Multi-robot Coverage Using Self-organized Networks for Central Coordination

Aryo Jamshidpey(✉), Weixu Zhu, Mostafa Wahby, Michael Allwright, Mary Katherine Heinrich, and Marco Dorigo

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{aryo.jamshidpey,weixu.zhu,mostafa.wahby,michael.allwright, mary.katherine.heinrich,mdorigo}@ulb.ac.be

Abstract. We propose an approach to multi-robot coverage that combines aspects of centralized and decentralized control, based on the existing 'mergeable nervous systems' concept. In our approach, robots self-organize a dynamic ad-hoc communication network for distributed asymmetric control, enabling a degree of central coordination. In the coverage task, simulated ground robots coordinate with UAVs to explore an arena as uniformly as possible. Compared to strictly centralized and decentralized approaches, we test our approach in terms of coverage percentage, coverage uniformity, scalability, and fault tolerance.

1 Introduction

Multi-robot coverage control targets the systematic, uniform observation of a physical area or terrain. A widely studied approach is coverage path planning, in which the motion of robots is often centrally planned and coordinated, sometimes with prior knowledge of the size and shape of the environment [4]. Centralized approaches to coverage path planning have high performance, but are limited in terms of scalability and fault tolerance, due to a lack of redundancy that results in single points of failure and communication bottlenecks. Self-organized approaches to coverage, by contrast, are typically scalable and fault-tolerant, but are slow and inefficient compared to centralized approaches (e.g., [10]).

We propose a novel approach to multi-robot coverage control that seeks to combine aspects of centralized and decentralized approaches. Our approach is based on the existing concept of 'mergeable nervous systems' (MNS) [15], where robots assemble and physically connect, and temporarily yield control of their sensors and actuators to a single brain robot. In our prior work [25], we have extended the MNS concept to strictly wireless communication, rather than making use of physical connections. In our approach, robots establish asymmetric control over a dynamic ad-hoc communication network that is established and managed exclusively through self-organization. In this way, a self-organized network is used to implement some degree of central coordination, combining aspects of centralized and decentralized control. In this paper, we apply our

M. Dorigo et al. (Eds.): ANTS 2020, LNCS 12421, pp. 216–228, 2020.
https://doi.org/10.1007/978-3-030-60376-2_17

MNS approach to the task of multi-robot coverage. We also define comparable centralized and decentralized approaches to the considered coverage task and compare their performance with our hybrid solution. Increased decentralization in multi-robot systems typically involves increased parallelization and redundancy, such that a group of robots governed by centralized control is likely to be faster and more efficient than those governed by decentralized control. Therefore, we would expect a centralized approach to outperform the MNS approach, which in turn should outperform a decentralized approach. We test this using experiments in simulation. While decentralization might cause a decrease in efficiency and speed, it also provides desirable features such as increased scalability and fault tolerance. We therefore test the MNS approach to assess how well these features typical of decentralization have been preserved in our approach. Specifically, we assess scalability in terms of robot–robot communication and interference, and fault tolerance in terms of performance after robot failures.

1.1 Related Work

Coverage control has been widely studied in sensor networks (e.g., [6,23]), and has also been studied for search and exploration tasks with single robots and multi-robot systems. In the task of single robot coverage, the robot should gather information about the environment as efficiently as possible [7,11]. The overall time for a coverage task can be decreased by using multiple robots, but multi-robot approaches require solutions to efficient coordination. In centralized approaches, multi-robot coverage control is often approached as a path planning problem [4,5,24] making use of optimization or learning techniques. These approaches sometimes incorporate aspects of decentralized control. For instance, in [20], decentralized path planning relies on reinforcement and imitation learning through a centralized planner. In [14], robots use decentralized control to initially spread out in the environment, and then use a centralized approach for online learning of a density function.

