

Chapter 1

Fundamental Collective Behaviours in Swarm Robotics

Contents

1	Fundamental Collective Behaviours	1
1.1	Introduction	4
1.2	Getting together: Aggregation	6
1.2.1	Variants of the aggregation behaviour	6
1.2.2	Self-organised aggregation in biological systems	7
1.2.3	Self-organised aggregation in swarm robotics	8
1.2.4	Other studies	9
1.3	Acting together: Synchronisation	9
1.3.1	Variants of the synchronisation behaviour	10
1.3.2	Self-organised synchronisation in biology	10
1.3.3	Self-organised synchronisation in swarm robotics	11
1.3.4	Other studies	12
1.4	Staying together: Coordinated motion	13
1.4.1	Variants of the coordinated motion behaviour	13
1.4.2	Coordinated motion in biology	13
1.4.3	Coordinated motion in swarm robotics	15
1.4.4	Other studies	16
1.5	Searching together: Collective exploration	16
1.5.1	Variants of the collective exploration behaviour	16
1.5.2	Collective exploration in biology	17
1.5.3	Collective exploration in swarm robotics	17
1.6	Deciding together: Collective decision making	19
1.6.1	Variants of the collective decision making behaviour	20
1.6.2	Collective decision making in biology	21
1.6.3	Collective decision making in swarm robotics	22
1.6.4	Other studies	23
1.7	Conclusions	23

Summary

In this chapter, we present and discuss a number of fundamental collective behaviours studied within the swarm robotics domain. Swarm robotics is a particular approach to the design and study of multi-robot systems, which emphasises decentralised and self-organising behaviours that deal with limited individual abilities, local sensing and local communication. The desired features for a swarm robotics system are flexibility to variable environmental conditions, robustness to failure and scalability to large groups. These can be achieved thanks to well-designed collective behaviours—often obtained via some sort of bio-inspired approach—that rely on cooperation among redundant components. In this chapter, we discuss the solutions proposed for a limited number of problems common to many swarm robotics systems—namely aggregation, synchronisation, coordinated motion, collective exploration and decision-making. We believe that many real-world applications subsume one or more of these problems, and tailored solutions can be developed starting from the studies we review in this chapter. Finally, we propose possible directions for future research and discuss the relevant challenges to be addressed in order to push forward the study and the applications of swarm robotics systems.

1.1 Introduction

Imagine the following scenario: in a large area there are multiple items that must be reached, and possibly moved elsewhere or processed in some particular way. There is no map of the area to be searched, and the area is rather unknown, unstructured, and possibly dangerous for the intervention of humans or any valuable asset. The items must be reached and processed as quickly as possible, as a timely intervention would correspond to a higher overall performance. This is the typical scenario to be tackled with swarm robotics. It contains all the properties and complexity issues that make a swarm robotics solution particularly appropriate. Parallelism, scalability, robustness, flexibility and adaptability to unknown conditions are features that are required from a system confronted with such a scenario, and exactly those features are sought in swarm robotics research.

Put in other terms, swarm robotics promises the solution of complex problems through robotic systems made up of multiple cooperating robots. With respect to other approaches in which multiple robots are exploited at the same time, swarm robotics emphasises aspects like decentralisation of control, limited individual abilities, lack of global knowledge and scalability to large groups.

One important aspect that characterises a swarm robotics system concerns the robotic units, which are unable to solve the given problem individually. The limitation is given either by physical constraints that would prevent the single robot to individually tackle the problem (e.g., the robot has to move some items that are too heavy), or by time constraints that would make a solitary action very inefficient (e.g., there are too many items to be collected in a limited time).

Another source of limitation for the individual robot comes from its inability to acquire a global picture of the problem, having only access to partial (local) information about the environment and about the collective activity. These limitations imply the need for cooperation to ensure task achievement and better efficiency. Groups of autonomous cooperating robots can be exploited to synergistically achieve a complex task, by joining forces and sharing information, and to distributedly undertake the given task and achieve higher efficiency through parallelism.

The second important aspect in swarm robotics is redundancy in the system, which is intimately connected with robustness and scalability. Swarm robotics systems are made by homogeneous robots (or by relatively few heterogeneous groups of homogeneous robots). This means that the failure of a single or few robots is not a relevant fact for the system as a whole, because the failing robot can be easily replaced by another teammate. Differently from a centralised system, in a swarm robotics system there is no single point of failure, and every component is interchangeable with other components. Redundancy, distributed control and local interactions also allow for scalability, enabling the robotic system to seamlessly adapt to varying group sizes. This is a significant advantage with respect to centralised systems, which would present an exponential increase in complexity for larger group sizes.

All the above features being desiderata, the problem remains on how to design and implement such a robotic system. The common starting point in swarm robotics is the biological metaphor, for which the fundamental mechanisms that govern the organisation of animal societies can be distilled in simple rules to be implemented in the robotic swarm. This approach allowed to extract the basic working principles for many collective behaviours, and several examples will be presented in this chapter. However, it is worth noting that swarm robotics systems are not constrained to mimicking nature. Indeed, in many cases there is no biological example to be taken as reference, or the mechanisms observed in the natural system are too difficult to be implemented in the robotic swarm (e.g., odour perception is an open problem in robotics, preventing to easily exploit pheromone-based mechanisms by using real chemicals). Still, even in those systems that have no natural counterpart, the relevant property that should be present is self-organisation, for which group behaviour is the emergent result of the numerous interactions among different individuals. Thanks to self-organisation, simple control rules repeatedly executed by the individual robots may result in complex group behaviours.

If we consider the scenario presented at the beginning of this chapter, it is possible to recognise a number of problems common to many swarm robotics systems, which need to be addressed in order to develop suitable controllers. One first problem in swarm robotics is having robots get together in some place, especially when the robotic system is composed by potentially many individuals. Getting together (i.e., aggregation) is the precondition for many collective behaviours, and needs to be addressed according to the particular characteristics of the robotic system and of the environment in which it must take place. The aggregation problem is discussed in Section 1.2. Once groups are formed,

robots need some mechanism to stay together and to keep a coherent organisation while performing their task. A typical problem is therefore how to maintain such coherence, which corresponds to ensuring the synchronisation of the group activities (Section 1.3), and to keep the group in coordinated motion when the swarm must move across the environment (Section 1.4). Another common problem in swarm robotics corresponds to searching together and processing some items in the environment. To this aim, different strategies can be adopted to cover the available space, and to identify relevant navigation routes without resorting to maps and global knowledge (see Section 1.5). Finally, to maintain coherence and efficiency, the swarm robotics system is often confronted with the necessity to behave as a single whole. Therefore, it must be endowed with collective perception and collective decision mechanisms. Some examples are discussed in Section 1.6. For each of these problems, we describe some seminal work that produced solutions in a swarm robotics context. In each section, we describe the problem along with some possible variants, the biological inspiration and the theoretical background, the relevant studies in swarm robotics, and a number of other works that are relevant for some particular contribution given to the specific problem.

1.2 Getting together: Aggregation

Aggregation is a task of fundamental importance in many biological systems. It is the basic behaviour for the creation of functional groups of individuals, and therefore supports the emergence of various forms of cooperation. Indeed, it can be considered a prerequisite for the accomplishment of many collective tasks. In swarm robotics too, aggregation has been widely studied, both as a standing-alone problem or within a broader context. Speaking in general terms, aggregation is a collective behaviour that leads a group of agents to gather in some place. Therefore, from a (more or less) uniform distribution of agents in the available space, the system converges to a varied distribution, with the formation of well recognisable aggregates. In other words, during aggregation there is a transition from a homogeneous to a heterogeneous distribution of agents.

