## 8.7 SWARM ROBOTICS

### ABOUT THIS CHAPTER

**_Why do we need this chapter?_**

*Swarm robotics is an important application of Organic Computing because it exploits the main advantages of OC in an embodied system with a vast variety of potential applications (e.g., monitoring of difficult to access areas, self-organised construction, automatic farming, exploration, cleaning). On the one hand, swarm robotics can serve as a very accessible use case of OC that motivates studying OC. On the other hand, swarm robotics in itself poses new challenges that require sophisticated, novel methods of self-organisation to solve them. Moreover, there are not many textbooks available that cover the state-of-the-art of swarm robotics.*

**_The content in a nutshell_**

- *Swarm robotics has been used, for example, for collective construction and mixed societies.*
- *The challenge is to switch between the micro-level (individual robot) and macro-level (whole swarm) and to deal with local information only.*
- *Swarm robotics systems can be modelled by rate equations and methods from opinion dynamics to support the design of swarm robotic systems.*
- *The task of collective decision-making is solved by the majority rule, the voter model, or the Hegselmann-Krause model.*

**_Who should read it?_**

*Students, young researchers, and interested engineers should read this chapter either as an interesting use case of OC or as a promising robotic concept applying methods of OC to real world problems. Interested readers will enjoy this chapter because it is one of the few available textbook texts on the state-of-the-art of swarm robotics.*

### INTRODUCTION

The occurrence of large numbers of biological organisms that effortlessly coordinate their actions, for instance in the realm of social insects (see also section 2.2.2), inspired the field of swarm robotics. Here, large numbers of robots interact and collaborate to achieve common goals. Technically speaking, robotic swarms are embodied distributed systems comprised of great numbers of autonomous elements. Yet, the term "swarm" not only refers to the great number of robots that interact. It carries a deeper meaning. To this end, it is interesting to note that contributions, for instance, from the field of swarm intelligence, do not attempt to close this conceptual gap. Instead, they focus on elaborating on behavioural descriptions of the swarm individuals and of the swarm as a whole. A common biological definition refers to swarms as an "aggregation often combined with collective motion", which combines both the great number of individuals of a swarm as well as their coordination in movement. As swarm robotics has a wider scope than collective movement, it also considers other goals that can only be accomplished by robotic collectives, which is appropriately captured by the following definition by Dorigo and Sahin [DORIGO 2004]: "Swarm robotics is the study of how a large number of relatively simple physically embodied agents can be designed such that a desired collective behaviour emerges from the local interactions among agents and between the agents and the environment." As such, robotic swarms can be considered a subset of OC systems that specifically consider large numbers of interacting embodied system components. Chapter 3 of this book helps to understand this notion of robotic swarms as OC systems. Chapter 4 provides concrete methods to quantify the emergent effects that may occur based on the collective behaviour exhibited by swarms of robots. Sometimes robotic units that are described as “relatively simple” are also referred to as reactive, i.e. they do not pursue elaborate strategies and as individuals they might not be very efficient with respect to a given task. In order to locally interact, the robots need the capabilities to sense and emit signals (see chapter 5). The resulting communication processes are considered a key feature in swarm systems. Another aspect of swarms is the need to cope with uncertainties as only local information is obtained and global states need to be estimated. Robotic swarms can approach this requirement at the population-level or at the level of its individuals (see chapter 7).

Three main benefits of swarm robotic systems are typically emphasised [DORIGO 2004]:

