

# SAGA: Drone Swarms in the Field

Dario Albani<sup>1,2</sup> and Vito Trianni<sup>1</sup>

**Abstract**—Wants to have the grass greener on ‘your’ side? Reduce the costs, improve the crops and say goodbye to annoying weeds, the objective of the SAGA (Swarm Robotics for Agricultural Applications) project is clear. Robots are the new frontier for agriculture, and robot swarms are emerging on the horizon. The SAGA project is among the firsts in deploying robot swarms in the field: gathering inspiration from the behaviours of ants and bees [1] [2], it uses a swarm of small Unmanned Aerial Vehicles (UAVs) to process information on-board and to constantly monitor the presence of weeds and the status of the crops. In this paper, we introduce the concept behind the SAGA project. We start with the problem of uniform coverage and mapping, showing the advantages of swarms in term of accuracy. Then, we introduce the non-uniform coverage problem, where swarms greatly improve in terms of mission completion time.

## I. INTRODUCTION

Robots are the new frontier for agriculture and unmanned aerial vehicles (UAVs) are powerful tools for several applications where aerial monitoring is required [3]. The SAGA project tackles the problem of sugar beets fields infested by volunteer potatoes, a common benchmark in agricultural robotics. Volunteer potatoes originate from tubers that remained in the soil after harvesting. The next season sugar beets are grown in the same field but volunteer potatoes are a major threat. Remained in the soil, they spread diseases and facilitate harmful soil nematodes. Regulations of some countries obligate farmers to control this weed, a costly operation—e.g. in the Netherlands it has an estimated cost of 300€ per hectare per growing season. Such a cost can be drastically reduced if farmers can rely on an efficient and automated weeding system as a system composed by multiple robots. Multi-robot systems are clearly faster and possibly more accurate than single robot operations. Additionally, redundancy and cooperation within a distributed robotic system can provide resilience and robustness to faults. Swarms can provide super-linear performance, so as to maximise the effectiveness of the group as a whole beyond the sensing and information-processing abilities of individual units (e.g., exploiting biological models of information retrieval and integration [4], [5]). Given the problem at hand, we believe that drones represent a valid solution and, in general, are expected to have a big impact on precision agriculture. UAVs can inspect the field at variable altitudes, obtaining data at different levels of detail. Flying robots avoid problems

from soil compaction given by unmanned ground vehicles (UGVs), they move faster and might care less about not harming the crops. A single UAV traversing and inspecting the field could provide the farmer with huge streams of data, two or three drones do it faster, a swarm, not only guarantees the same output in a fraction of the time but, thanks to distributed algorithms, drastically increases accuracy.

Controlling a bevy of robots is not trivial and requires specific techniques to ensure robustness to faults and scalability to large groups. To this end, SAGA’s UAVs employ an algorithm relying on random walks, an exploration technique based on stochastic movements [1] having the advantage of being simple, efficient and easily tunable. First, the field is discretized into a two dimensional grid with fixed size cells. At every decision step, a choice is made for the next cell to visit according to probabilities assigned to each cells in a way to optimally deal with individual effort and social influence. To avoid costly maneuvers, both in terms of energy and time, bigger probabilities are assigned to cells that are in the same direction of motion of an UAV. To obviate long relocations, only a subset of cells in the vicinity is taken into consideration. Moreover, drones collaborate and send information about their position and already visited locations. Such information is translated into attracting and repulsive forces: attraction toward points of interest and repulsion from each other. Altogether, these mechanisms can be controlled with few parameters that allow the system to be tuned and adapted to specific working conditions, achieving an efficient, reliable and fast mapping of the field [1].