1.2.1 Variants of the aggregation behaviour

Aggregation can be achieved in many different ways. The main issue to be considered is whether or not the environment contains pre-existing heterogeneities that can be exploited for aggregation: light or humidity gradients (think of flies or sow bugs), corners, shelters and so forth represent heterogeneities that can be easily exploited. Their presence can therefore be at the basis of a collective aggregation behaviour, which however may not exploit interactions between different agents. Instead, whenever heterogeneities are not present (or cannot be exploited for the aggregation behaviour), the problem is more complex. The agents must behave in order to create the heterogeneities that support the for-




Figure 1.1: Aggregation process based on a diffusing signal that creates an intensity gradient. (a) Agents individually emit a signal, and move in the direction of higher gradient. (b) The individual signals sum up to form a stronger intensity gradient in correspondence of forming aggregates. (c) A positive feedback loop amplifies the aggregation process, until all agents are in the same cluster (d).

mation of aggregates. In this case, the basic mechanism of aggregation relies on a self-organising process based on a positive feedback mechanism. Agents are sources of some small heterogeneity in the environment (e.g., being the source of some signal that can be chemical, tactile or visual). The more aggregated agents, the higher the probability to be attracted by the signal. This mechanism leads to amplification of small heterogeneities, leading to the formation of large aggregates.

1.2.2 Self-organised aggregation in biological systems

Several biological systems present self-organised aggregation behaviours. One of the best studied examples is given by the cellular slime mold *Dictyostelium discoideum*, in which aggregation is enabled by self-generated biochemical signals that support the migration of cells and the formation of a multicellular body [10, 79]. A similar aggregation process can be observed in many other unicellular organisms [14]. Social and pre-social insects also present multiple forms of aggregation [28, 44]. In all these systems, it is possible to recognise two main variants of the aggregation process. On the one hand, the agents can emit a signal that creates an intensity gradient in the surrounding space. This gradient enables the aggregation process: agents react by moving in the direction of higher intensity, therefore aggregating with their neighbours (see Figure 1.1). On the other hand, aggregation may result from agents modulating their stopping time in response to social cues. Agents have a certain probability to stop and remain still for some time. The vicinity to other agents increases the probability of stopping and of remaining within the aggregate, eventually producing an aggregation process mediated by social influences (see Figure 1.2). In both cases, the same general principle is at work. Aggregation is dependent on two main probabilities: the probability to enter an aggregate, which increases with the aggregate size, and the probability to leave an aggregate, which decreases




Figure 1.2: Aggregation process based on variable probability of stopping within an aggregate. (a) Agents move randomly, and may stop for some time (grey agent). (b) When encountering a stopped agent, other agents stop as well, therefore increasing the size of the aggregate. (c) The probability to meet an aggregate increases with the aggregate size for geometric reasons. Social interactions modulate the probability of leaving the aggregate, which diminishes with increasing number of individuals. (d) Eventually, all agents are in the same aggregate.

accordingly. This creates a positive feedback loop that makes larger aggregates more and more attractive with respect to small ones. Some randomness in the system helps in breaking the symmetry and reaching a stable configuration.

1.2.3 Self-organised aggregation in swarm robotics

On the basis of the studies of aggregation in biological systems, various robotic implementations have been presented, based on either of the two behavioural models described above. Of particular interest is the work presented in [33], in which the robotic system was developed to accurately replicate the dynamics observed in the cockroach aggregation experiments presented in [44]. In this work, a group of Alice robots [18] was used and their controller was implemented by closely following the behavioural model derived from experiments with cockroaches. The behavioural model consists of four main conditions: (i) moving in the arena centre, (ii) moving in the arena periphery, (iii) stopping in the centre, and (iv) stopping in the periphery. When stopping, the mean waiting time is influenced by the number of perceived neighbours (for more details, see [33]). The group behaviour resulting from the interaction among Alice robots was analysed with the same tools used for cockroaches (see [44, 33]). The comparison of the robotic system with the biological model shows a very good correspondence, demonstrating that the mechanisms identified by the behavioural model are sufficient to support aggregation in a group of robots, with dynamics that are comparable to the one observed in the biological system. Additionally, the robotic model constitutes a constructive proof that the identified mechanisms really work as suggested.

This study demonstrates the approach of distilling in terms of simple rules the relevant mechanisms that produce a given self-organising behaviour. A

different approach consists in exploiting artificial evolution to synthesise the controllers for the robotic swarm. This allows the user to simply define some performance metric for the group, and let the evolutionary algorithm find the controllers capable of producing the desired behaviour. This generic approach has been exploited to evolve various self-organising behaviours, including aggregation [30]. In this case, robots were rewarded to minimise their distance from the geometric centre of the group, and to keep moving. The analysis of the evolved behaviours revealed that in all cases robots are attracted by teammates and repelled by obstacles. When a small aggregate forms, robots keep on moving thanks to the delicate balance between attractive and repelling forces. This makes the aggregate continuously expand and shrink, slightly moving across the arena. This slow motion of the aggregate makes it possible to attract other robots or other formed aggregates in the vicinity, and result in a very good scalability of the aggregation behaviour with respect to the group size. This experiment revealed a possible alternative mechanism for aggregation, which is not dependent on the probability of joining or leaving an aggregate. In fact, robots here never quit an aggregate in which they are attracted. Rather, the aggregates themselves are dynamic structures capable of moving within the environment, and in doing so they can be attracted by neighbouring aggregates, until all robots belong to the same group.

1.2.4 Other studies

The seminal papers described above are representative of other studies, which either exploit a probabilistic approach [45, 68], or rely on artificial evolution [4]. Approaches grounded on mathematical models and control theory are also worth mentioning [2, 35]. Other variants of the aggregation behaviour can be considered. The aggregate may be characterised by an internal structure, that is, agents in the aggregate are distributed on a regular lattice or form a specific shape. In such cases, we talk about *pattern/shape formation* [69]. Another possibility is given by the admissibility of multiple aggregates. In the studies mentioned so far, multiple aggregates may form at the beginning of the aggregation process, but as time goes by smaller aggregates are disbanded in favour of larger ones, eventually leading to a single aggregate for the whole swarm. However, it could be desirable to obtain multiple aggregates forming functional groups of a specific size. In this case, it is necessary to devise mechanisms for controlling the group size [47].

1.3 Acting together: Synchronisation

Synchronisation is a common phenomenon observed both in the animate and inanimate world. In a synchronous system, the various components present a strong time coherence between the individual behaviours. In robotics, synchronisation can be exploited for the coordination of actions, both within a single- or a multi-robot domain. In the latter case, synchronisation may be particu-

larly useful to enhance the system efficiency and/or to reduce the interferences among robots.

1.3.1 Variants of the synchronisation behaviour

Synchronisation in a multi-agent system can be of mainly two forms: “loose” and “tight”. In the case of loose synchronisation, we observe a generic coordination in time of the activities brought forth by different agents. In this case, single individuals do not present a periodic behaviour, but as a group it is possible to observe bursts of synchronised activities. Often in this case there are external cues that influence synchrony, such as the daylight rhythm. On the other hand, it is possible to observe tight synchronisation when the individual actions are perfectly coherent. To ensure a tight synchronisation in a group, it is possible to rely on either a centralised or a distributed approach. In the former, one agent acts as a reference (e.g., a conductor for the orchestra or the music theme for a ballet) and drives the behaviour of the other system components. In the latter, a self-organising process is in place, and the system shows the ability to synchronise without an externally-imposed rhythm. It is worth noting that tight synchronisation does not necessitate individual periodic behaviour, neither in the centralised nor in the self-organised case. For instance, synchronisation has also been studied between coupled chaotic systems [58]. In the following, we focus on self-organised synchronisation of periodic behaviours, which is the most studied phenomenon as it is commonly observed in many different systems.