1. Robotic swarms are robust, i.e. they are fault tolerant and fail-safe due to massive redundancy and avoidance of single-points of failure. As the individual robots are usually identical in construction, each robot can substitute any other one. There is no central instance controlling the swarm. Rather, its behaviour emerges from the many control mechanisms of all the involved robots. In OC terminology, swarms implement strongly self-organised systems. Therefore, single failing units only have local impact, which is quickly compensated for by the remainder of the swarm. In fact, depending on the robot swarm configuration, a certain percentage of units may fail without jeopardising its effectiveness, i.e. it will still accomplish its task but may be less efficient in doing so. In order to better understand robustness of robotic swarms, one can consider them as networks of interacting units. These networks can be analysed by means of graph theory. For instance, for a specific task, it might be essential for a swarm to maintain a formation in which all the units can communicate with each other, either through intermediaries or directly. In this situation, any removal of such communication links between pairs of robots has to be avoided. Otherwise, the network would be split into two graph partitions and robustness could not be guaranteed. Determining exactly whether or not a single link is crucial for maintaining stability of a system is a non-trivial task that is aggravated by the decentralised setting. As an alternative, one can ensure a certain minimal density of robots within a specified area or volume, which statistically implies a certain average number of neighbours and, thus, a certain average degree of connectivity of each node. Based on this constraint, a probabilistic statement about the robustness of a robotic swarm can be provided. Space flight provides an illustrative example for robustness by redundancy: Nowadays, space probes are sent to remote places one at a time. Only one fatal failure of such a probe renders the whole exploration mission a failure as well. Therefore, a lot of effort is invested into the perfection of the probes’ design and into its tests. As an alternative, large numbers of small, simple, and cheap probes could be manufactured and launched into space. The simplicity of these multiple drones (for instance their limited battery life, noisy transmitters or lower resolution sensors) is overcome by collaboration. At the same time, the loss of some of the probes would not endanger the whole mission for as long as the remaining swarm can effectively pursue its task.
2. Robotic swarms are flexible. By means of collaboration, robot swarms can adapt to a wide range of tasks without the need for individuals with specialised capabilities. Due to this paradigm of identically designed swarm individuals, each robot can take anyone else’s place. This local flexibility also translates to the whole swarm, as its adaptation to new or changing goals merely requires the propagation of the respective information and taking the corresponding local actions. When coordinating accordingly, a robot swarm can flexibly accomplish tasks that no single unit would be able to achieve such as the collective transport of objects or the formation of multi-hop communication lines to compensate for short transmission ranges. Swarms of self-assembling robots can even flexibly adapt to challenges of the robots’ physical dimensions, for instance by joining one large robotic body that can bridge across cavities or climb over obstacles.

3. Robotic swarms scale well. As each swarm robot only interacts with its local environment, the algorithms to drive the robot swarm's behaviour can run on arbitrarily large numbers of entities. Algorithms targeting global interactions, such as broadcasting messages from one individual to the whole swarm, are precluded. In fact, scalability of some algorithms necessitates constant densities of the robots. However, physical dimensions and hardware limitations often provide clear boundaries regarding densities. Also, the swarm can be scaled to arbitrary numbers, if the operational space is increased proportionally to the size of the robot swarm.

### EXAMPLE TASKS

Many tasks for swarm robotics have been proposed and investigated. The body of literature on swarm robotics grows continuously. We refer to the survey paper of Brambilla et al. [BRAMBILLA 2013] for an exhaustive overview. In the following we only look into two scenarios as example tasks: cooperative construction and mixed societies.

In *cooperative construction*, the swarm robots are asked to construct a building or a structure [GERLING 2016]. The advantage of multiple robots in construction is that the construction process can be parallelised. Several robots pick up building material, transport it, and place it at its destination. Without central control and based on local information only, it is not straightforward to design appropriate control algorithms. In addition, the task description should also specify whether the desired structure is fully predefined down to the lowest level or only roughly and the construction robots can still influence the structure at site. For the case of a fully predetermined structure, that is, there is a blueprint, an approach has been proposed that takes such a blueprint as input and generates appropriate rules for an individual robot [WERFEL 2014]. These rules describe to where a robot should move for a given local perception, when to pick up building material, and when to place it. These rules need to be defined carefully to prevent deadlocks that might emerge due to parallelism and the decentralised approach. In the approach of Werfel et al. [WERFEL 2014], this challenge is resolved by an offline algorithm that operates as compiler, which takes the blueprint as input and then calculates individual robot controllers that guarantee to prevent deadlocks.