Furthermore, as pointed out, drones have the ability to inspect the field at varying altitude drastically decreasing operation time. The presented application, as many others, requires high-resolution data only in certain areas while other areas can receive lower attention. Hence, UAVs can reduce the altitude to increase the resolution only where needed, making non-uniform strategies efficient both in time and energy expenditure [6], [7]. A close inspection guarantees higher resolution and bigger accuracy, but at the cost of higher inspection time due to the limited footprint of the camera. To obtain the best trade-off between accuracy and speed, the SAGA swarm implements a decentralised deployment strategy inspired by honeybees behaviour [8], that dynamically assigns UAVs to different areas to be monitored, and suitably re-assigns them to other areas when needed. In this way, UAVs inspect at low altitude only some areas, while mildly monitoring other parts of the field. Thanks to this self-organising strategy, the swarm can partition the monitoring task in an optimal way and autonomously allocates the required resources only where and when needed.

<sup>1</sup> Institute of Cognitive Sciences and Technologies, National Research Council of Italy (ISTC-CNR). Via San Martino della Battaglia 44, 00185 Rome, Italy. name.surname@istc.cnr.it

<sup>2</sup> Department of Computer, Control, and Management Engineering Antonio Ruberti at Sapienza University of Rome. Via Ariosto 25, 00185 Rome, Italy. albanidiag.uniroma1.it

## II. BACKGROUND

The SAGA project involves several state-of-the-art features that span through swarming, decision making and task allocation. Within this section, we provide a concise overview of these topics. Coverage and mapping are interrelated problems largely studied in multi-agent systems and robotics, although not specifically linked to agricultural applications [9], [10], [11]. Generally speaking, a certain number of points of interest (POIs, i.e., weeds) must be reached (coverage) and processed (mapping). This makes the problem similar to task allocation in multi-agent systems, which has been approached with different methods like Distributed Constraint Optimization (DCOP [12]), Distributed Pseudotree Optimization (DPOP [13]), the Contract Net Protocol [14] or bounty-based auction methods [15], to cite some. Few studies take a swarm robotics approach to coverage and mapping. In [16], the coverage problem is tackled through a flocking algorithm ensuring permanent connectivity among robots. Similarly in [17], multi-robot coverage is performed to maximise spread while maintaining connectivity among robots and reducing the communication overhead. The focus on connectivity is justified by the need to largely spread information within the group. In a previous study, we relax this constraint to provide more efficient coverage solutions, and we introduce re-broadcast protocols to account for limited communication ranges [1]. In the same work, in order to provide scalability and robustness to the system, we devise a stochastic exploration strategy based on a reinforced random walk [18], [19], [20] where individual UAVs follow a correlated random walk and interact with neighbours to avoid interferences. All previous cited works deal with the coverage and mapping problem using a two dimensional approach ignoring the third dimension or using it only partly. Non-uniform coverage strategies for UAVs have proven to be efficient [7], [6], they assume that a coarse estimation of the monitoring effort can be obtained by UAVs flying at a high altitude, so that images of the area are obtained at a somewhat low resolution, however sufficient to get hints about the requirements of that region. By flying at lower altitudes, inspection can be conducted with a higher resolution only in correspondence of interesting areas, resulting in substantial savings in the coverage time and effort.

## III. APPROACH

We consider the case of a bounded field with no obstacles. We divide the field in cells where the size of a cell is fixed and is related to the camera footprint installed on the robot—assuming fixed focal length. This said, we allow the drones in the swarm to fly at different, fixed altitudes translating in different overlapping field layers having cells of different size, bigger cells are associated to higher altitudes and includes several smaller cells that, on the other hand are generated for lower altitudes. Figure 1 shows an example of a two layer simulated environment used in our experiments: on the left, the overall field divided in a  $9 \times 9$  grid, on the right image details of a single high-altitude cell of the field

