

Efficient Decision-Making in a Self-Organizing Robot Swarm: On the Speed Versus Accuracy Trade-Off

Gabriele Valentini¹, Heiko Hamann² and Marco Dorigo²

¹ IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{gvalenti, mdorigo}@ulb.ac.be

² Department of Computer Science,
University of Paderborn, Paderborn, Germany
heiko.hamann@uni-paderborn.de

April 18, 2015

Abstract

We study a self-organized collective decision-making strategy to solve the best-of- n decision problem in a swarm of robots. We define a distributed and iterative decision-making strategy. Using this strategy, robots explore the available options, determine the options' qualities, decide autonomously which option to take, and communicate their decision to neighboring robots. We study the effectiveness and robustness of the proposed strategy using a swarm of 100 Kilobots. We study the well-known speed versus accuracy trade-off analytically by developing a mean-field model. Compared to a previously published simpler method, our decision-making strategy shows a considerable speed-up but has lower accuracy. We analyze our decision-making strategy with particular focus on how the spatial density of robots impacts the dynamics of decisions. The number of neighboring robots is found to influence the speed and accuracy of the decision-making process. Larger neighborhoods speed up the decision but lower its accuracy. We observe that the parity of the neighborhood cardinality determines whether the system will over- or under-perform.

1 Introduction

Governing the increasing complexity of man-made systems in terms of reliability and robustness requires new paradigms of systems engineering. Simplicity might be the key to enable the design of large complex systems based on many simple and autonomous components. In robotics, the design of the Kilobot robot [27] shows how production costs can be drastically reduced by limiting sensory, actuation, and computational requirements. Yet, the control of large collections of

autonomous robots remains an under-investigated challenge that demands novel approaches. As a result of local interactions between individual agents and their environment, swarm intelligence enables the design of simple controllers that are highly scalable and robust to noise and component failures [1, 2]. Swarms of up to 1000 robots have recently been shown to successfully complete tasks such as aggregation [15], collective transport [28] and self-assembly [29]. Here, we focus on the more general task of collective decision-making which is a prerequisite for a number of different swarm applications [22, 11]. Whether the swarm needs to identify the shorter path to traverse, the most suitable shape to form, or the most favorable working location, it first needs to address a quality-dependent collective decision-making problem [9, 18].

We study a self-organized collective decision-making strategy to solve the best-of- n decision problem in a swarm of robots [21]. This problem requires the swarm to establish a collective agreement on the highest-valued option among a finite set of n alternatives. A collective decision-making strategy that solves the best-of- n problem is a distributed decision mechanism capable of processing the information provided by the options' quality in order to drive the swarm towards the formation of a majority for the best option. This mechanism—generally known as the modulation of positive feedback [10]—acts by amplifying or reducing the period of time in which an individual agent actively participates in the decision-making process as a function of the option's quality. Preferences for different options are promoted proportionally to their quality with the best option being favored.

Previous studies focused on the solution of particular instances of the best-of- n problem and rely on domain-specific choices for the modulation of positive feedback; being strictly coupled to the asymmetry of the environment, these design choices cannot be easily transferred to other scenarios. In [9, 4], the quality of an option corresponds to the size of the associated area in the environment. Two decision rules are specifically designed to adjust an agent's probability to change the preference for an area as a function of the area's size. These rules are based on ad hoc defined functions. Their free parameters are determined either empirically or using a genetic algorithm with the aim of rapid and stable collective decisions. In [18, 31, 35] the authors apply the more general majority rule to the problem of finding the shorter of two paths connecting a starting and a target location. Once 3 agents meet in the starting location, they form a team, they mutually share their preferences for a particular path, and eventually apply the majority rule to determine the team's path towards the target location. In this case, the modulation of positive feedback is provided indirectly to the agents via the environment: the shorter the path to traverse, the more frequently agents participate in the decision process favoring that alternative. The same reasoning applies to the study presented in [32, 3] where the majority rule is substituted by the k -unanimity rule. When using the k -unanimity rule, an agent changes its preference towards a particular option only after consecutively encountering k other agents favoring that option. The value of k determines the speed and accuracy of the decision-making process. Reina et al. [25] take inspiration from unifying models of decision-making in honeybees swarms

and vertebrate brains [17] and use direct recruitment and cross-inhibition between pairs of agents. Finally, consensus achievement problems are also widely investigated in the field of control engineering [12, 26, 30]. However, researchers generally focus on continuous decisions-making problems, i.e., problems with an infinite number of equally-valued alternatives which do not require a quality-based discrimination process.