As second example scenario, we discuss the so-called *mixed societies*. One speaks of mixed societies once the robot swarm coexists with another group of agents, such as animals or natural plants. One of the first proposals of a mixed society was the work by Caprari et al. [CAPRARI 2005]. They mixed a group of cockroaches with small mobile robots. At first glance such a setup may seem rather absurd and academic, however, there are useful applications of similar setups. Caprari et al. showed that they were able to control the group of cockroaches using the robot swarm. The robots were accepted as mates by the cockroaches because they were clad in paper treated with pheromones. The investigated cockroach behaviour was an aggregation behaviour. They provided the cockroaches with two types of potential shelters (small, circular, opaque plastic roofs). Shelter $A$ had about the desired size for the given cockroach group and shelter $B$ was too big. The natural behaviour of the cockroaches would be to agree on shelter $A$ and aggregate beneath it. It was shown that the robots were able to influence the cockroaches by moving between the two shelters. The cockroaches ended up aggregating beneath the undesired shelter $B$ – an unnatural behaviour. Other approaches to mixed societies work with honeybees and fish [SCHMICKL 2013] and even natural plants [HAMANN 2015].

## MICRO-MACRO PROBLEM AND LOCAL SAMPLING[1]

Swarms naturally illustrate the transition from the microscopic level, i.e. the level of the individual robot, to the macroscopic level, i.e. the level of the whole robotic swarm. The microscopic level is determined by a robot's state, its knowledge, its view, its actions, and possibly some uncertainty. The macroscopic level can be described as the states of all the swarm robots and their interrelations. Therefore, the macroscopic level assumes a global view, it has full knowledge about the system, a comprehensive overview of the state of accomplishment with respect to a given, overall task. The duality of levels, and especially their mutual interdependency, challenges the designer of robot swarms, as well as the user to command a swarm at presumably high levels of control [VON MAMMEN 2016]. As a robotic swarm is meant to pursue global goals, its tasks are defined at the macro-level, whereas the implementation of the swarm's individuals' behaviours happens at the micro-level. Descriptions of robot swarms at both levels are also utilised to classify a swarm model: Macroscopic models do not flesh out the details of individual swarm robots, but rather they focus on detailing the overall goals. In contrast, microscopic models explicitly provide details of all the involved robots.

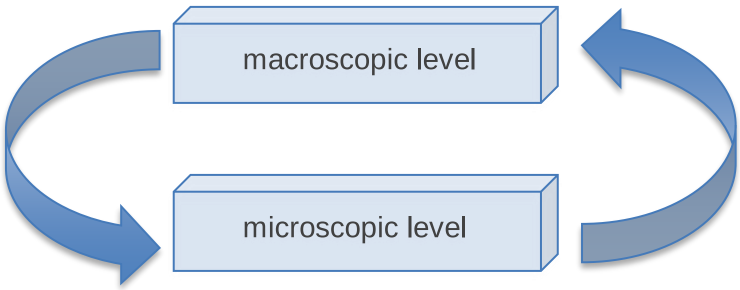


*Figure 1: Micro-macro problem.*

The design of concrete controllers for swarm robots considers, of course, the micro-level, where the lack of global information poses a difficult challenge. For instance, it might be crucial for the individual robot to know, where the currently greatest cluster of robots is located. Yet, an individual robot is limited in terms of perception and because of the exclusive communication to its local neighbourhood. This dilemma also applies to agents with greater cognitive capabilities, as for instance experienced by everyone who has ever tried to orientate himself at the corner of a street but without a city map. Similarly, human society struggled in the 16th century in the attempt to determine earth's coordinates relative to its surrounding stars and planets. Predicting the results of democratic elections poses another analogous problem. Only a small number of voters can be interviewed and it is of great importance to ask a representative crowd. A comparable situation occurs when predicting the market success of a new product. Here, too, only a few potential customers can be interviewed.

