are now the ideal platforms to adopt this philosophy because of the otherwise complexity of designing for centralized awareness and because of the additional pressure to drive down cost when building many entities. The assumption that centralized situational awareness is needed to deal with the complexity of the real world may be holding back the field from deploying many-robot systems.

5.2. Individual Robots Have to Be Minimal

It is often thought that individual robots in many-robot systems have to be relatively simple or even minimal.[7] While this is often the case, increased capabilities on the individuals, driven by the advent of low-cost hardware, could also be used to great effect at a local level, allowing sophisticated modelling of the local environment and interpretation of sensory data. The resulting distributed situational awareness could result in richer information at the level of the swarm than with central situational awareness, with actions happening at the right location and time, leading to desired collective behaviors.[8] Distributed situational awareness fundamentally is more than the sum of its parts. Pooling all the information from the robots to a central control and then using this central control to direct the actions of the individual robots would invariably lead to a loss of information. Instead, robots can assess all the information locally, leading to reduction in communication and processing requirements.

5.3. Users Can't Work with Swarms

Because of their decentralized nature, swarms are seen as too difficult to interact with as there is no central control to receive and send commands. This does not have to be the case. By being decentralized, users can arguably interact with robots interchangeably, query the system for up-to-date snapshots of their state, and interact closely with the swarm by being themselves part of the collective system and its complex environment. Likewise, their decentralized nature may make them easier to install out-of-the-box by nonexperts, as installation can be iterative (one robot installed after the other) while still being operational and does not require heavy infrastructure and computational skills.[9]

6. Overcoming Barriers

The challenges to the deployment of distributed situational awareness include designing these systems (hardware and algorithms), ensuring that they are safe and effective, and deciding how to interface with them as users.

6.1. Designing for Distributed Situational Awareness

Systems that make use of distributed situational awareness are not obvious to design as they result from many robots operating using limited local information without any central control. Designing hardware and software for such systems can be challenging.[10] Hardware has to allow for local sensing and interactions with other robots and their environments, the ability to convey state to the outside world, and sufficient capability to conduct their task. Yet few swarm platforms have moved out of the laboratory and beyond toy scenarios. The recent review by Schranz et al.[2] provides an excellent overview of platforms that are now available and their potential applications. Algorithmic tools are also allowing for the automatic discovery of swarm controllers that give rise to desired collective behaviors in a variety of scenarios.[5,9] Artificial evolution of behavior trees for example allows for automatic design of swarm controllers that are human readable.[11]

6.2. Demonstrating Distributed Situational Awareness Works and Is Safe

Our recent work[12] showed that potential users of swarm technology are broadly positive about its use in their professions but would have to be shown that it works reliably, can be trusted, and that information will not be lost, or made impenetrable. One example was a museum that worried that the robots might lose the information about the items stored, which would waste time. With distributed situational awareness, the "database" however can constantly be recreated and updated through robots making local inventories of objects and pooling them to the user. Because of their decentralized nature and resulting emergent collective behaviors, demonstrating that systems that use distributed situational awareness will work reliably is challenging and may rely on a combination of formal verification where possible and thorough testing. Safety considerations are needed to determine if distributed situational awareness is safe for the users, safe for the environment, and ethical as outlined in our recent work.[13]

6.3. Making Distributed Situational Awareness Intuitive for Users

Understanding how users and distributed systems interact will be key to their successful deployment. The challenge is both inputting user requests in the system and conveying the state of the system to the user. This has led to a new field of human–swarm interactions, focused on studying how humans will interact with swarms and how to marry needs of oversight with the benefits of distributed situational awareness.[14] Ideally, when the user is external to the swarm, the interface with the swarm would be similar to a centralized system, although the interworkings are distributed. The user can also be an active agent with distributed situational awareness in the swarm, in which case the local interactions between robots and individual robots should be intuitive. For the cloakroom scenario, for example, the user might query the system using the app to check if a luggage is there (external agent) or may enter the cloakroom to deposit and retrieve a suitcase (embedded agent).