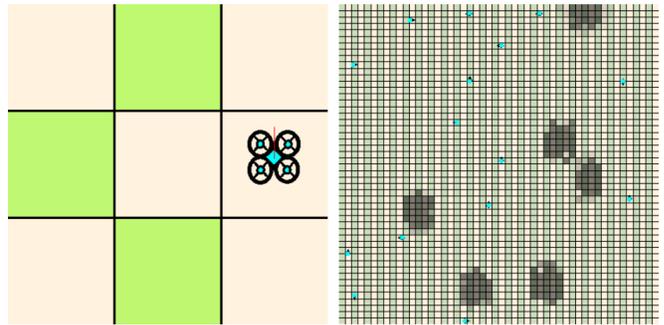


Fig. 1. A visualisation of the simulation environment. (Left) The high-altitude partition of the field in  $M=9$  areas, in green areas in which weed patches are present. (Right) The low-altitude partition of an area in  $M_c=50 \times 50$  cells, grey zones represent weed patches and blue dots represent drones.

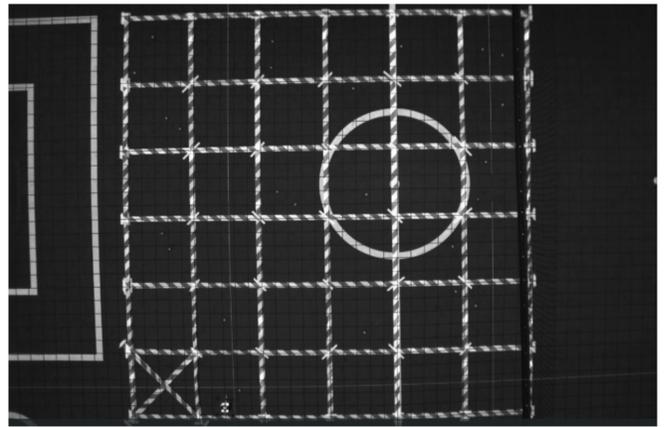


Fig. 2. A top-down view of some preliminary experiments developed at the Technical University of Eindhoven. The barricade tape has been used to aid a human viewer to reproduce the two dimensional grid although the drone – situated in the bottom left part of the field – is not using it and it is autonomously navigating the environment following the algorithm presented in [1].

as shown in the left image. The collective-level control is responsible for the overall mission accomplishment. Instead of a-priori planning the mission for the whole group, we will exploit swarm robotics techniques in which the group behaviour emerges from self-organisation, hence providing flexibility, robustness to faults and scalability with group size. Our goal is to devise collective strategies with an optimal trade-off between distributed exploration and timely weed recognition.

For activity coordination and efficient operations, communication among UAVs is essential. By broadcasting their absolute position, UAVs can implement collaborative avoidance strategies (e.g., the hybrid reciprocal velocity obstacle method [21]). Additionally, UAVs broadcast information about the visited cells and also about virtual “beacons” that UAVs place in areas requiring additional attention. In this work, we neglect collision avoidance issues, and we focus on the information required for monitoring and mapping. We assume that an agent  $h$  broadcasts a message after visiting a new cell. Such a message contains: (i) the current cell  $c_{ij}$ , (ii) the estimated weed density  $\rho_{ij}$  and confidence  $z_{ij}$ , (iii) the

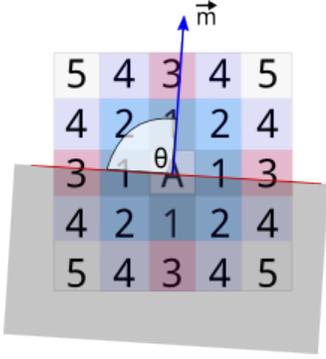


Fig. 3. Distance-based exploration pattern. In the image the central cell marked with an A represents the agent's position. Numbers indicate the priority of each location assigned according to the distance. Vector  $\vec{m}$  is the momentum of the agent, representing a possible directional bias. Cells that fall in the shadowed region are accepted only if no valid cell is found in the other semi-plane

next cell  $c'_{ij}$  chosen for visiting, (iv) the list of known active beacons  $\mathcal{B}$ , and (v) a unique identifier of the message. The ability to spread such information widely within the swarm is key for efficient coordination. It is therefore important to study how communication affects both the behaviour of the agents and the overall swarm efficiency. Information spreading depends on the communication range  $R_c$ , and on the chosen re-broadcasting protocol. The communication range and the density of agents in the field determines the features of the static interaction network, and the speed at which information propagates [22]. The re-broadcasting protocol determines how efficient the communication is in reaching agents via multi-hop forwarding. In this and in our previous work, we adopt concepts from information-centric mobile ad-hoc networks (ICMANET [23]) where we favour data content.