In swarm robotics, the individuals can only gather data from their immediate neighbourhood, and global information is not readily available. Often, swarm robots do not even directly communicate with their robot neighbours to extend their knowledge. Rather, they obtain information by looking at their environment, where they perceive and place signals to communicate indirectly. This form of

[1] The micro-macro problem is also discussed in section 4.6 in the context of "controlled emergence".

communication is referred to as “stigmergy” and it is frequently found in social insect swarms that, for instance, communicate through pheromone trails to direct their foraging efforts (see also section 2.2.2). Swarm robots may also make decisions based on their neighbours’ states rather than on messages sent back and forth. For instance, in order to move in formations, individual robots align themselves in accordance with their neighbours’ orientation and speed. Such local sampling works, as the velocities of one’s neighbours usually provide a good estimate of the movement of the overall swarm. Concisely showing the benefit of this assumption and also that it does not lead to catastrophic results, however, is a non-trivial problem.

Self-organisation is a key aspect to support the effectiveness of such local sampling processes (see also chapter 4). Consider, for instance, the emergence of order in pedestrian flows. Although the majority of pedestrians are preoccupied with something completely different, pedestrians move collectively, seamlessly and effortlessly. Frequently, two streams of pedestrians even emerge that follow the same general walking direction. The locally sampled data that determines your movement is the walking direction of the person in front of you – if he approaches you, you step aside. If he is walking in your direction, you follow. As a result, ordered flows emerge based on simple rules that rely on local sampling.

### MODELLING APPROACHES

Modelling of swarm robotic systems is an important aspect of swarm robotics. Due to the micro-macro problem, it is challenging to design control algorithms and to implement swarm systems. Models of swarm systems can help to support the algorithm design phase. However, modelling itself is a scientific challenge due to the micro-macro problem. In the following we quickly introduce three representative modelling techniques for swarm robotics.

Rate equations are one of the first proposed modelling approaches [MARTINOLI 2004]. The technique originates from chemistry. A rate equation can be as simple as r=kAB. It describes the rate of a chemical reaction between two reactants. Say, the concentrations of the chemical species are A and B. k is the rate coefficient. Say, we have $A+B\rightarrow C$ then an ordinary differential equation (ODE) is defined by the reaction rate:

$\frac{dC}{dt}=kAB$,

for concentrations A, B, and C. We are not going into details here but the main underlying concept is that of the law of mass action. The application of rate equations to swarm robotics is based on interpreting the concentrations differently. Instead of concentrations of chemical species we have swarm fractions of robots that are in certain states. Instead of chemical reactions we have state transitions as effects of robot-robot interactions. For example, if two robots in state exploring (E) approach each other both of them make a transition to state collision avoiding (C). So, we would model that by

$\frac{dC}{dt}=2rE$ and $\frac{dE}{dt}=-2rE$

for an appropriate reaction rate r. Obviously there should also be a reaction back to the exploring state, which could be modelled with a time-delay equation.

Given an appropriate ODE or an ODE system together with a given initial state of the system (e.g., the initial state of the robots), we can integrate forward in time and try to predict future states of the swarm system. This modelling approach is macroscopic, probabilistic, and non-spatial because we assume a well-mixed system (i.e., robot states are not assumed to be correlated in space). Using the rate equation approach, it is often difficult to relate parameters and other features of the model to features of the control algorithm, that is, at the microscopic level of the system. Several examples of how to use the rate equation approach are given in Lerman et al. [LERMAN 2005]. The rate equation approach supports the designer during the algorithm design phase to estimate the macroscopic effects starting from a microscopic robot control. Essential features that can be predicted include time to convergence, estimated swarm fractions in certain system states over time, and sensitivity of the swarm system to oscillations.

As discussed above, the rate equations assume well-mixed systems, that is, spatial information is considered to be irrelevant. In a rather rough abstraction step, space is therefore not modelled. In a second example of how swarm systems can be modelled, we investigate a model that represents space. In particular, we focus on approaches from the field of graph theory. A simple definition of a graph is G=(V,E) for a set V of vertices vi and a set E of edges ej which, in turn, are defined by pairs e=(v0,v1) of vertices $v0, v1 \in V$. A special kind of graphs are random graphs. They are generated in a stochastic way. For example, for a given set of vertices V we can go through the list of all potential edges $e\in E_pot$. We can define a probability p for an edge to be included in a given graph. When going through the list of potential edges $E_pot$, we make use of a stochastic process and check against probability p to decide whether they are included. The class of graphs that we have just defined is called Erdos-Renyi graphs. They are the most popular variant of random graphs. The connection to swarm systems is straightforward because we model robots by vertices and the neighbourhood relation (i.e., two robots perceive each other) by edges.