7. Conclusion

Swarms are now ready to move to real-world applications, yet they suffer from the misconception that their decentralized nature makes them too simple to be useful, too unpredictable, and difficult to interact with. Central to these assumptions is the idea that to operate in the real-world, multirobot solutions need to have centralized situational awareness—a synthesized integrated model of the environment and the robots, which drives individual robot actions. Distributed situational awareness instead builds on individual robots' local awareness to effectively work as a collective toward global tasks that would typically require global information and central controllers. Tools to discover suitable robot local awareness and actions, giving rise to desired collective behaviors, are at the core of swarm engineering. The distributed nature of these swarms makes them useable out of the box, with little infrastructure, in a scalable and adaptable way. This makes them accessible to a wider audience, typically locked out from their use due to high costs, long setup times, and need for infrastructure. Designing these systems however requires new hardware and algorithms that are tailored to the applications at hand, new solutions to demonstrate that they are safe and reliable, and studies on how best to interact with swarms as part of the distributed system.

Acknowledgements

S.J. and S.H. were supported by EPSRC IAA Award EP/R511663/1. E.M. was supported by the EPSRC Centre for Doctoral Training in Future Autonomous and Robotic Systems (FARSCOPE) and an EPSRC ICASE Award.

Conflict of Interest

The authors declare no conflict of interest.

Keywords

distributed situational awareness, logistics, swarm robotics

Received: May 25, 2020
Revised: July 2, 2020
Published online:

[1] M. R. Endsley, *Hum. Factors* **1995**, *37*, 32.
[2] M. Schranz, M. Umlauft, M. Sende, W. Elmenreich, *Front. Robot. AI.* **2020**, *7*, 36.
[3] M. Brambilla, E. Ferrante, M. Birattari, M. Dorigo, *Swarm Intel.* **2013**, *7*, 1.
[4] H. Hamann, *Swarm Robotics*, Springer International Publishing, Cham, Switzerland **2018**.
[5] M. Birattari, A. Ligot, D. Bozhinoski, M. Brambilla, G. Francesca, L. Garattoni, D. Garzón Ramos, K. Hasselmann, M. Kegeleirs, J. Kuckling, F. Pagnozzi, *Front. Robot. AI.* **2019**, *6*, 59.
[6] R. A. Brooks, *IEEE. J. Robot. Autom.* **1986**, *2*, 14.
[7] E. Şahin, *International Workshop On Swarm Robotics*, Springer, Santa Monica, USA **2004**, pp. 10–20
[8] S. Jones, A. F. Winfield, S. Hauert, M. Studley, *Adv. Intel. Sys.* **2019**, *1*, 1900031.
[9] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, M. Lewis, *IEEE Trans. Human-Machine Syst.* **2016**, *46*, 9.
[10] G. Francesca, M. Birattari, *Front. Robot. AI.* **2016**, *3*, 29.
[11] S. Jones, M. Studley, S. Hauert, A. Winfield, *Proc. Distributed Autonomous Robotic Systems*, Springer, London, UK **2018**, pp. 487–501.
[12] D. Carrillo-Zapata, E. Milner, J. Hird, G. Tzoumas, P. J. Vardanega, M. Sooriyabandara, M. Giuliani, A. F. T. Winfield, S. Hauert, *Front. Robot. AI.* **2020**, *7*, 53.
[13] E. R. Hunt, S. Hauert, *Nat. Mach. Int.*, **2020**, *2*, 420.
[14] G. Kapellmann-Zafra, N. Salomons, A. Kolling, R. Groß, *Int. Conf. on Swarm Intelligence*, Springer, Cham, Switzerland **2016**. 125–136.

Sabine Hauert is associate professor in robotics at the University of Bristol and Bristol Robotics Laboratory. Before joining Bristol, she was a Human Frontier Science Program Cross-Disciplinary fellow at the Koch Institute for Integrative Cancer Research at MIT and Ph.D. student at EPFL. Her expertise is in engineering swarms across scales, from nanoparticles for cancer treatment, to robots for environmental monitoring. Central to her work is the use of bioinspiration and machine learning to engineer desired swarm behaviors. She is also president and cofounder of Robohub, a nonprofit organization dedicated to connecting the robotics community to the world.