In this paper, we abstract the quality of an option from its particular origins (e.g., brightness level, chemicals concentration). When designing the decision-making strategy, we assume that the robot is equipped with appropriate sensors to determine the quality of the different options and that this is always a bounded measurement. With this abstraction we are able to design more general and portable solutions for scenarios where the options' quality is an absolute metric.

We propose a self-organized decision-making strategy that overcomes the limitations of previous studies by decoupling the modulation of positive feedback from the decision rule. In our strategy, agents directly modulate positive feedback by advertising their preferences for a time proportional to the option quality. This feature is shared with the weighted voter model by Valentini et al. [36]. When using the voter model as a decision rule, agents change their preferences by copying the preference of a randomly chosen neighbor within a limited interaction range. In contrast to the weighted voter model, we implement individual agent decisions using the majority rule. The majority rule speeds up the decision process and enables the implementation of the strategy in very large swarms. We study the well-known speed versus accuracy trade-off in collective decision-making [6, 24]. We test the effectiveness of the proposed strategy with more than 30 hours of robotic experiments using a swarm of 100 Kilobots. The dynamics of collective decisions are modeled by a system of ordinary differential equations (ODEs). We show analytically that the majority rule allows the system to take faster decisions compared to the weighted voter model but with lower accuracy. An influential feature is found to be the spatial density of the robots, which defines the cardinality of a robot's neighborhood (i.e., number of agents within perception range) and affects both the speed and the accuracy of the decision process.

2 Decision-Making Strategy

Our aim is to design a self-organized decision-making strategy that is applicable to different instances of the best-of- n decision problem. Here, the quality of an available option i is a value $\rho_i \in (0, 1]$; that is, we abstract away the particular features of an option that determine its quality. In the considered scenario, the swarm options correspond to regions in the space representing resources with a certain quality. A swarm of N agents is initially located in the *nest*, which is an area functioning as a decision-making hub. The nest provides access to $n = 2$ equally distant *sites*, A and B . Agents in the swarm may perceive the quality $\rho_i, i \in \{A, B\}$, of a site by exploring it. Once returned to the nest,

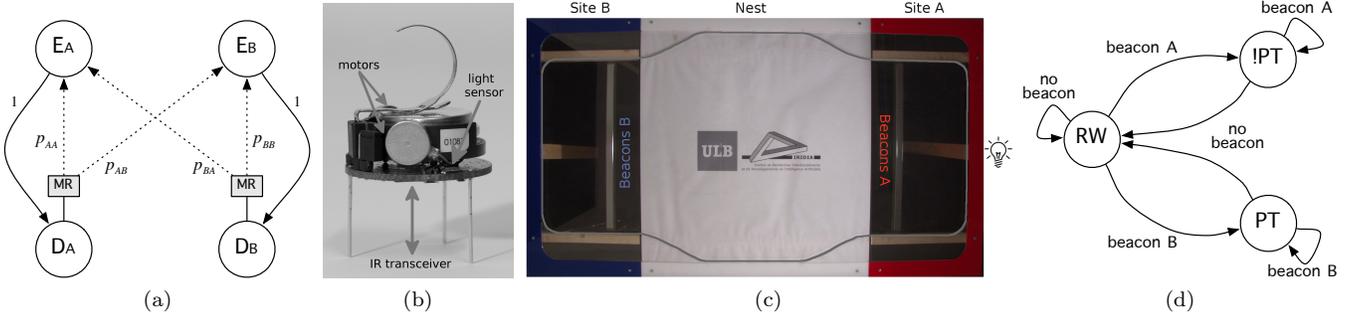


Figure 1: Illustrations of: a) the probabilistic finite-state machine of the individual agent (solid and dotted lines represent respectively deterministic and stochastic transitions; MR, majority rule); b) the Kilobot robot; c) robot arena; d) finite-state machine of robot motion control during the dissemination state (RW, random walk; PT, phototaxis; !PT, antiphototaxis).

the acquired information is exploited during the decision-making process by modulating the positive feedback for the respective site.

We propose a simple, self-organized decision-making strategy. Agents always have a preference for a particular site referred to as the agent’s opinion. The agent control algorithm consists of four control states (Fig. 1a). In the *dissemination* states (D_A, D_B) agents broadcast their opinion within a limited spatial range. Concurrently, they perform defined movements to maintain a well-mixed spatial distribution within the nest [20]. The purpose of the dissemination states is to spread agents’ opinions and to prevent their spatial fragmentation that may create deadlocks [5]. In the *exploration* states (E_A, E_B) agents travel from the nest to the site associated with their current opinion, explore the site, estimate its quality, return to the nest, and proceed to the respective dissemination state. The procedure enables the agents to estimate the quality of a site, that is, to collect a sample measurement. Despite the noisiness of this individual measurement, an average over many agents’ measurements determines the global behavior due to the self-organized nature of the decision-making strategy [21, 36].