However, concerning our application as model for swarm robotic systems, they have an unfortunate feature. There are no constraints on edges besides probability p. Problematic is, for example, the following. Say $e0=(a,b)\in E$ and $e1=(b,c)\in E$, then one would expect that it is quite likely that we also have $e2=(a,c)\in E$. However, in an Erdos-Renyi graph the existence of e2 is statistically independent from e0 and e1. A better model, hence, is that of geometric random graphs. They are generated in the following way. Say, we have a two-dimensional plane – the unit square. Vertices are now defined as points on the unit square, that is, they have x- and y-coordinates. Edges are now defined by distances between vertices on the plane. For example, we can use the Euclidean distance d(v0,v1) of vertices v0, v1 and a threshold r. That threshold could be a sensor range. If $d(v0,v1)<r$ then $e=(v0,v1)\in E$. Geometric random graphs can be used in two ways as models for swarm robotics. If we make use of the information about vertex positions on the unit square, then we are close to a full spatial description of a swarm robot system. If we ignore that information and instead only make use of the edge information E, then we have abstracted away individual agent positions while still keeping an abstracted spatial representation (a model that only speaks about neighbourhood relationships). However, once we want to have a dynamic model, that is, a model that also describes how, for example, the edge property develops in time, then the mere edge information will make it difficult to formulate the dynamics.

While the model that makes use of point positions is microscopic, the model that relies only on neighbourhood relations has already macroscopic aspects. For example, one can speak of macroscopic concepts, such as connected components. In principle that is also possible in the purely microscopic model but it is incomprehensible because of too much information and it does not natively provide the explanatory concept of graphs.

### COLLECTIVE DECISION-MAKING

A fundamental capability of swarms is collective decision-making (see section 5.2). A swarm is comprised of many robots, which are collaborating autonomous entities on the microscopic level. On the macro-level, the swarm as a whole has to establish autonomy as well in order to operate properly. An essential feature of autonomy is to make decisions, which happens collectively at the macroscopic level.

For instance, the agreement of a flock on flying in a specific direction $\alpha$ can be considered a choice from an infinite number of alternatives, $\alpha \in [0^\circ, 360^\circ)$. The swarm's motion can also be constrained to a predefined path, for instance along a circle, leaving only one degree of freedom, namely whether to move clockwise (CW) or counter-clockwise (CCW). Consider, for example, the desert locusts, Schistocerca gregaria. When they are in the growth stage of a wingless nymph, they may exhibit a certain kind of collective motion. It is referred to as “marching bands”, where individuals seemingly change their movement direction in response to their neighbours. The emergence of this motion pattern depends on the density of locust individuals. Even in small swarms, the locusts might all change their direction, despite the majority being aligned beforehand. It has been empirically observed that immature locusts, which are still in the process of metamorphosis towards adulthood, march, highly aligned, in one direction within circular arenas for two to three hours when they occur in low densities. At this point, they spontaneously change their direction of choice and within only a few minutes they have collectively switched, now marching in the opposite direction for several hours. Locust nymphs that occur at high densities do not change their direction. In fact, they have been shown to march in one and the same direction for eight hours straight.

This behaviour and many other examples of collective decision-making can be modelled with methods from opinion dynamics. The microscopic process, that is, the decision-making rule of an individual robot can, for example, be implemented by the local majority rule, the voter model, or the Hegselmann--Krause model.

We start with the majority rule. A robot requests the current opinion of its neighbours. For example, within a short range the robot is able to perceive the current direction (CW or CCW) of close-by robots. The neighbours' directions are considered their opinion. The robot counts how many are in favour of CW compared to how many are in favour of CCW (including its own current opinion), and then it switches to the majority (or stays with the majority if it was already in favour of the majority opinion). The macroscopic interpretation of this behaviour is that it creates a positive feedback effect. Once there is a majority on a global scale, it is probable to take over the full population. Noise and potential spatial correlations in the system, however, can complicate both modelling and predictions.