1.3.2 Self-organised synchronisation in biology

Although synchronisation has always been a well known phenomenon [71], its study did not arouse much interest until the late 1960s, when Arthur Winfree began investigating the mechanisms underlying biological rhythms [82]. He observed that many systems in biology present periodic oscillations, which can get entrained when there is some coupling between the oscillators. A mathematical description of this phenomenon was first introduced by Yoshiki Kuramoto [46], who developed a very influential model that was afterwards refined and applied to various domains [71].

Similar mechanisms are at the base of the synchronous signalling behaviour observed in various animal species [14]. “Chorusing” is a term commonly used to refer to the coordinated emission of acoustic communication signals by large groups of animals. To cite a few, chorusing has been observed in frogs, crickets and spiders. However, the probably-most-fascinating synchronous display is the synchronous flashing of fireflies from South-East Asia. This phenomenon has been thoroughly studied until a self-organising explanation was proposed to account for the emergence of synchrony [11].

A rather simple model describes the behaviour of fireflies as the interactions between pulse-coupled oscillators [49]. In Figure 1.3, the activity of two oscillators is represented as a function of time. Each oscillator is of the integrate-and-fire type, which well represents a biological oscillator such as the one of




Figure 1.3: The synchronisation between pulse-coupled oscillators. The oscillator emit pulses each time its state variable reaches the threshold level (corresponding to 1 in the plot). When one oscillator emits a pulse, its state is reset while the state of the other oscillator is advanced by a constant amount, which corresponds to a phase shift, or to the oscillator firing if it overcomes the threshold.

fireflies. The oscillator is characterised by a voltage-like variable that is integrated over time until a threshold is reached. At this point, a pulse is fired and the variable is reset to the base level (see Figure 1.3). Interactions between oscillators take the form of constant phase shifts induced by incoming pulses, which bring other oscillators close to the firing state, or make them directly fire. These simple interactions are sufficient for synchronisation: in a group of similarly pulse-coupled oscillators, constant adjustments of the phase made by all the individuals lead to a global synchronisation of pulses (for a detailed description of this model, see [49]).

1.3.3 Self-organised synchronisation in swarm robotics

The main purpose of synchronisation in swarm robotics is the coordination of the activities in a group. This can be achieved in different ways, and mechanisms inspired by the behaviour of pulse-coupled oscillators have been developed. In [47], synchronisation is exploited to regulate the size of travelling robotic aggregates. Robots can emit a short sound signal (a chirp), and enter a refractory state for a short time after signalling. Then, robots enter an active state in which they may signal at any time, on the basis of a constant probability per time-step. Therefore, the chirping period is not constant and depends on the chirping probability. In this state, robots also listen to external signals and react by immediately emitting a chirp. This mechanism, similar to chorusing in frogs and crickets, leads to synchronised emission of signals. Thanks to this

simple synchronisation mechanism, the size of an aggregate can be somehow estimated: given the probabilistic nature of chirping, a robot has a probability of independently initiating signalling that depends on the number of individuals in the group; estimating this probability by listening to own and others' chirps allows an approximate group size estimation. Synchronisation therefore ensures a mechanism to keep coherence in the group, which is the precondition for group size estimation.

In [83], synchronisation is instead necessary to reduce the interferences between robots, which periodically perform foraging and homing movements in a cluttered environment. Without coordination, the physical interferences between robots going toward and away from the home location lead to a reduced overall performance. Therefore, a synchronisation mechanism based on the firefly behaviour was devised. Robots emit a signal in correspondence to the switch from foraging to homing. This signal can be perceived by neighbouring robots within a limited radius, and induces a reset of the internal rhythm that corresponds to a behavioural shift to homing. Despite the limited range of communication among robots, a global synchronisation is quickly achieved, which leads the group to reduce interferences and increase the system performance [83].

A different approach to the study of synchronisation is described in [76]. Here, artificial evolution is exploited to synthesise the behaviour of a group of robots, with the objective of obtaining minimal communication strategies for synchronisation. Robots were rewarded to present an individual periodic movement and to signal in order to synchronise the individual oscillations. The results obtained through artificial evolution are then analysed to understand the mechanisms that can support synchronisation, showing that two types of strategies were evolved: one is based on a modulation of the oscillation frequency, the other relies on a phase reset. These two strategies are also observed in biological oscillators: for instance, different species of fireflies present different synchronisation mechanisms, based on delayed or advanced phase responses [11].

1.3.4 Other studies

While self-organised synchronisation is a well known phenomenon, its application in collective and swarm robotics has not been largely exploited. The coupled-oscillator synchronisation mechanism was applied to a cleaning task to be performed by a swarm of micro robots [40]. Another interesting implementation of the basic model can be found in [20]. Here, synchronisation is exploited to detect and correct faults in a swarm robotics system. It is assumed that robots can synchronise a periodic flashing behaviour while moving in the arena and accomplishing their task. If a robot incurs some fault, it will be forced to stop synchronising. This fault can be detected and recovered by neighbouring robots. Similar to the heartbeat in distributed computing, correct synchronisation corresponds to a well-functioning system, while the lack of synchronisation corresponds to a faulty condition.

Finally, synchronisation behaviours may spontaneously emerge in an evolutionary robotics setup, even if they are not explicitly rewarded. In [84], synchro-

nisation of group activities evolved spontaneously as a result of the need to limit the interferences among robots in a foraging task. In [70], robots were rewarded to maximise the mean mutual information between their motor actions. Mutual information is a statistical measure derived in Information Theory, and roughly corresponds to the correlation between the output of two stochastic processes. Evolution therefore produced synchronous movements among the robots, which could actually maximise the mutual information while maintaining a varied behaviour.

1.4 Staying together: Coordinated motion

Another fundamental problem for a swarm is ensuring coherence in space. This means that the individuals in the swarm must display coordinated movement in order to maintain a consistent spatial structure. Coordinated motion is often observed in groups of animals. Flocks of birds or schools of fish are fascinating examples of self-organised behaviours producing a collective motion of the group. Similar problems need to be tackled in robotics, for instance for moving in formation or for distributedly deciding a common direction of motion.

1.4.1 Variants of the coordinated motion behaviour

The coordinated motion of a group of agents can be achieved in different ways. Also in this case, we can distinguish mainly between a centralised and a distributed approach. In a centralised approach, one agent can be considered the leader and the other agents follow (e.g., the mother duck with her ducklings). In the distributed approach, instead, there is no single leader and some coordination mechanism must be found to let the group move in a common direction. Of particular interest for swarm robotics are the coordinated motion models based on self-organisation. Such models consider multi-agent systems that are normally homogeneous and characterised by a uniform distribution of information: no agent is more informed than the others, and there exists no a priori preference for any direction of motion (i.e., agents start being uniformly distributed in space). However, through self-organisation and amplification of shared information, the system can break the symmetry and converge to a common direction of motion. A possible variant of the self-organised coordinated motion consists in having a non-uniform distribution of information, which corresponds to having some agents that are more informed than the others on a preferred direction of motion. In this case, a few informed agents may influence the motion of the entire group.