A key feature of the proposed strategy is the modulation of positive feedback that promotes the spread of the best opinion. Agents adjust the average time spent in dissemination states proportionally to the opinion’s quality by the product $\rho_i g, i \in \{A, B\}$. The parameter g is the unbiased dissemination time and is set by the system designer. In this way, agents directly control the positive feedback effect of the majority rule by modulating their participation in the decision process and driving it towards the best opinion. A similar modulation mechanism is adopted by house-hunting honeybee swarms in their waggle dance behavior [19, 33]. Before agents leave the dissemination states, they record the opinions of their neighbors. They add their own current opinion to the record and apply the majority rule to determine their next state (gray boxes in Fig. 1a).

Depending on the outcome of the majority rule, agents switch either to E_A or E_B . In the case of ties, agents keep their current opinion.

3 Robotic Scenario

Large-scale robotic experiments are a valuable tool for the validation of a self-organized control strategy. Physics-based multi-agent simulations require simple models of robots and interactions to ensure computational tractability and may limit accuracy. In contrast, the constraints imposed by the real world and real robots allow a more convincing verification of a strategy’s scalability and robustness. We therefore implemented the collective decision-making strategy proposed above in a swarm of 100 Kilobots. A video describing the strategy and the robot experiments can be found in [37].

The Kilobot is a small, low-cost robot commercially available for about €110 that enables experimentation with large groups of embodied agents (see Fig. 1b). It has a diameter of 3.3 cm and a battery that allows for a few hours of autonomy. The Kilobot implements stick-slip motion using three legs and a pair of vibrating motors positioned at its sides. It achieves forward motion at a nominal speed of 1 cm/s and turns in place at up to $\pi/4$ rad/s. The Kilobot is equipped with a light sensor that enables the robot to perceive the intensity of ambient light. Finally, the robot is able to communicate infrared messages of 3 bytes with nearby robots and to sense the distance to the transmitting robot at a range of up to 10–20 cm depending on the reflection properties of the ground surface.

We implemented the proposed self-organized, collective decision-making strategy in a site-selection scenario. $N = 100$ robots are placed in a rectangular arena of 100×190 cm² (Fig. 1c), which is bigger than a single robot’s footprint by a factor of approximately 2×10^3 . The arena is partitioned into three regions: at the two sides there are sites of 80×45 cm² (respectively, site A on the right side and site B on the left side); the nest is at the center and has dimensions 100×100 cm². The goal of the swarm is to reach a large majority of individuals that favor the better site (henceforth site A). A light source positioned at the right side of the arena creates a light gradient and allows for directed navigation between the three areas.

As a consequence of the Kilobots’ limited perception capabilities, we emulated the identification of sites and the estimation of their quality using infrared beacons. Robots perceive the two borders between the sites and the nest from two arrays of beacons positioned under the Perspex surface of the arena. For each border, 5 Kilobots are positioned upside-down under the surface and function as beacons. Each beacon repeatedly broadcasts a message containing the type (A or B) and the quality (ρ_A or ρ_B) of a site. Robots perceive these messages only within the sites in the proximity of the borders (approximately 15 cm) because the area under the nest is covered by lightproof paper and beacons have a limited communication range.

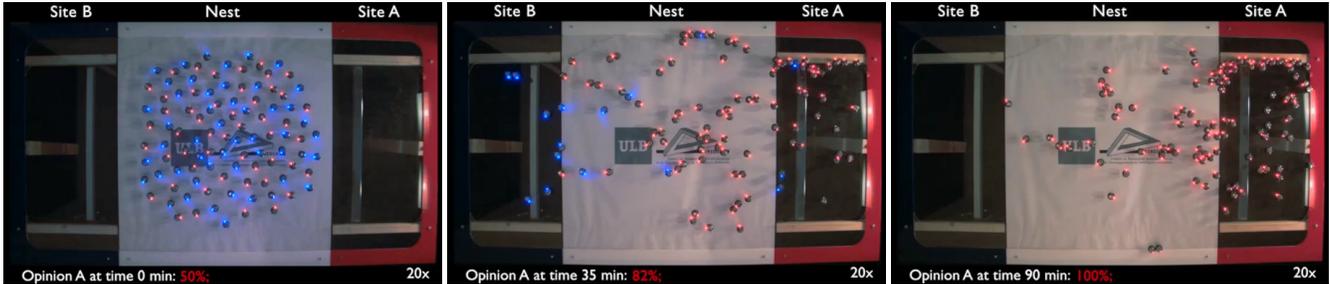


Figure 2: Illustration of an experiment with 100 Kilobots. Screen-shots taken respectively at the beginning of the experiment, $t = 0$ min (left); at approximately half of the decision process, $t = 35$ min (center); and at the end, $t = 90$ min (right).