In decentralized approaches, solutions to spatial coordination include leaving markings during exploration (e.g., artificial pheromones [12]), or maintaining communication (e.g., via line-of-sight [18]). Connectivity maintenance has been investigated in [13], seeking to maximize coverage and minimize communication overhead, and also has been investigated in [22], using a Voronoi tessellation approach to add fault tolerance. Connectivity maintenance during task parallelization has also been studied—using a distributed navigation controller and a global layer for task scheduling, in [16] it is shown that an hybrid centralized/decentralized approach can maintain connectivity in a scenario in which robots are deployed towards certain task-specific locations in the environment. Efficiency is also a key challenge for decentralized approaches, as they are prone to redundancy. In [10], a large number of robots perform coverage by simple collision avoidance, but full coverage is not guaranteed and efficiency is low, as robots frequently revisit areas. In a similar approach that reduces repeated coverage [8], robots leave markings during exploration; in another, a pheromone-based approach is used to achieve coverage efficiency [21]. Similar to pheromone-based approaches, activated beacons are used in [1] to guide coverage in a swarm of UAVs. Finally, coverage has also been studied in a heterogeneous swarm of robots with different sensing capabilities [19].

2 Methods

We investigate the applicability of the 'mergeable nervous systems' (MNS) [15] concept to the task of multi-robot coverage control. As the MNS concept combines aspects of centralized and decentralized control, we compare the performance of our approach to that of fully centralized control (i.e., all robots are controlled using global communication by a single robot with a global view) or fully decentralized control (i.e., all robots are controlled independently). In the decentralized approach, robots explore the environment by means of a random walk without any centralized coordination. In the MNS and centralized approaches, robots maintain a target formation while exploring the environment in a coordinated way. In the centralized approach, all robots are given motion instructions by one robot, whose identity as the central coordinating entity (i.e., the *master*) is predetermined and static. In our MNS approach, robots form a self-organized communication network—specifically, a directed rooted tree, where each link connects a parent robot to a child robot. One robot in the MNS is dynamically assigned the role of the *brain*, through self-organization (for the details of this process, see our prior work [25]). The robots use the network to receive motion instructions from their respective parents in the communication topology—except for the brain robot, which defines its own motion.

In this section, we describe the methods for our experiments. First, we define the coverage task. Second, we define the two motion behaviors that the three approaches can utilize during the coverage task. One is collision avoidance that is performed by robots independently, and is used in all three approaches. The other is perimeter following, which directs the motion of one robot, in both the approaches where robots are coordinated. The perimeter-following behavior is used by the brain robot in the MNS approach, and by the master robot in the centralized approach. Third, we define the target formations that robots maintain in the MNS and centralized approaches. Fourth, we give the implementation details of the three approaches. Overall, we keep the implementation details of the centralized and decentralized approaches as similar as possible to those of the MNS approach, to facilitate direct comparability. Finally, we describe the details of our simulation setup and the types of experiments conducted.

2.1 Coverage Task

We define the coverage task as uniform environment exploration—the robots should collectively visit every portion of the environment, and spend equal time visiting each portion. The environment is an enclosed square arena with randomly distributed small obstacles. The portions of the arena that need to be visited are the cells of a 16×16 overlay grid (i.e., 256 cells of equal size).

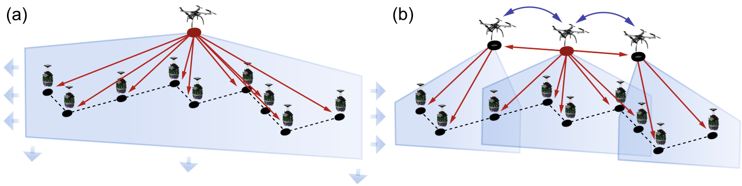

Fig. 1. Communication network topologies, and robot positions in the formations (centralized and MNS approaches). Red arrows show network connections. Light blue zones indicate approximate UAV field of view. Dashed black lines are not connections, but help to visualize the zigzag line that ground robots form. (a) Centralized approach. Connections are predetermined and static. The UAV is the master. (b) MNS approach. Dark blue arrows indicate UAV interchangeability. Network connections are self-organized; the UAV at the center is the brain. (Color figure online)

The environment is explored by differential drive ground robots, which are capable of detecting obstacles and other ground robots. In the MNS and centralized approaches, the ground robots are accompanied by camera-equipped UAVs that send motion instructions to other robots (both ground robots and other UAVs).

In all approaches, ground robots independently avoid obstacles and other ground robots. They are equipped with a ring of short-range proximity sensors. If a robot senses an object in a direction within 60° of its heading, it performs collision avoidance: it turns right if it senses objects only to the left of its heading; otherwise, it turns left. When not avoiding collisions, a robot follows its default motion behavior in the respective control approach.