1.4.2 Coordinated motion in biology

Many animal species present coordinated motion behaviours, ranging from bacteria to fish and birds. Not all animal species employ the same mechanisms, but

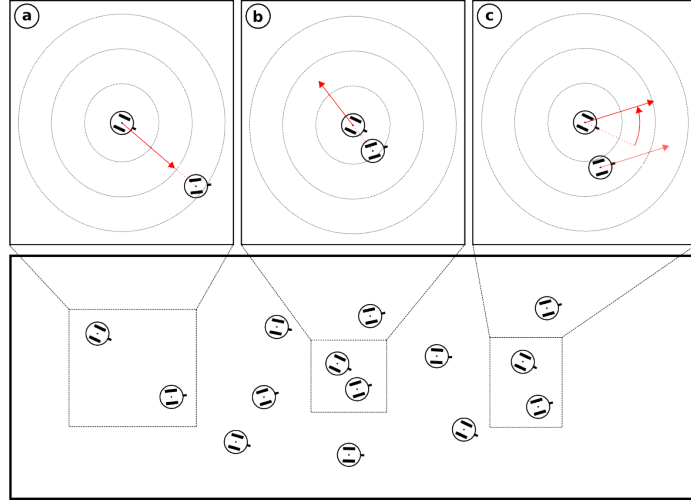


Figure 1.4: Self-organised coordinated motion in a group of agents. In the bottom part, a group of agents is moving in roughly the same direction. According to the model presented in [43], agents react to the closest neighbour within their perception range, and follow three main rules: (a) agents move toward a neighbour when it is too far; (b) agents move away from a neighbour when it is too close; (c) agents rotate and align with a neighbour situated at intermediate distances. The iterated application of these rules leads the group to move in a same direction.

in general it is possible to recognise various types of interactions among individuals that have a bearing on the choice of the motion direction. Coordinated motion has mainly been studied in various species of fish, in birds and in insect swarms [3, 51]. The most influential model has been introduced by Huth and Wissel to describe the behaviour of various observed species of fish [43]. In this model, it is assumed that each fish is influenced solely by its nearest neighbour. Also, the movement of each fish is based on the same behavioural model, which includes also some inherent random fluctuation. According to the proposed behavioural model, each fish follows essentially three rules: (i) approach a faraway individual, (ii) get away from individuals that are too close and (iii) align with the neighbour direction (see Figure 1.4). When the nearest neighbour is within the closest region, the fish reacts by moving away. When the nearest neighbour is in the farthest region, the fish reacts by approaching. Otherwise, if the neighbour is within the intermediate region, the fish reacts by aligning. These simple rules are sufficient to produce collective group motion, and the final direction emerges from the interactions among the individuals.

Starting from the above model, a number of variants have been proposed, which take into account different parameters and different numbers of individuals. In [25], a model including all individuals in the perceptual range was

introduced, and a broad analysis of the parameters was performed, showing how minor differences at the individual level correspond to large differences at the group level. In [6], an experimental study on bird flocks in the field was performed, obtaining position and velocity data for each bird in a real flock through stereo-photography and 3D mapping. The obtained data were used to verify the assumption about the number of individuals that each bird monitors during flocking, showing that this number is constant (and corresponds to about 7 individuals) notwithstanding the varying density of the flock. Finally, in [24], a model is developed in which some of the group members have individual knowledge on a preferential direction. The model describes the outcome of a consensus decision in the flock as a result of the interaction between informed and uninformed individuals.

1.4.3 Coordinated motion in swarm robotics

The models introduced for characterising the self-organised behaviour of fish schools or bird flocks have also inspired a number of interesting studies. The most influential work is definitely the one of Craig Reynolds, who developed virtual creatures called *boids* [62]. In this work, each creature executes three simple behaviours: (i) collision avoidance, to avoid crashing with nearby flockmates; (ii) velocity matching, to move in the same way of nearby flockmates; and (iii) flock centering, to stay close to nearby flockmates. Notice that the behavioural model corresponds to the models proposed in biological studies. The merit of this work is being the first implementation of the rules studied for real flocks in a virtual 3D world, showing a close correspondence of the behaviour of boids with those of flocks, herds and schools. The research of Reynolds has been taken as inspiration by many other studies on coordinated motion, mainly in simulation. In [78], an implementation of the flocking behavioural model has been proposed and tested on real robots. Robots use infrared proximity sensors to recognise the presence of other robots and their distance, necessary to perform collision avoidance and flock centering behaviours. Additionally, a dedicated sensor to perceive the heading of neighbours was developed to support the aligning behaviour. This system, called Virtual Heading System (VHS), is based on a digital compass and wireless communication. Despite the fact that digital compass cannot reliably work in an indoor environment, it is assumed that neighbouring robots have similar perceptions. The perceived heading with respect to the local north is communicated over the wireless channel, and it is exploited for the alignment behaviour. This system allowed to test the flocking behaviour of small robotic groups in a physical setting, and to study the dynamics of flocking with up to 1000 simulated robots. This work was later extended in [13], by having a subgroup of informed individuals which could steer the whole flock, following the model presented in [24]. The dynamics of steered flocking have been studied by varying the percentage of informed robots in simulation, and tests with real robots have been performed as well.

1.4.4 Other studies

As mentioned above, there exist numerous studies inspired by the schooling/flocking models. All these studies adopt some variants of the behavioural rules described above, or analyse the group dynamics under some particular perspective. A different approach to coordinated motion can be found in [16]. In this work, robots have to transport a heavy object and have imperfect knowledge of the direction of motion. They can, however, negotiate the goal direction by displaying their own preferred direction of motion and by adjusting it on the basis of the direction displayed by others. On the whole, this mechanism implements similar dynamics to the alignment behaviour of the classical flocking model. Here, however, robots are connected together to the object to be transported, adding a further constraint to the system that obliges a good negotiation to allow motion. A similar constraint characterises the coordinated motion studies with physically assembled robots presented in [5, 75]. Here, robots form a physical structure of varying shape, and can rotate their chassis in order to match the direction of motion of the other robots. In this case, there is no direct detection of the motion direction of neighbours. Instead, robots can sense the pulling and pushing forces that are exerted by the other connected robots through the physical connections. These pulling/pushing forces are naturally averaged by the force sensor, which returns their resultant. Artificial evolution was exploited to synthesise an artificial neural network that could transform the sensed forces to motor commands. The obtained results show an impressive capability of self-organised coordination between the robots, as well as scalability and generalisation to different size and shapes [5], and the ability to cope with obstacles and to avoid falling outside the borders of the arena [75].

1.5 Searching together: Collective exploration

Exploring and searching the environment is an important behaviour for robot swarms. In many tasks, the swarm must interact with the environment, sometimes only to monitor it, but sometimes also to process materials or other kinds of resources. Usually, the swarm cannot completely perceive the environment, and the environment may also change during the operation of the robots. Hence, robots need to explore and search the environment to monitor for changes or in order to detect new resources.

To cope with its partial perception of the environment, a swarm can move, for instance using flocking, in order to explore new places (some locations may be unavailable though). Hence, most of the environment can be perceived, but not at the same time. As in many other artificial system, a tradeoff between exploration and exploitation exists, and requires careful design choices.

1.5.1 Variants of the collective exploration behaviour

There is no perfect exploration and search strategy because the structure of the environment in which the swarm is placed can take many different shapes.

Strategies only perform more or less well as a function of the situation with which they are faced [85]. For instance, the swarm could be in a maze, in a open environment with few obstacles, or in an environment with many obstacles.

We identified a restricted number of environmental characteristics that play an important role in the choice of the searching behaviour in swarm robotics. These characteristics are commonly found in swarm robotics scenarios, and are the presence of a central place, the size of the environment, the presence of obstacles.

The central place is a specific location where robots must come back regularly, for instance for maintenance or to deposit foraged items. A scenario that involves a central place requires a swarm able to either remember or keep track of that location.

If the environment is closed (finite area) and not too large, the swarm may use random motion to explore, with fair chances to rapidly locate resources (or even the central place). In an open environment, robots can get lost very quickly. In this type of environment, it is necessary to use behaviours that allow robots to stay together and maintain connectivity.

Obstacles are environmental elements that constrain the motion of the swarm. If the configuration of the obstacles is known in advance, the swarm can move in the environment following appropriate patterns. In most cases, however, obstacles are unexpected or might be dynamic and may prevent the swarm to explore parts of the environment.