The voter model is even simpler. A robot picks a neighbour randomly and switches to the opinion of that neighbour. The macroscopic effect possibly looks to be not easily predictable because it seems a random process. However, the macroscopic effect is positive feedback again, and the voter model implements an effective decision-making process.

Majority rule and voter model allow for an intriguing comparison. It is known that collective-decision making comes with a trade-off between either fast decisions or accurate decisions [FRANKS 2003]. Both at the same time seem not possible. It turns out that the majority rule is faster while the voter model is more accurate. Hence, a designer of a swarm robotic system has to decide based on the requirements [Valentini 2015].

The Hegselmann—Krause model differs from the above examples because it allows for a continuum of options [HEGSELMANN 2002]. A robot is allowed to pick any real number $x\in[0,1]$ as its opinion. A similar situation arises in flocking when robots can choose any direction as discussed above. The decision rule is then more difficult to define. For example, a robot checks the opinions of its neighbours and then takes the arithmetic average of all opinions as its new opinion. However, that does not necessarily result in a consensus (i.e., all robots end up with the same or similar opinions). Robots could form several `clusters’ in the opinion space. The main contributions of studies using the Hegselmann —Krause model are theoretical analyses that give upper bounds, for example, for the time it will take to converge on a solution. For swarm robotics, these results do not always have an application but the Hegselmann—Krause model allows for a good intuitive understanding of the challenges of collective decision-making. One of the main challenges is the problem of how to reach a consensus, which is a macroscopic feature, based on local information and local actions only. So, we are back to the micro-macro problem, and playing with the Hegselmann—Krause model is a good start for the interested reader to gain a deeper understanding of the micro-macro challenges in swarm robotics.

### Example Implementations and Projects

The Swarm-bots project was one of the first bigger projects in the history of swarm robotics. It was a European project running from 2001 to 2005. The concept was to develop a small mobile robot with the capability of connecting physically to other swarm-bots. A ring around the robot chassis and a gripper, that it can hold on, implements this. The upper part of the swarm-bot and its locomotion system can be rotated, such that attached robots can agree on a direction of travel and move as aggregate. During the project, it was shown, that in this aggregated form the swarm-bots are able to cross gaps or steep slopes that could not be crossed by individual robots. Hence, it nicely showed the potential of swarm robotics, where the robot group considerably extends the capabilities of the individual robot. The project also showed the potential of applying methods from evolutionary computation to synthesise robot controllers automatically. Hence, they started the research on evolutionary swarm robotics.

*Figure 3: The I-SWARM robot.*

Another relevant European project was the I-SWARM project. It was active from 2004 to 2008. The vision of I-SWARM was to build the artificial ant. The extremely ambitious objective was to build about 1000 robots of approximate dimension of 3x3x3mm3. The robots stand on three legs that can vibrate for locomotion. Piezoelectric motors implement the vibrations, and the technology is a flexible printed circuit board. The robots have four infrared sensors for proximity sensing. On this size scale, there is no possibility to include a battery, hence the robots have very efficient solar cells to drive the robot. The robot controllers run on an ASIC (application-specific integrated circuits). Besides this grand hardware challenge the project also included a lot of research on software approaches of how to control such robots. Output from this side included, for example, the BEECLUST algorithm, which adaptively aggregates the I-SWARM depending on environmental features.

## FURTHER READING

Unfortunately, there are not many textbooks that cover swarm robotics. For a start it is important to understand the basics of the underlying concept of swarm intelligence. A good read for that is the book of Bonabeau et al. [BONABEAU 1999]. The book by Floreano and Mattiussi [FLOREANO 2008] on bio-inspired AI has a few sections dedicated to swarm robotics, however, also other parts of the book are of interest here, such as behavioral robotics. The book by Hamann [HAMANN 2010] gives a good introduction to the challenge of combining the two levels of microscopic and macroscopic control and modeling approaches.

Bonabeau, Eric, Marco Dorigo, and Guy Theraulaz. Swarm intelligence: from natural to artificial systems. No. 1. Oxford university press, 1999.