In the centralized and MNS approaches, one robot—the master UAV and brain UAV respectively—is equipped with a simple motion controller to follow the arena perimeter. This controller moves the UAV forward in a straight line unless it detects a boundary, in which case it turns 90° to the left, and then moves forward again. This results in a counter-clockwise motion around the arena. As this perimeter-following behavior is deterministic, one loop of the master or brain UAV around the perimeter always takes the same amount of time, in both the centralized and MNS approaches. A master or brain UAV begins this perimeter-following behavior after the ground robots have established the target formation from their randomly distributed starting positions. The master or brain UAV then continues the behavior at a constant speed, irrespective of the speed of the other robots, until experiment termination. Given the size and shape of the target formation and the arena, this simple counter-clockwise path is sufficient to enable coverage.

In the centralized and MNS approaches, robots establish a target formation from randomly distributed starting positions, and then maintain that formation during coverage. In the target formation used here, ground robots are positioned in a zigzag line (see Fig. 1). The zigzag line formation is selected to reduce the occurrence of robot-robot collisions, compared to a straight line formation with smaller gaps between robots. The target formation of ground robots is identical in the centralized and MNS approaches. In the centralized approach, the master UAV is positioned above the center of the ground robots (see Fig. 1(a)). All robots in the centralized approach are always wirelessly connected to the master. In the MNS approach, the brain UAV is in the same position as the master UAV of the centralized approach; the other UAVs in the MNS approach are in a straight row above the ground robots (see Fig. 1(b)). In the MNS approach, the communication network topology is a caterpillar tree—i.e., a tree in which all inner nodes are on one central path, to which each leaf node is connected (see Fig. 1(b)). The centralized and MNS approaches use one and three UAVs, respectively, and each use nine ground robots. The decentralized approach uses nine ground robots, which perform a random walk without UAV guidance.

2.2 Approaches to Multi-robot Coverage

In the centralized approach, the default motion behavior for all robots is directly controlled by the *master* UAV that acts as central coordinating entity. The master UAV can directly communicate with all robots constantly, and can always see all robots and the full environment, regardless of its position. At the beginning of an experiment, the master sends all robots motion instructions, to move them into the target formation. The master then uses its perimeter-following behavior to follow the arena perimeter, while simultaneously sending all robots motion instructions, to maintain the target formation (relative to the master UAV).

In the decentralized approach, the default motion behavior for all ground robots is simply forward motion. At initiation, the robots are distributed randomly and begin moving in random directions. They only change direction as a result of collision avoidance. Due to the density of obstacles in the environment, collision avoidance is sufficient to change the robots' directions frequently enough for environment exploration.

In our MNS approach, the default motion behavior for non-brain robots is received from parents in the communication network. Our MNS approach is based on the existing concept of 'mergeable nervous systems' [15], for physically connected robots that we have extended to wireless connections in prior work [25]. In this approach, a heterogeneous swarm of UAVs and ground robots forms a target communication network topology through a self-organized process, and then uses this network to pass motion instructions between neighbors, moving robots into positions and orientations that match a given target formation. Please refer to [25] for details of the process by which the MNS is established and maintained. In the approach, a UAV can establish links with ground robots in its field of view, and can establish links with other UAVs when there is a shared ground robot in both their fields of view. In the experiments here, robots initially use the MNS process to establish the communication network and target formation. Then, the UAV that has become the brain (one of the three UAVs) begins to follow the arena perimeter. As the brain moves, it sends motion instructions to each of its children, which subsequently send motion instructions to their own children, thereby moving the whole formation.