1.5.2 Collective exploration in biology

In nature, animals are constantly looking for resources such as food, sexual partners, or nesting sites. Group living animals may use several behaviours to explore their environment and locate these resources.

For instance, fish can take advantage of the number of individuals in a shoal to improve their capabilities to find food [59, 60, 41]. To do so, they move and maintain large interdistances between individuals. In this way, fish increase their perceptual coverage as well as their chances to find new resources.

Animals also heavily rely on random motion to explore their environment [22, 77, 52]. Usually the exploratory pattern is not fully random (that is, isotropic), because animals use all possible environmental cues at hand to guide themselves. Random motion can be biased towards a given direction, or it can be constrained in a specific area, for instance around a previously memorised location [9]. Some desert ants achieve high localization performance with odometry (counting their footsteps) and relying on gravity and the polarization of natural light. They may move randomly to look for resources but they are able to quickly return to their nest, and also to return to an interesting location previously identified.

1.5.3 Collective exploration in swarm robotics

One of the most common exploration strategies used in robotics is the random exploration. In a typical implementation, robots wander in the environment un-




Figure 1.5: Gas expansion behaviour to monitor the surroundings of a central place. (a) The swarm starts aggregated around a central place (represented by a black spot). (b,c) Robots try to move as far as possible from their neighbours, while maintaining some visual or radio connection. (d) As a result, the whole swarm expands in the environment, like a gas, covering part of the environment.




Figure 1.6: Chaining behaviour in action with a central place represented by a large black dot, bottom left. (a) Robots start aggregated around the central place. (b) While maintaining visual or radio contact with neighbours, some robots change role and become part of a chain (greyed out). (c) Other robots move around the central place and encounter the early chain of robots. With some probability, they also turn into new parts of the chain. (d) At the end of the iterative process, robots are forming a long chain that spans through the environment and maintains a physical link to the central place.

til they perceive a feature of interest [27, 65, 37]. By doing this, robots possibly lose contact with each other and therefore their ability to work together, hence this strategy is not suited for large or open environments. Due to the stochastic nature of the strategy, its performance can only be evaluated statistically. On average, the time to locate a feature is proportional to the squared distance with robots [22].

Systematic exploration strategies are very different. Robots use some a priori knowledge about the structure of the environment in order to methodically sweep it and find features. To ensure that robots do not repeatedly cover the same places, they may need to memorize which places have already been explored. This is often implemented with localization techniques and mapping of the environment [12]. The advantage of this technique is that an answer will be found with certainty, and time of exploration has a lower and upper bound if the

environment is not open. However, the memory requirements may be excessive, and the strategy is not suited for open environments.

Between the two extreme strategies reported in the previous paragraph lie a number of more specialised strategies that present advantages and drawbacks depending on the structure of the environment and the distribution of the resources.

Collective motion (already detailed in Section 1.4) allows swarms to maintain their cohesion while moving through the environment. The flocking behaviour can be employed in an open environment with a limited risk of losing contact between robots. The swarm behaves like a sort of physical mesh that covers part of the environment; to maximize the area covered during exploration robots can increase as much as possible their interdistance during motion.

The gas expansion behaviour (see Figure 1.5) allows robots to quickly and exhaustively explore the surroundings of a central place [55, 56, 42, 7]. While one or several robots keep track of a central place, other robots try to move as far as possible from their neighbours, while still maintaining direct line of sight with at least one of them. The swarm behaves like a fluid or gas that penetrates the asperities of the environment. The exploration is very effective and any change or new resource within the perception range of the swarm is immediately perceived. However, since robots are bound to the central place, the area that they can explore is limited by the number of robots in the swarm. If robots do not stick to a central place, the resulting behaviour shifts to a type of flocking or moving formation.

With the chaining behaviour, swarms can form a chain with one end that sticks to a central place and the other end that freely moves through the environment (see Figure 1.6). In [81], a minimalistic behaviour produces a static chain, but different types of chain motions can be imagined. In [50] for instance, chains can build up, move, and disaggregate until a resource is found. Contrary to the gas expansion behaviour, a chaining swarm may not immediately perceive changes in the environment because it has to constantly sweep the space. Chaining allows robots to cover a more important area than the gas expansion behaviour, ideally a disc of radius proportional to the number of robots.

1.6 Deciding together: Collective decision making

Decision making is a behaviour used by any artificial system that must produce an adapted response when facing new or unexpected situations. Because the best action depends on the situation encountered, a swarm cannot rely on a pre-programmed and systematic reaction. Monolithic artificial systems make decisions all the time, by gathering information and then evaluating the different options at hand. However, when it comes to swarms, each group member might have its own opinion about the correct decision. If all individuals perceive the same information, and process it in the same way, then they might

independently make the same decision. But in practice, the more common case is that individuals perceive partial and noisy information about the situation. Thus, if no coordination among group members occurs, a segregation based on differing opinions might take place, thereby removing the advantages of being a swarm. Therefore, the challenge is to have the whole group collaborate to make a collective decision, and take action accordingly.

1.6.1 Variants of the collective decision making behaviour

There are mainly three mechanisms reported in the literature that allow swarms to make collective decisions. The first and most simple mechanism is based on opinion propagation. As soon as a group member has enough information about a situation to make up its mind, it propagates its opinion through the whole group.

The second mechanism is based on opinion averaging. All individuals constantly share their opinion with their neighbors and also adjust their own opinion in consequence. This iterative process leads to the emergence of a collective decision. The adjustment of the opinion is typically achieved with an average function, especially if opinions are about quantitative values such as a location, a distance, or a weight for instance.

The third and last mechanism relies on amplification to produce a collective decision. In a nutshell, all individuals start with an opinion, and may decide to change their opinion for another one. The switch to a new opinion happens with a probability calculated on the basis of the frequency of this opinion in the swarm. Practically, this means that if an opinion is more represented in the group, it has also more chances to be adopted by an individual, which is why the term amplification is used.

Each of the three aforementioned mechanisms has some advantages over the others and may be preferred depending on the situation faced. The factors that play an important role in collective decision processes include the speed needed to make the choice, the robustness of communication, and the reliability of individual information.

In terms of speed, opinion propagation allows fast collective decisions, in contrast with the the two other mechanisms which require numerous interactions among individuals. However, this speed generally comes at the cost of robustness or accuracy [19]. If communication is not robust enough, messages can be corrupted. The mechanism of opinion propagation is particularly sensitive to such effects, and a wrong or random collective decision might be made by a swarm in that case.

The averaging mechanism would produce a more robust decision because wrong information from erroneous messages is diluted in the larger amount of information present in the swarm [38]. However, opinion averaging works best if all individuals have roughly identical knowledge. If a small proportion of individuals have excellent knowledge to make the decision, while the remaining individuals have poor information, opinion propagation may produce better results than opinion averaging [23].

Lastly, the amplification mechanism is the main choice for a gradually emerging collective decision if opinions cannot be merged with some averaging function. Instead of adjusting opinions, individuals simply adopt new opinions with some probability. It is worth noting that this mechanism can produce good decisions even if individuals have poor information.

1.6.2 Collective decision making in biology

The powerful possibilities of decision making in groups were already suggested by Francis Galton back in 1905 [31]. In this paper, Galton reports the results of a weight-judging competition in which competitors had to estimate the weight of a fat ox. With slightly less than 800 independent estimates, Galton observed that the average estimate was accurate to one percent of the real weight of the ox. This early observation opened interesting perspectives about the accuracy of collective estimations, but it was not describing a collective decision mechanism since Galton himself had to gather the estimates and apply some calculation to evaluate the estimate of the crowd.