#### A. Uniform Field Coverage and Mapping

The basic strategy for the single layer (i.e. uniform) field coverage adopted by the agents is a simple correlated random walk [19], [20], on top of which we implement several mechanisms to improve the exploration efficiency.

The random walk for each UAV is implemented by selecting a random cell to be visited in the neighbourhood. To minimise the distance covered, the choice is made within sets of cells at increasing distance, as depicted in Fig. 3 left, where cells belonging to the same set are identified by the same ordinal number  $R$ . For instance, the set with  $R=1$  corresponds to the cells directly adjacent to the agent's current one. To choose the next cell, an agent makes a two-step decision. In the first step, it looks for a sufficient number of *valid* cells, that is, cells that have not been previously visited or that are not occupied/targeted by other agents, to the best of the local knowledge available. Validity of cells is checked sequentially for sets of increasing distance  $R$ , until a given number  $V$  of cells is discovered (in this study,  $V=1$ ). At this point, a set  $\mathcal{V}$  is defined including all valid cells within the maximum distance reached. Note that, in general,

$|\mathcal{V}| > V$ , as all the cells within a given distance are included, not limited to those that allow reaching the threshold  $V$ .

To avoid that the selection of valid cells interferes with possible directional biases for cell selection (i.e., when the first valid cells lay in the opposite direction), an agent divides the exploration plane in two parts (semi-planes), hence choosing valid cells first in the semi-plane complying with the directional bias, and only if unsuccessful, in the opposite semi-plane. More precisely, given the agent  $h$  position  $\vec{x}_h$ , a directional bias  $\vec{b}$  and the relative position of the cell  $\vec{r}_{ij} = \vec{x}_{ij} - \vec{x}_h$ , the semi-plane with priority is given by all cells  $c_{ij}$  for which  $\vec{b} \cdot \vec{r}_{ij} \geq 0$ . If no valid cell is available satisfying this conditions, then also the remaining cells are evaluated for inclusion in  $\mathcal{V}$

In the second step, a random choice is performed within  $\mathcal{V}$  to select the target cell. To implement a correlated random walk, we consider a unit vector  $\vec{m}_h$  representing the momentum of the agent  $h$ . Additionally, the influence from neighbouring agents is considered to avoid interferences and to bias random movements in areas with low agent density. To this end, we compute for each agent  $h$  a repulsion vector  $\vec{r}_h$  summing up repulsive forces from all agents  $k$  in the neighbourhood of agent  $h$ , with the module decaying according to a gaussian function with width  $\sigma_a$ . To account for both momentum and repulsion, a bias vector is computed as  $\vec{b}_h = \vec{m}_h + \vec{r}_h$  and a utility is assigned to each cell  $c_{ij} \in \mathcal{V}$  according to the angular difference  $\theta_{ij}$  between the cell and  $\vec{b}_h$  as follows:

$$u_{ij} = C(\theta_{ij}, p), \quad C(\theta, p) = \frac{1}{2\pi} \frac{1-p^2}{1+p^2-2p\cos\theta} \quad (1)$$

where  $C(\cdot)$  is the wrapped Cauchy density function with persistence  $p \in [0, 1[$  which determines the function skewness. For high values of persistence, the utility is very high only for cells aligned with  $\vec{b}_h$ , while for a low value of persistence the utility is more uniformly assigned despite the angular deviation.