2.3 Experiment Setup

The experiments are conducted using the ARGoS multi-robot simulator [17], with robot models implemented using an extension [2,3]. The 4×4 m^2 arena is fully enclosed, with its bottom-left corner at $(0, 0)$ of the coordinate frame. Static $4 \times 4 \times 2$ cm^3 obstacles are positioned randomly in the 3.7×3.7 m^2 center of the arena (with uniform distribution). The arena has a 16×16 overlay grid (with 0.25×0.25 m^2 cells). The UAV model has a maximum speed of 7.4 cm/s, and is equipped with a downward-facing camera. In the MNS approach, each UAV views a 1.5×1.75 m^2 rectangular ground area, at the 1.5 m flight altitude used in the experiments. Collectively, the three UAVs in the default MNS formation have a 1.5×2.75 m^2 view. By contrast, the UAV in the centralized approach has a full view of the arena at all times. The ground robot model has an average speed of 6.8 cm/s, and is equipped with a ring of 12 outward-facing proximity sensors with a 5.0 cm range. Ground robots are topped with fiducial markers encoding unique IDs, which the UAVs use to detect the relative positions and orientations of the ground robots. In our setup, UAVs are unable to detect obstacles. In the MNS approach, the communication range for UAVs and ground robots is 1 m. In the centralized approach, the master UAV has unlimited communication with all ground robots. The mechanical bodies of UAVs and ground robots are represented by simple 2.5 cm radius cylinders. In all approaches, if a ground robot reaches the arena boundary, its normal motion behavior is temporarily overridden—it turns to a random direction in the 180° range facing away from the boundary, then drives straight forward.

3 Results

In this section, we give the results of our experiments testing performance, scalability, and fault tolerance. In all experiments, we record robot positions. In all three approaches, ground robots initially face random directions and are positioned randomly in a 1.0×1.25 m^2 rectangular area against the southern arena boundary, following a uniform distribution. In the centralized and MNS approaches, the UAVs are positioned above the ground robots, near the southern boundary. Once the formation is established, the master or brain UAV has the southern boundary in view, and therefore turns left to start following the arena perimeter. For the centralized and MNS approaches, we define a *round* as one complete loop around the arena perimeter.

3.1 Performance

The performance experiments compare the three approaches in terms of *coverage percentage* (i.e., the percentage of grid cells visited by at least one ground robot) and *coverage uniformity* (i.e., the uniformity of the total time robots spend in each grid cell), and in terms of the time and energy expended (according to potential consumption rates). We test the performance of the three approaches (centralized, MNS, decentralized), with three different numbers of obstacles (100, 200, 300)—in total nine performance experiments.

Real-world energy consumption of UAVs and ground robots can vary considerably. Therefore, we test five possible ratios of UAV-to-ground-robot energy consumption $\{0.5, 1, 2, 3, 4\}$, with the ground robot consumption rate fixed at 30 units per step. Ratios over 1 represent scenarios with small simple ground robots and powerful UAVs, and ratios of 1 and 0.5 represent scenarios with large complex ground robots (e.g., quadruped robots) and minimal lightweight UAVs. Experiments testing the MNS approach terminate at step 4710, at the completion of one round; others terminate when the robots have consumed the same total energy as the MNS approach, under the same energy ratio. For example, under energy ratio 0.5, if the MNS approach has consumed 100 energy units at step 4710, then 100 energy units is the energy budget for the other approaches under that ratio. For each experiment type, we execute 10 experiment runs for each energy ratio termination time.

Table 1. The coverage percentage results of the performance experiments.

		MNS		Centralized		Decentralized	
Ratio	Energy: $/10^6$ Units	Time (steps)	#Obstacles : Coverage percentage	Time (steps)	#Obstacles : Coverage Percentage	Time (steps)	#Obstacles : Coverage percentage
0.5	1.48365	4710	100 : 96.9% 200 : 95.7% 300 : 92.2%	5206	100 : 98.4% 200 : 97.7% 300 : 96.1%	5495	100 : 86.7% 200 : 80.5% 300 : 76.2%
1	1.6956	4710	100 : 96.9% 200 : 95.7% 300 : 92.2%	5652	100 : 98.8% 200 : 98.1% 300 : 96.7%	6280	100 : 90.2% 200 : 84.8% 300 : 80.1%
2	2.1195	4710	100 : 96.9% 200 : 95.7% 300 : 92.2%	6423	100 : 98.8% 200 : 98.1% 300 : 97.7%	7850	100 : 94.9% 200 : 92.2% 300 : 87.1%
3	2.5434	4710	100 : 96.9% 200 : 95.7% 300 : 92.2%	7065	100 : 98.8% 200 : 98.1% 300 : 97.7%	9420	100 : 97.2% 200 : 95.3% 300 : 91.8%
4	2.9673	4710	100 : 96.9% 200 : 95.7% 300 : 92.2%	7609	100 : 99.2% 200 : 98.4% 300 : 98.1%	10990	100 : 98.4% 200 : 97.3% 300 : 94.5%