More recent studies about group navigation have shown that groups of animals cohesively moving together towards a goal direction reach their objective faster than independent individuals [67, 21]. The mechanism of collective navigation not only allows the individuals to move and stay together, but it also acts as a distributed averaging function that locally fuses the opinions of individuals about the direction of motion, allowing them to improve their navigation performance.

In the last decades, the amplification mechanism has been identified as a source of collective decision in a broad range of animal species such as ants [26, 8], spiders [64], cockroaches [1], monkeys [57], and sheeps [48].

Ants that choose one route to a resource probably constitute the most well known example of the amplification mechanism. In [8], an ant colony is offered two paths to two identical resource sites. Initially the two resources are exploited equally, but after a short time ants focus on a single resource. This collective choice happens because ants that have found the resource come back to the nest, marking the ground with a pheromone trail. The next ants that try to reach the resource are sensitive to this odor and have higher chances to follow the path with higher pheromone concentration. As a result of this amplified response, a collective decision rapidly emerges. In addition, it was shown in [36] that when ants are presented two paths of different lengths to the same resource, the same pheromone-based mechanism allows them to choose the shortest path. This can be explained by the fact that ants using the shortest path need less time to make round trips, making the pheromone concentration on this path grow faster.

Quorum sensing is a special case of the amplification mechanism which has been notably used to explain nest site selection in ants, bees, and fish [73, 80, 61]. The most basic example of quorum sensing uses a threshold to dictate if individuals should change their opinion. If an individual perceives enough neighbours (above the threshold) that already share the opposite opinion, then it will in turn adopt this opinion. It has been shown that this threshold makes




Figure 1.7: A swarm of robots is presented with two resource sites in its environment and must collectively choose one. (a) Initially, robots are randomly scattered. (b) Using a random walk, they move until a resource site is found. On average, the swarm is split in the two sites. (c) The more neighbours they perceive, the longer the robots stay. A competition between the two sites takes place and any random event may change the situation. Here a robot just left the top right site, further reducing the chances that other robots stay there. (d) The swarm has made a choice in favour of the bottom left site. The choice is stable, although some robots may frequently leave the site for exploration.

quorum sensing more robust to the propagation of erroneous information during the decision process. In addition, the accuracy of collective decisions made with quorum sensing may improve with group size, and cognitive capabilities of groups may outperform the ones of single individuals [74, 17].

1.6.3 Collective decision making in swarm robotics

In swarm robotics, opinion averaging has been used to improve the localization capabilities of robots. In [37], a swarm of robots carries out a foraging task between a central place and a resource site. The robots have to navigate back and forth between the two places, and use odometry to estimate their location. As odometry provides noisy estimates, robots using solely this technique may quickly get lost. Here, robots can share and merge their localization opinions when they meet, by means of local infrared communication. By doing so, robots manage better localization and improve their performance in the foraging task. Moreover, robots associate a confidence level to their estimates, which is used to decide how information is merged. If a robot advertises an opinion with a very high confidence, then the mechanism produces opinion propagation. Hence the two mechanisms of averaging and propagation are blended in a single behaviour, and the balance between them is tuned by the user with a control parameter.

The aggregation behaviour previously mentioned in Section 1.2.3 can be exploited to trigger collective decision making in situations where there are several environmental heterogeneities. In [32], the robots are presented two shelters and they choose one of them as a resting site by aggregating there. The behaviour of the robots closely follows the one observed in cockroaches (see Figure 1.7). In [39], both robots and cockroaches are introduced in an

arena with two shelters, demonstrating the influence of the two groups on each other when making the collective decision. The collective decision is the result of an amplification mechanism, implemented via the probability for a robot to leave an aggregate. This probability diminishes with the number of perceived neighbours, allowing larger aggregates to attract more robots.

1.6.4 Other studies

The opinion averaging mechanism has been deeply investigated with a general mathematical approach in [54, 53]. These studies demonstrate convergence of the mechanism, and emphasize the importance of the topology of the communication network through which interactions take place.

Another amplification mechanism inspired from the behaviour of honeybees has been implemented in [66]. With this mechanism, it was shown that robots were able to make a collective decision and reliably choose between two sites the one offering the best illumination conditions.

The amplification mechanism based on pheromone trails and used by ants has also inspired several swarm robotics studies. In [72, 34], the pheromone is replaced by light projected by a beamer. This implementation is limited to laboratory studies, but it allowed demonstrating path selection with robot swarms. In [15], the process is abstracted inside a network of robots that are deployed in the environment. Virtual ants hop from robots to robots and deposit pheromone inside them. Eventually, the shortest path to a resource is marked out by robots with high and sustained levels of pheromone.

1.7 Conclusions

In this chapter, we have presented a broad overview of the common problems to be faced in a swarm robotics context, and we pointed to possible approaches to obtain solutions based on a self-organising process. We have discussed aggregation, synchronisation, coordinated motion, collective exploration and decision-making, and we argued that many application scenarios could be solved by a mix of the above solutions. So, are we done with swarm robotics research? Definitely not.

First of all, the fact that possible solutions exist does not mean that they are the most suitable for any possible application scenario. Hardware constraints, miniaturisation, environmental contingencies and performance issues may require the design of different solutions, which may strongly depart from the examples given above. Still, the approaches we presented constitute a logical starting point, as well as a valid benchmark against which novel approaches can be compared.

Another important research direction consists in characterising the self-organising behaviours we presented in terms of abstract properties, such as time of convergence toward a stable state, sensitivity to parameter changes, robustness to failures and so forth. From this perspective, the main problem is

ensuring a certain functionality of the system with respect to the needs of the application, and predicting the system features before actual development and testing. In many cases, a precise characterisation of the system is not possible, and only a statistical description can be achieved. Still, such an enterprise would bring swarm robotics closer to an engineering practice, eventually allowing to guarantee a certain performance of the developed system, as well as other properties that engineering commonly deals with.

The examples we presented all refer to homogeneous systems, in which all individuals are physically identical and follow exactly the same rules. This is however a strong simplification, which follows the tradition of biological modelling of self-organising behaviours. However, instead of being a limitation, heterogeneity is potentially a richness to be exploited in a swarm robotics system, which can lead to more complex group behaviours. For instance, different individual reactions to features in the environment can be at the basis of optimal decision making at the group level [63]. Otherwise, heterogeneity between groups of individuals can be exploited for performing tasks that require specialised abilities, but maintaining an overall redundancy of the system that ensures robustness and scalability [29].

In conclusion, swarm robotics research has still many challenges to address, that range from the need for more theoretical understanding of the relation between individual behaviour and group dynamics, to the autonomy and adaptation to varied real world conditions in order to face complex application scenarios (e.g., due to harsh environmental conditions such as planetary or underwater exploration, or to strong miniaturisation down to the micro scale). Whatever is the theoretical or practical driver, we believe that the studies we have presented in this chapter constitute fundamental reference points that teach us how self-organisation can be obtained in a swarm robotics system.

Acknowledgments

Vito Trianni acknowledges support from the H²Swarm project, founded within the EUROCORES Programme “EuroBioSAS” of the European Science Foundation. The project is partially supported by funds from the Italian CNR and the Belgian F.R.S.-FNRS. Alexandre Campo acknowledges support from the CoCoRo project, funded by the Information and Communication Technologies programme (call FP7-ICT-2009-6) of the European Commission under grant number 270382.