Floreano, Dario, and Mattiussi, Claudio. Bio-inspired artificial intelligence: theories, methods, and technologies. MIT press, 2008.

Hamann, Heiko. Space-Time Continuous Models of Swarm Robotic Systems: Supporting Global-to-Local Programming. Springer Science & Business Media, 2010.

## REFERENCES

[DORIGO 20014] Dorigo M, Sahin E (2004) Guest editorial: Swarm robotics. Autonomous Robots 17(2-3):111–113

[BRAMBILLA 2013] Brambilla M, Ferrante E, Birattari M, Dorigo M (2013) Swarm robotics: a review from the swarm engineering perspective. Swarm Intelligence 7(1):1–41, DOI 10.1007/s11721-012-0075

[WERFEL 2014] Werfel J, Petersen K, Nagpal R (2014) Designing collective behavior in a termite-inspired robot construction team. Science 343(6172):754–758, DOI 10.1126/science.1245842, URL http://www.sciencemag.org/content/343/6172/754.abstract, http://www.sciencemag.org/content/343/6172/754.full.pdf

[CAPRARI 2005] Caprari G, Colot A, Siegwart R, Halloy J, Deneubourg JL (2005) Animal and robot mixed societies: building cooperation between microrobots and cockroaches. IEEE Robotics & Automation Magazine DOI 10.1109/MRA.2005.1458325

[GERLING 2016] Victor Gerling, Sebastian von Mammen. Robotics for Self-Organised Construction. In Proceedings of the 2016 IEEE 1st International Workshops on Foundations and Applications of Self* Systems, pages 156--165}, Augsburg, Germany, 2016, IEEE Computer Society.

[HAMANN 2015] Heiko Hamann, Mostafa Wahby, Thomas Schmickl, Payam Zahadat, Daniel Hofstadler, Kasper Stoy, Sebastian Risi, Andres Faiña, Frank Veenstra, Serge Kernbach, Igor Kuksin, Olga Kernbach, Phil Ayres, Przemyslaw Wojtaszek, **"flora robotica – Mixed Societies of Symbiotic Robot-Plant Bio-Hybrids",** Proc. of IEEE Symposium on Artificial Life IEEE ALIFE 2015, Cape Town, South Africa, pp. 1102-1109, IEEE, 2015

[SCHMICKL 2013] Thomas Schmickl, Stjepan Bogdan, Luís Correia, Serge Kernbach, Francesco Mondada, Michael Bodi, Alexey Gribovskiy, Sibylle Hahshold, Damjan Miklic, Martina Szopek, Ronald Thenius, José Halloy: *ASSISI: mixing animals with robots in a hybrid society*. Biomimetic and Biohybrid Systems Lecture Notes in Computer Science 8064 (2013), 441-443

[MARTINOLI 2004] Alcherio Martinoli, Kjerstin Easton, and William Agassounon. Modeling swarm robotic systems: A case study in collaborative distributed manipulation. Int. Journal of Robotics Research, 23(4):415–436, 2004.

[LERMAN 2005] Kristina Lerman, Alcherio Martinoli, and Aram Galstyan. A review of probabilistic macroscopic models for swarm robotic systems. In Erol Sahin and William M. Spears, editors, Swarm Robotics – SAB 2004 International Workshop, pages 143–152, Berlin, Germany, 2005. Springer.

[HEGSELMANN 2002] Hegselmann, Rainer, and Ulrich Krause. "Opinion dynamics and bounded confidence models, analysis, and simulation." *Journal of Artificial Societies and Social Simulation* 5.3 (2002).

[FRANKS 2003] Franks, Nigel R., et al. "Speed versus accuracy in collective decision making." *Proceedings of the Royal Society of London B: Biological Sciences* 270.1532 (2003): 2457-2463.

[Valentini 2015] Valentini, Gabriele, Heiko Hamann, and Marco Dorigo. "Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off." *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[VON MAMMEN 2016] von Mammen S, Tomforde S and Hähner J (2016) An Organic Computing Approach to Self-Organizing Robot Ensembles. *Front. Robot. AI*.