Coverage Percentage. We compare the coverage percentage of the three approaches under equal energy expenditure. Table 1 shows that the centralized approach outperforms the other approaches for all energy ratios and obstacle densities, although its performance is only slightly better than that of the MNS approach. When the energy ratio is less than 3, the MNS approach outperforms the decentralized approach for all obstacle densities. If time expenditure is considered, Table 1 shows that, for energy ratios of 3 and 4, the decentralized approach takes more than twice as much time as the MNS approach and achieves only slightly better coverage percentage. Because of the high UAV energy cost in these cases, the decentralized approach is allowed much longer exploration time than the MNS approach. Table 1 also shows that coverage percentage becomes lower for all approaches as obstacle density increases. Figure 2 shows the coverage percentage over time for the three approaches, in two obstacle setups at energy ratio 4. We report results only for one energy ratio because the graphs are similar for all energy ratios—the only difference being in the change in the performance gaps between the three approaches. We have chosen to report energy ratio 4 because it is the worst energy ratio for the MNS approach; the gap between the MNS approach and the better-performing centralized approach is largest in this ratio, and the gap between the MNS approach and the worse-performing decentralized is smallest in this ratio. Figure 2 therefore shows that, in all cases, the MNS approach substantially outperforms the decentralized approach, and the centralized approach slightly outperforms the MNS approach. Energy ratio 4 bears the worst performance for the MNS approach because the MNS uses three UAVs, as opposed to one UAV or no UAVs.

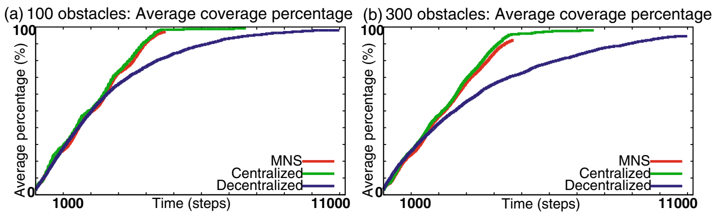

Fig. 2. Average coverage percentage for MNS, centralized, and decentralized approaches, with an energy ratio of 4. (a) 100 obstacles setup; (b) 300 obstacles setup.

Coverage Uniformity. For coverage uniformity, we assess the centralized, MNS and decentralized approaches at the timestep of the first complete MNS round, and additionally assess the decentralized experiments at energy exhaustion. For each run, $v_i \in \mathbf{v}$ is defined as the total time spent by all robots in cell i. The coverage uniformity p is the norm of $\mathbf{v}$, calculated as follows:

$$p = \sum_{i=1}^{256} \sqrt{|v_i - M(\mathbf{v})|}, \tag{1}$$

where $M(\mathbf{v})$ is the median of $\mathbf{v}$. The smaller the value of p, the more uniformity between cells; the most uniform case is $p = 0$. Figure 3(a) shows the coverage uniformity p of all three approaches, at the step of the first MNS round completion (step 4710). The centralized approach is the most uniform (i.e., the smallest p, on average). Figure 3(a) shows that, in terms of coverage uniformity, the MNS approach substantially outperforms the decentralized approach, and the centralized approach slightly outperforms the MNS approach. While the MNS and centralized approaches have approximately similar uniformity in later rounds (i.e., later in time), the uniformity of the decentralized approach becomes steadily worse over time, as shown in Fig. 3(b). The worsening of uniformity over time, in the decentralized approach, is most pronounced in the highest obstacle density. Figure 3(a) also shows that, at the completion of the first MNS round, the uniformity of all three approaches worsens as obstacle density increases.