Bibliography

- [1] J.M. Amé, J. Halloy, C. Rivault, C. Detrain, and J.-L. Deneubourg. Collegial decision making based on social amplification leads to optimal group formation. *PNAS*, 103(15):5835–5840, 2006.
- [2] Hideki Ando, Yoshinobu Oasa, Ichiro Suzuki, and Masafumi Yamashita. Distributed memoryless point convergence algorithm for mobile robots with limited visibility. *IEEE Transactions on Robotics and Automation*, 15(5):818–828, 1999.
- [3] I. Aoki. A simulation study on the schooling mechanism in fish. *Bulletin of the Japanese Society of Scientific Fisheries*, v. 48(8) p. 1081-1088, 1982.
- [4] E. Bahçeci and E. Şahin. Evolving aggregation behaviors for swarm robotic systems: a systematic case study. In *Swarm Intelligence Symposium, 2005. SIS 2005. Proceedings 2005 IEEE*, pages 333–340. IEEE Press, Piscataway, NJ, 2005.
- [5] G. Baldassarre, V. Trianni, M. Bonani, F. Mondada, M. Dorigo, and S. Nolfi. Self-organised coordinated motion in groups of physically connected robots. *IEEE Transactions on Systems, Man and Cybernetics - Part B: Cybernetics*, 37(1):224–239, 2007.
- [6] M. Ballerini, N. Calbibbo, R. Candeir, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4):1232–1237, 2008.
- [7] M. Batalin and G. Sukhatme. Spreading out: A local approach to multi-robot coverage. In Hajime Asama, Tamio Arai, Toshio Fukuda, and Tsutomu Hasegawa, editors, *Distributed Autonomous Robotic Systems 5*, pages 373–382. Springer Japan, 2002.
- [8] R. Beckers, J.-L. Deneubourg, S. Goss, and J.M. Pasteels. Collective decision making through food recruitment. *Insectes sociaux*, 37(3):258–267, 1990.

- [9] Simon Benhamou. Spatial memory and searching efficiency. *Animal Behaviour*, 47(6):1423 – 1433, 1994.
- [10] J. T. Bonner. Chemical signals of social amoebae. *Scientific American*, 248:114–120, 1983.
- [11] J. Buck. Synchronous rhythmic flashing of fireflies. II. *The Quarterly Review of Biology*, 63(3):256–289, 1988.
- [12] Wolfram Burgard, Mark Moors, Dieter Fox, Reid Simmons, and Sebastian Thrun. Collaborative multi-robot exploration. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2000 (ICRA '00)*, volume 1, pages 476–481. IEEE, 2000.
- [13] H. Çelikkanat and E. Şahin. Steering self-organized robot flocks through externally guided individuals. *Neural computing & applications*, 19(6):849–865, 2010.
- [14] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- [15] A. Campo, Á. Gutiérrez, S. Nouyan, C. Pinciroli, V. Longchamp, S. Garnier, and M. Dorigo. Artificial pheromone for path selection by a foraging swarm of robots. *Biological cybernetics*, 103(5):339–352, 2010.
- [16] A. Campo, S. Nouyan, M. Birattari, R. Groß, and M. Dorigo. Negotiation of goal direction for cooperative transport. In Marco Dorigo, Luca M. Gambardella, Mauro Birattari, A. Martinoli, R. Poli, and Thomas Stützle, editors, *Ant Colony Optimization and Swarm Intelligence: 5th International Workshop, ANTS 2006*, volume 4150 of *Lecture Notes in Computer Science*, pages 191–202. Springer Verlag, Berlin, Germany, 2006.
- [17] Stéphane Canonge, J.-L. Deneubourg, and Grégory Sempo. Group living enhances individual resources discrimination: The use of public information by cockroaches to assess shelter quality. *PLoS ONE*, 6(6):e19748, 06 2011.
- [18] G. Caprari and R. Siegwart. Mobile micro-robots ready to use: Alice. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005)*, pages 3295–3300. IEEE Press, Piscataway, NJ, 2005.
- [19] Lars Chittka, Peter Skorupski, and Nigel E Raine. Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7):400–407, 2009.
- [20] A. L. Christensen, R. O’Grady, and M. Dorigo. From fireflies to fault-tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766, 2009.

- [21] E.A. Codling, J.W. Pitchford, and S.D. Simpson. Group navigation and the many-wrongs principle in models of animal movement. *Ecology*, 88(7):1864–1870, 2007.
- [22] 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.
- [23] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon A Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [24] I.D. Couzin, J. Krause, N.R. Franks, and S.A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [25] I.D. Couzin, J Krause, R James, G.D. Ruxton, and N.R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1–11, 2002.
- [26] J.-L. Deneubourg and S. Goss. Collective patterns and decision-making. *Ethol. Ecol. Evol*, 1:295–311, 1989.
- [27] J.-L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chrétien. The dynamics of collective sorting robot-like ants and ant-like robots. In *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*, pages 356–363, 1991.
- [28] J.-L. Deneubourg, J. C. Grégoire, and E. Le Fort. Kinetics of larval gregarious behavior in the bark beetle *dendroctonus micans* (coleoptera, scolytidae). *Journal of Insect Behavior*, 3(2):169–182, 1990.
- [29] M. Dorigo, D. Floreano, L.M. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A.L. Christensen, A. Decugnire, G.A. Di Caro, F. Ducatelle, E. Ferrante, A. Fröster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M.A. Montes de Oca, R. O’Grady, C. Pinciroli, G. Pini, P. Rétornaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A.E. Turgut, and F. Vaussard. Swarmanoid: a novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics and Automation Magazine*, 2012. Accepted for publication.
- [30] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2–3):223–245, 2004.
- [31] F. Galton. Vox populi. *Nature*, 75:450–451, 1907.

- [32] S. Garnier, J. Gautrais, M. Asadpour, C. Jost, and G. Theraulaz. Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2):109–133, 2009.
- [33] S. Garnier, C. Jost, J. Gautrais, M. Asadpour, G. Caprari, R. Jeanson, A. Grimal, and G. Theraulaz. The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial Life*, 14(4):387–408, 2008.
- [34] S. Garnier, F. Tâche, M. Combe, A. Grimal, and G. Theraulaz. Alice in pheromone land: An experimental setup for the study of ant-like robots. In *Swarm Intelligence Symposium, 2007. SIS 2007. IEEE*, pages 37–44. IEEE, 2007.
- [35] V. Gazi. Swarm aggregations using artificial potentials and sliding-mode control. *IEEE Transactions on Robotics*, 21(6):1208–1214, 2005.
- [36] S. Goss, S. Aron, J.-L. Deneubourg, and J.M. Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76(12):579–581, 1989.
- [37] Á. Gutiérrez, A. Campo, F. Santos, F. Monasterio-Huelin Maciá, and M. Dorigo. Social odometry: imitation based odometry in collective robotics. *International Journal of Advanced Robotic Systems*, 6(2):129–136, 2009.
- [38] Alvaro Gutiérrez, Alexandre Campo, Francisco C. Santos, Félix Monasterio-Huelin, and Marco Dorigo. Social odometry: Imitation based odometry in collective robotics. *International Journal of Advanced Robotic Systems*, 6(2):1–8, 2009.
- [39] José Halloy, Grégory Sempo, Gilles Caprari, Colette Rivault, M Asadpour, F Tache, I Said, V Durier, S Canonge, JM Ame, et al. Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853):1155, 2007.
- [40] M. Hartbauer and H. Roemer. A novel distributed swarm control strategy based on coupled signal oscillators. *Bioinspiration & Biomimetics*, 2(3):42–56, 2007.
- [41] D.J. Hoare, I.D. Couzin, J.G.J. Godin, and J. Krause. Context-dependent group size choice in fish. *Animal Behaviour*, 67(1):155–164, 2004.
- [42] A. Howard, M.J. Matarić, and G.S. Sukhatme. An incremental self-deployment algorithm for mobile sensor networks. *Autonomous Robots*, 13(2):113–126, 2002.
- [43] A. Huth and C. Wissel. The simulation of the movement of fish schools. *Journal of Theoretical Biology*, 156(3):365–385, 1992.
- [44] R. Jeanson, C. Rivault, J.-L. Deneubourg, S. Blanco, R. Fournier, C. Jost, and G. Theraulaz. Self-organized aggregation in cockroaches. *Animal Behaviour*, 69(1):169–180, 2005.