#### B. Non-Uniform Field Coverage and Mapping

The non-uniform coverage strategy exploits all the degree of motion of the UAVs, and incorporates the uniform strategy only for the lowest layer. First of all, UAVs need to identify the areas of interest and estimate their initial utility, hence they start navigating the field at high altitude obtaining a coarse estimation of big portion of the field. On such a basis, they have to choose an area to inspect, possibly in collaboration with other agents so as to allocate part of the swarm to sub-portion of the field, overall minimising the completion time. The problem contains elements of collective decisions—when agents need to converge to a single area to inspect—and task allocation—when agents need to disperse on different areas, and can be suitably tackled following a design pattern inspired by the honeybee nest-site selection behaviour [5]. The design pattern defines a macroscopic model of the system dynamics, which—in its

most general form [2]—reads as follows:

$$\begin{cases} \dot{x}_i = \gamma_i x_u - \alpha_i x_i + \rho_i x_u x_i - \sum_{j=1}^M x_j \beta_{ji} x_i \\ x_u = 1 - \sum_{i=1}^M x_i, \quad \gamma_i, \alpha_i, \rho_i, \beta_{ij} \geq 0 \end{cases} \quad (2)$$

where  $x_i = n_i/N$  represents the fraction of agents deployed on area  $A_i$  (also referred to as “committed” agents [5]), and  $x_u$  represents the fraction of “uncommitted” agents, that is, agents available for deployment in any area. The macroscopic dynamics are determined by four concurrent processes: (i) uncommitted agents spontaneously enroll to area  $A_i$  at rate  $\gamma_i$ ; (ii) committed agents spontaneously abandon an area  $A_i$  at rate  $\alpha_i$ ; (iii) agents committed to area  $A_i$  can recruit uncommitted agents at rate  $\rho_i$ ; (iv) agents committed to area  $A_j$  “inhibit” agents committed to  $A_i$  at rate  $\beta_{ji}$ , so that the latter become uncommitted. Differently from previous studies, we extended the inhibition paradigm to act between any pair of committed agents, hence including both cross-inhibition ( $\beta_{ji} \neq 0$  for each  $j \neq i$ ) and self-inhibition ( $\psi_i \triangleq \beta_{ii} \neq 0$ ). Self-inhibition allows to finely control the number of agents working in the same area, procedure that comes in hand when moving out from abstract environments and considering real world applications with crowding, occlusions and collisions. This allow us to control the overall behaviour of the swarm and to avoid overcrowding with few parameters.

#### IV. CONCLUSION

In this extend abstract, we have discussed how the problem of coverage and mapping in precision agriculture can be addressed through a swarm robotics approach. We have designed a complete framework, from dynamic allocation to stochastic exploration, allowing UAVs to autonomously tackle the problem at hand. Having implemented such framework on real robots, we are currently moving toward field tests and believe that we will soon be able to demonstrate the capabilities of our approach on real platforms. The successful demonstration of the SAGA project raises the bar for swarm robotics research. The natural step forward is the extension of the agricultural swarms concept toward heterogeneous systems made of ground and flying robots, tackling all sort of tasks within the farm, from weed removal to pest control, from optimal usage of fertilisers to harvesting.

#### REFERENCES

- [1] Dario Albani, Daniele Nardi, and Vito Trianni. Field coverage and weed mapping by uav swarms. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 4319–4325. IEEE, 2017.
- [2] Andreagiovanni Reina, James AR Marshall, Vito Trianni, and Thomas Bose. Model of the best-of-n nest-site selection process in honeybees. *Physical Review E*, 95(5):052411, 2017.
- [3] G. Pajares. Overview and Current Status of Remote Sensing Applications Based on Unmanned Aerial Vehicles (UAVs). *Photogrammetric Engineering & Remote Sensing*, 81(4):281–330, 2015.
- [4] Andrea Baronchelli, Tao Gong, Andrea Puglisi, and Vittorio Loreto. Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences of the United States of America*, 107(6):2403–2407, February 2010.