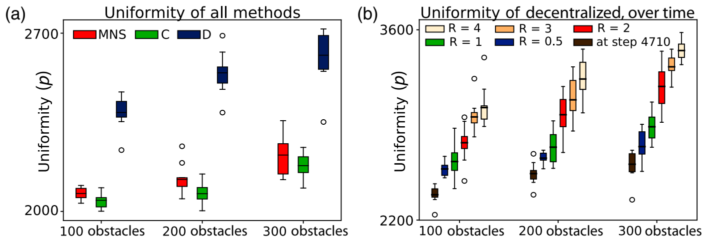

Fig. 3. Coverage uniformity (lowest p is the most uniform). (a) Uniformity p of all approaches (MNS; centralized, C; and decentralized, D), at the timestep of the first MNS round completion (step 4710). (b) Uniformity p of the decentralized approach, over time. Uniformity p at the step of energy exhaustion for each energy ratio (R)—i.e., p at termination—is compared to p at the step of the first MNS round completion (step 4710) for all five energy ratios—i.e., p early in the run.

3.2 Scalability

The scalability and fault tolerance experiments test whether the MNS approach displays features that would typically be observed in decentralized robot systems (cf. [9]). We evaluate scalability in the MNS approach in terms of communication (i.e., the number of messages exchanged) and interference (i.e., the number of robot-robot collisions). The scalability experiments are conducted in an arena without obstacles, with three different swarm sizes that are arranged in the same type of target formation as the default (caterpillar tree, zigzag line—see Fig. 1), with 10 runs per swarm size. The sizes are: 1) two UAVs, four ground robots; 2) four UAVs, eight ground robots; 3) six UAVs, twelve ground robots. Figure 4(a) shows that the number of messages increases linearly with increasing swarm size, as robots communicate only with their neighbors in the network. Figure 4(b) shows that the number of collisions increases steadily at the beginning of the experiments, as the MNS formation is being established. Once the MNS is established and begins to explore the environment, no further robot–robot collisions are observed.

3.3 Fault Tolerance

The ability of the MNS approach to replace a robot after failure, or to repair a broken network connection, has been demonstrated in [25]. In this paper, we investigate the fault tolerance of the MNS approach when ground robots fail and cannot be replaced or repaired, evaluated according to connectivity and coverage percentage. When a ground robot fails, its network link(s) are broken, and it can no longer move or communicate with any robots. We use the default MNS formation (see Fig. 1) and the setup with 100 obstacles, and impose failure at step 400. We test failure of the following numbers of ground robots (10 runs each), out of 12 total robots in the swarm: 1, 3, 5, 7, 8. We assess the impact of the failures on coverage percentage results. As UAV-to-UAV connections are established indirectly when mutually viewing ground robots, we also assess the impact of the ground robot failures on parent connectivity (i.e., whether the brain UAV maintains communication with the other UAVs, which are parent nodes in the communication network). Figure 4(c) shows that coverage percentage decreases as the number of failures increases. Figure 4(d) shows that parent connectivity is maintained in all cases of 5, 3, or 1 failure(s). In cases of 7 or 8 failures—in which more than half of the swarm fails—connectivity is maintained in 70% and 50% of runs, respectively.

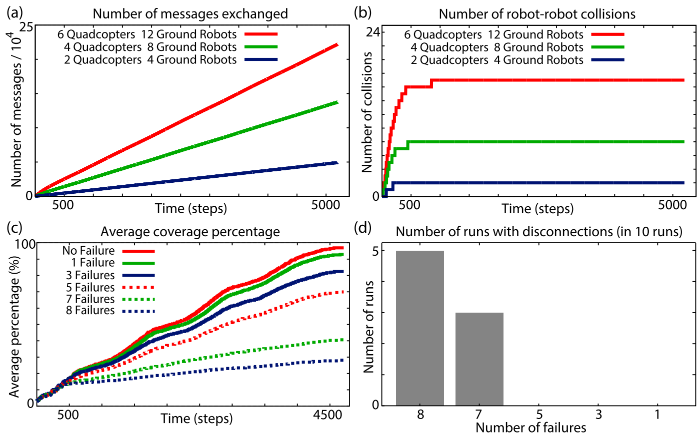

Fig. 4. (a–b) Scalability. (a) Number of messages exchanged in the MNS over time. (b) Number of robot-robot collisions over time. (c–d) Fault tolerance. (c) Coverage percentage over time, in MNSs with varying number of failing ground robots. (d) Number of runs that suffer a brain–parent disconnection, according to number of ground robot failures (out of 10 total runs for each number of failures).