- [45] S. Kernbach, R. Thenius, O. Kernbach, and T. Schmickl. Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. *Adaptive Behavior*, 17(3):237–259, 2009.
- [46] Y. Kuramoto. Phase dynamics of weakly unstable periodic structures. *Progress of Theoretical Physics*, 71(6):1182–1196, 1984.
- [47] C. Melhuish, O. Holland, and S. Hoddell. Convoying: using chorusing to form travelling groups of minimal agents. *Robotics and Autonomous Systems*, 28:207–216, 1999.
- [48] P. Michelená, R. Jeanson, J.-L. Deneubourg, and A.M. Sibbald. Personality and collective decision-making in foraging herbivores. *Proceedings of the Royal Society B: Biological Sciences*, 277(1684):1093, 2010.
- [49] R. E. Mirollo and S. H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [50] S. Nouyan, R. Groß, M. Bonani, F. Mondada, and M. Dorigo. Teamwork in self-organized robot colonies. *Evolutionary Computation, IEEE Transactions on*, 13(4):695–711, 2009.
- [51] A. Okubo. Dynamical aspects of animal grouping: Swarms, schools, flocks, and herds. *Advances in Biophysics*, 22:1–94, 1986.
- [52] A. Okubo and S.A. Levin. *Diffusion and ecological problems: modern perspectives*, volume 14. Springer Verlag, 2001.
- [53] R. Olfati-Saber, J.A. Fax, and R.M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [54] R. Olfati-Saber and R.M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *Automatic Control, IEEE Transactions on*, 49(9):1520–1533, 2004.
- [55] D. Payton, M. Daily, R. Estowski, M. Howard, and C. Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, 2001.
- [56] D. Payton, R. Estowski, and M. Howard. Progress in pheromone robotics. In M. Gini, W.-M. Shen, C. Torras, and H. Yuasa, editors, *Intelligent Autonomous Systems 7*, pages 256–264, 2002.
- [57] O. Petit, J. Gautrais, J.B. Leca, G. Theraulaz, and J.-L. Deneubourg. Collective decision-making in white-faced capuchin monkeys. *Proceedings of the Royal Society B: Biological Sciences*, 276(1672):3495, 2009.
- [58] A. Pikovsky, M. Rosenblum, and J. Kurths. Phase synchronization in regular and chaotic systems. *International Journal of Bifurcations and Chaos*, 10(10):2291–2305, 2000.

- [59] T.J. Pitcher, A.E. Magurran, and I.J. Winfield. Fish in larger shoals find food faster. *Behavioral Ecology and Sociobiology*, 10(2):149–151, 1982.
- [60] T.J. Pitcher and J.K. Parrish. Functions of shoaling behaviour in teleosts. *Behaviour of teleost fishes*, 2:369–439, 1993.
- [61] S.C. Pratt, E.B. Mallon, D.J. Sumpter, and N.R. Franks. Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52(2):117–127, 2002.
- [62] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [63] E.J.H. Robinson, N.R. Franks, S. Ellis, S. Okuda, and J.A.R. Marshall. A simple threshold rule is sufficient to explain sophisticated collective decision-making. *PLoS ONE*, 6(5):e19981, 2011.
- [64] F. Saffre, R. Furey, B. Krafft, and J.-L. Deneubourg. Collective decision-making in social spiders: dragline-mediated amplification process acts as a recruitment mechanism. *Journal of theoretical Biology*, 198(4):507–517, 1999.
- [65] T. Schmickl and K. Crailsheim. Trophallaxis within a robotic swarm: bio-inspired communication among robots in a swarm. *Autonomous Robots*, 25(1):171–188, 2008.
- [66] T. Schmickl, R. Thenius, C. Moeslinger, G. Radspieler, S. Kernbach, M. Szymanski, and K. Crailsheim. Get in touch: cooperative decision making based on robot-to-robot collisions. *Autonomous Agents and Multi-Agent Systems*, 18(1):133–155, 2009.
- [67] A.M. Simons. Many wrongs: the advantage of group navigation. *Trends in ecology & evolution*, 19(9):453–455, 2004.
- [68] O. Soysal and E. Şahin. A macroscopic model for self-organized aggregation in swarm robotic systems. In E. Şahin, W. Spears, and A. Winfield, editors, *Swarm Robotics, Second International Workshop, SAB 2006, Rome, Italy, September 30-October 1, 2006, Revised Selected Papers*, volume 4433 of *Lecture Notes in Computer Science*, pages 27–42. Springer Verlag, Berlin, Germany, 2007.
- [69] W. M. Spears, D. F. Spears, J. C. Hamann, and R. Heil. Distributed, physics-based control of swarms of vehicles. *Autonomous Robots*, 17(2):137–162, 2004.
- [70] V. Sperati, V. Trianni, and S. Nolfi. Evolving coordinated group behaviours through maximization of mean mutual information. *Swarm Intelligence*, 2(2–4):73–95, 2008.

- [71] S. H. Strogatz. *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press, New York, NY, 2003.
- [72] K. Sugawara, T. Kazama, and T. Watanabe. Foraging behavior of interacting robots with virtual pheromone. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 3074–3079. IEEE, 2004.
- [73] D.J.T. Sumpter, J. Krause, R. James, I.D. Couzin, and A.J.W. Ward. Consensus decision making by fish. *Current Biology*, 18(22):1773–1777, 2008.
- [74] D.J.T. Sumpter and S.C. Pratt. Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):743–753, 2009.
- [75] V. Trianni and M. Dorigo. Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231, 2006.
- [76] V. Trianni and S. Nolfi. Self-organising sync in a robotic swarm. A dynamical system view. *IEEE Transactions on Evolutionary Computation*, 13(4):722–741, 2009.
- [77] P. Turchin. *Quantitative analysis of movement: measuring and modeling population redistribution in animals and plants*. Sinauer Associates Sunderland, Massachusetts, 1998.
- [78] A.E. Turgut, H. Çelikkanat, F. Gökçe, and E. Şahin. Self-organized flocking in mobile robot swarms. *Swarm Intelligence*, 2(2-4):97–120, 2008.
- [79] C. van Oss, A. V. Panfilov, P. Hogeweg, F. Siegert, and C. J. Weijer. Spatial pattern formation during aggregation of the slime mould *Dictyostelium discoideum*. *Journal of Theoretical Biology*, 181(3):203–213, 1996.
- [80] A.J.W. Ward, D.J.T. Sumpter, I.D. Couzin, P.J.B. Hart, and J. Krause. Quorum decision-making facilitates information transfer in fish shoals. *Proceedings of the National Academy of Sciences*, 105(19):6948, 2008.
- [81] B.B. Werger and M.J. Matarić. Robotic “food” chains: Externalization of state and program for minimal-agent foraging. In P. Maes, M.J. Matarić, J. Meyer, J. Pollack, and S. Wilson, editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, MA, 1996.
- [82] A. T. Winfree. Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology*, 16(1):15–42, 1967.
- [83] S. Wischmann, M. Hulse, J. F. Knabe, and F. Pasemann. Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14(2):117–127, 2006.

- [84] S. Wischmann and F. Pasemann. The emergence of communication by evolving dynamical systems. In S. Nolfi et al., editor, *From animals to animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*, volume 4095 of *LNAI*, pages 777–788. Springer Verlag, Berlin, Germany, 2006.
- [85] David H Wolpert and William G Macready. No free lunch theorems for search. Technical report, Technical Report SFI-TR-95-02-010, Santa Fe Institute, 1995.