- [5] Andreagiovanni Reina, Gabriele Valentini, Cristian Fernández-Oto, Marco Dorigo, and Vito Trianni. A Design Pattern for Decentralised Decision Making. *PLoS ONE*, 10(10):e0140950–18, October 2015.
- [6] S.A. Sadat, J. Wawerla, and R. Vaughan. Fractal trajectories for online non-uniform aerial coverage. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA 2015)*, pages 2971–2976. IEEE, 2015.
- [7] S.A. Sadat, J. Wawerla, and R.T. Vaughan. Recursive non-uniform coverage of unknown terrains for UAVs. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1742–1747. IEEE, 2014.
- [8] Dario Albani, Manoni Tiziano, Daniele Nardi, and Vito Trianni. Dynamic uav swarm deployment for non-uniform coverage. In *Seventeenth International Conference on Autonomous Agents and Multiagent Systems(AAMAS)*. ACM, 2018.
- [9] B. Hrotenok, S. Luke, K. Sullivan, and C. Vo. Collaborative foraging using beacons. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 3-Volume 3*, pages 1197–1204. International Foundation for Autonomous Agents and Multiagent Systems, 2010.
- [10] M. Paradzik and G. Ince. Multi-agent search strategy based on digital pheromones for uavs. In *Signal Processing and Communication Application Conference (SIU), 2016 24th*, pages 233–236. IEEE, 2016.
- [11] M. Popovic, G. Hitz, J. Nieto, I. Sa, R. Siegart, and E. Galceran. Online informative path planning for active classification using uavs. *arXiv preprint arXiv:1609.08446*, 2016.
- [12] M. Yokoo. *Distributed constraint satisfaction: foundations of cooperation in multi-agent systems*. Springer Science & Business Media, 2012.
- [13] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence, IJCAI’05*, pages 266–271. Morgan Kaufmann Publishers Inc., San Francisco, CA, 2005.
- [14] R.G. Smith. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.
- [15] D. Wicke, D. Freelan, and S. Luke. Bounty hunters and multiagent task allocation. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 387–394. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [16] E. Mathews, T. Graf, and K.S.S.B. Kulathunga. Biologically inspired swarm robotic network ensuring coverage and connectivity. In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, pages 84–90. IEEE, 2012.
- [17] Z. Laouici, M.A. Mami, and M.F. Khelifi. Cooperative approach for an optimal area coverage and connectivity in multi-robot systems. In *Advanced Robotics (ICAR), 2015 International Conference on*, pages 176–181. IEEE, 2015.
- [18] A. Stevens and H. Othmer. Aggregation, Blowup, and Collapse: The ABC’s of Taxis in Reinforced Random Walks. *Siam Journal on Applied Mathematics*, 57(4):1044–1081, 1997.
- [19] E.A. Codling, M.J. Plank, and S. Benhamou. Random walk models in biology. *Journal of The Royal Society Interface*, 5(25):813–834, 2008.
- [20] C. Dimidov, G. Oriolo, and V. Trianni. Random Walks in Swarm Robotics: An Experiment with Kilobots. In M. Dorigo, M. Birattari, X. Li, M. López-Ibáñez, K. Ohkura, C. Pinciroli, and T. Stützle, editors, *Proceedings of the 10th International Conference on Swarm Intelligence (ANTS 2016)*, pages 185–196. Springer International Publishing, 2016.
- [21] J. Snape, J. van den Berg, S.J. Guy, and D. Manocha. The Hybrid Reciprocal Velocity Obstacle. *IEEE Transactions on Robotics*, 27(4):696–706, 2011.
- [22] V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of Consensus in a Multi-Robot Network: From Abstract Models to Empirical Validation. *IEEE Robotics and Automation Letters*, 1(1):348–353, 2016.
- [23] X. Liu, Z. Li, P. Yang, and Y. Dong. Information-centric mobile ad hoc networks and content routing: A survey. *Ad Hoc Networks*, 58:255–268, 2017.