4 Discussion

In terms of coverage percentage and coverage uniformity, the MNS approach substantially outperforms the decentralized approach. The decentralized approach requires approximately twice as much time as the MNS approach to reach similar coverage percentage, when the energy consumption ratio is at least 3 (see Table 1). If the ratio is less than 3, the decentralized approach never reaches the coverage percentage of the MNS approach, due to exhaustion of the energy budget. The MNS approach also achieves better coverage uniformity than the decentralized approach; coverage uniformity in the decentralized approaches worsens over time. The lower performance of the decentralized approach is due to the uneven distribution of robots that occurs during a random walk. As expected, the centralized approach outperforms the other two approaches in coverage percentage and coverage uniformity. However, the performance difference between the centralized and MNS approaches is relatively small, compared to that between the MNS and decentralized.

The scalability of the MNS approach, in terms of number of messages exchanged and robot-robot collisions, is good (see Fig. 4(a,b)). The number of messages increases linearly with increasing swarm size, and no robot-robot collisions are observed after the MNS is established, in any swarm size. The MNS approach is also fault-tolerant, in terms of connectivity and coverage performance. Substantial drops in performance only occur when more than half of the swarm fails. Connectivity in the MNS approach might be an advantage over the decentralized approach in consensus achievement tasks (e.g. collective decision making or collective sensing). As the MNS approach recovers from brain failure (see [25]), it also has an advantage over centralized approaches, in which all robots fail if the master UAV fails.

A possible direction for future development of our MNS approach would be to adapt the target formation on the fly. In this case, if failures are detected, the MNS could switch to a new formation shape that is better suited to the remaining swarm size. In future work, we will extend our MNS approach to apply it to tasks such as collective sensing, or localization and mapping.

5 Conclusions

We have presented an MNS approach to multi-robot coverage, and tested its coverage performance against strictly centralized and strictly decentralized approaches. Our results indicate that the MNS approach significantly outperforms the decentralized approach, but is slightly outperformed by the centralized approach. We have also tested the MNS approach for its performance in terms of scalabilty and fault tolerance—two features that are difficult to obtain with a centralized approach. Our results show that the MNS approach scales linearly in terms of inter-robot communication, and that its performance and connectivity are robust to failures if less than 50% of the ground robots fail. Overall, the results demonstrate that the MNS approach successfully combines aspects of centralized and decentralized control in a coverage task, as it achieves high performance (similar to centralized approaches), and achieves scalability and fault tolerance (similar to decentralized approaches).

Acknowledgements. This work is partially supported by the Program of Concerted Research Actions (ARC) of the Université libre de Bruxelles; by the Ontario Trillium Scholarship Program through the University of Ottawa and the Government of Ontario, Canada; by the Office of Naval Research Global (Award N62909-19-1-2024); by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 846009; and by the China Scholarship Council (grant number 201706270186). Mary Katherine Heinrich and Marco Dorigo acknowledge support from the Belgian F.R.S.-FNRS, of which they are a Postdoctoral Researcher and a Research Director respectively.

References

1. Albani, D., Nardi, D., Trianni, V.: Field coverage and weed mapping by UAV swarms. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4319–4325. IEEE (2017)
2. Allwright, M., Bhalla, N., Pinciroli, C., Dorigo, M.: ARGoS plug-ins for experiments in autonomous construction. Technical reports, TR/IRIDIA/2018-007, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium (2018)
3. Allwright, M., Bhalla, N., Pinciroli, C., Dorigo, M.: Simulating multi-robot construction in ARGoS. In: Dorigo, M., Birattari, M., Blum, C., Christensen, A.L., Reina, A., Trianni, V. (eds.) ANTS 2018. LNCS, vol. 11172, pp. 188–200. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00533-7_15
4. Almadhoun, R., Taha, T., Seneviratne, L., Zweiri, Y.: A survey on multi-robot coverage path planning for model reconstruction and mapping. SN Appl. Sci. **1**(8), 1–24 (2019). https://doi.org/10.1007/s42452-019-0872-y
5. Dewangan, R.K., Shukla, A., Godfrey, W.W.: Survey on prioritized multi robot path planning. In: IEEE International Conference on Smart Technologies and Management for Computing, Communication, Controls, Energy and Materials (ICSTM), pp. 423–428. IEEE (2017)
6. Dieber, B., Micheloni, C., Rinner, B.: Resource-aware coverage and task assignment in visual sensor networks. IEEE Trans. Circuits Syst. Video Technol. **21**(10), 1424–1437 (2011)
7. Galceran, E., Carreras, M.: A survey on coverage path planning for robotics. Robot. Auton. Syst. **61**(12), 1258–1276 (2013)
8. Ge, S.S., Fua, C.H.: Complete multi-robot coverage of unknown environments with minimum repeated coverage. In: Proceedings of the 2005 IEEE International Conference on Robotics and Automation, pp. 715–720. IEEE (2005)
9. Hamann, H.: Swarm Robotics: A Formal Approach. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-74528-2
10. Ichikawa, S., Hara, F.: Characteristics of object-searching and object-fetching behaviors of multi-robot system using local communication. In: Proceedings of 1999 IEEE International Conference on Systems, Man, and Cybernetics, vol. 4, pp. 775–781. IEEE (1999)
11. Juliá, M., Gil, A., Reinoso, O.: A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. Auton. Robots **33**(4), 427–444 (2012)

12. Koenig, S., Liu, Y.: Terrain coverage with ant robots: a simulation study. In: Proceedings of the Fifth International Conference on Autonomous Agents, pp. 600–607. Association for Computing Machinery (2001)
13. Laouici, Z., Mami, M.A., Khelfi, M.F.: Cooperative approach for an optimal area coverage and connectivity in multi-robot systems. In: 2015 International Conference on Advanced Robotics (ICAR), pp. 176–181. IEEE (2015)
14. Luo, W., Sycara, K.: Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 6359–6364. IEEE (2018)
15. Mathews, N., Christensen, A.L., O'Grady, R., Mondada, F., Dorigo, M.: Mergeable nervous systems for robots. Nat. Commun. **8**, 439 (2017)
16. Panerati, J., Gianoli, L., Pinciroli, C., Shabah, A., Nicolescu, G., Beltrame, G.: From swarms to stars: task coverage in robot swarms with connectivity constraints. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 7674–7681. IEEE (2018)
17. Pinciroli, C., et al.: ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intell. **6**(4), 271–295 (2012)
18. Rekleitis, I., Lee-Shue, V., New, A.P., Choset, H.: Limited communication, multi-robot team based coverage. In: IEEE International Conference on Robotics and Automation, 2004, Proceedings. ICRA 2004, vol. 4, pp. 3462–3468. IEEE (2004)
19. Santos, M., Egerstedt, M.: Coverage control for multi-robot teams with heterogeneous sensing capabilities using limited communications. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5313–5319. IEEE (2018)
20. Sartoretti, G., et al.: PRIMAL: pathfinding via reinforcement and imitation multi-agent learning. IEEE Robot. Autom. Lett. **4**(3), 2378–2385 (2019)
21. Schroeder, A., Ramakrishnan, S., Manish, K., Trease, B.: Efficient spatial coverage by a robot swarm based on an ant foraging model and the lévy distribution. Swarm Intel. **11**(1), 39–69 (2017)
22. Siligardi, L., et al.: Robust area coverage with connectivity maintenance. In: 2019 International Conference on Robotics and Automation (ICRA), pp. 2202–2208. IEEE (2019)
23. Wang, X., Han, S., Wu, Y., Wang, X.: Coverage and energy consumption control in mobile heterogeneous wireless sensor networks. IEEE Trans. Autom. Control **58**(4), 975–988 (2012)
24. Zafar, M.N., Mohanta, J.C.: Methodology for path planning and optimization of mobile robots: a review. Procedia Comput. Sci. **133**, 141–152 (2018)
25. Zhu, W., Allwright, M., Heinrich, M.K., Oğuz, S., Christensen, A.L., Dorigo, M.: Formation control of UAVs and mobile robots using self-organized communication topologies. In: Swarm Intelligence - Proceedings of ANTS 2020 - Twelfth International Conference. Lecture Notes in Computer Science. Springer, Heidelberg (2020)