3.1 Robot Control Algorithm

The proposed collective decision-making strategy meets the requirements of simplicity imposed by the limited capabilities of the Kilobot robot. We develop the decision-making strategy using only the motors, the light sensor, and the infrared transceiver of the Kilobot. Depending on the current perceptions of the environment and on the control state, a robot alternates between three low-level motion behaviors: random motion (*random-walk*) and oriented motion towards or away from a light source (respectively, *phototaxis* and *anti-phototaxis*). In the random-walk behavior, robots perform an exploration-oriented random walk called Lévy flight [14]. The robot moves forward for an exponentially distributed period of time; then, it changes orientation by turning in place for a normally distributed time. In the phototaxis (anti-phototaxis) behavior, robots perform oriented motion towards (away from) the light source placed on the right side of the arena. The robot initially searches for the direction towards (away from) the light source by turning in place and sensing the ambient light. Once the correct direction is found, the robot moves straight until ambient light measurements fall outside a certain tolerance range. At this point the robot resumes the search for the correct direction.

In the dissemination state, robots move within the nest and repeatedly broadcast their opinion as well as a randomly generated 16-bit identifier that (with high probability) uniquely identifies the robot’s opinion in a local neighborhood. To modulate positive feedback, robots spend an exponentially distributed period of time in the dissemination state whose mean duration is given by either ρ_{Ag} or ρ_{Bg} . Note that the mean duration of the dissemination state is a combination of the scaling parameter g , set by the system designer, and of the robot’s current estimate of a site’s quality (either ρ_A or ρ_B), which depends on the problem at hand. Concurrently, robots perform a random walk aimed at stirring the spatial distribution of opinions within the boundaries of the nest (Fig. 1d). If, during this period of time, a robot perceives a message from a

border beacon, for example the border beacon of site A (site B), the robot recognizes that it is mistakenly leaving the nest and enters the anti-phototaxis (phototaxis) behavior with the aim to return to the nest. Before leaving the dissemination state, a robot records the opinions of its neighbors for three seconds, adds its own current opinion, and applies the majority rule to determine which site to explore (possibly switching opinion). The rather short time for opinion collection is required to reduce the time-correlation of the recorded opinions and to prevent individual decisions based on outdated information.

In the exploration state, robots move towards the site associated with their current opinion using the light source as a reference point; they explore the area for an exponentially distributed period of time (emulating the actual estimation of the site’s quality), and then return to the nest. We developed a simple mechanism based on infrared beacons to ensure that transitions between dissemination and exploration states happen only in the nest. This mechanism ensures that a robot which has to explore a site reaches it before moving back to the dissemination state. For a robot in state E_A (respectively, E_B) the phototaxis (anti-phototaxis) behavior is adopted in two stages. Firstly, the robot performs phototaxis (anti-phototaxis) until it perceives a message from beacon A (beacon B), thus ensuring that the robot is correctly entering site A (B). Secondly, the phototaxis (anti-phototaxis) behavior is prolonged for as long as the robot receives messages from beacons A (B). Following this simple mechanism robots enter the site and advance into it. The same mechanism is used by robots in order to return to the nest.

3.2 Robot Experiments

The spatial density of robots may influence the performance of collective systems either positively or negatively [13, 16]. To understand how spatiality affects the proposed self-organized decision-making strategy, we study the effects of different neighborhood sizes \mathcal{N} on the dynamics of the decision process. Specifically, we are interested in understanding how neighborhood size affects the time and the accuracy of the process. We perform two series of robotic experiments where we vary the maximum number of opinion messages that a robot is allowed to receive before applying the majority rule. In doing so, we emulate the effects of different spatial densities on the decision process. We refer in the analysis to the size $\mathcal{G} = \mathcal{N} + 1$ of the opinion group used by a robot which includes its own opinion ($\mathcal{G}_{\max} \in \{5, 25\}$). We consider a scenario where site A is twice as good as site B . We therefore set the sites quality broadcast by the boarder beacons to $\rho_A = 1$ and $\rho_B = 0.5$. Robots are initially located in the nest, initialized in the dissemination state with random opinion (either A or B) and unbiased quality estimation ($\hat{\rho}_A = \hat{\rho}_B = 1$). For each experiment, we perform 10 independent runs using $N = 100$ robots for a duration of 90 min each. The parameter g determines the average duration of the dissemination state prior to its modulation. Too small values for g may prevent robots’ opinions from being spatially well-mixed, hence resulting in opinion fragmentation, while too big values increase the time necessary for the swarm to take a collective decision. Based on

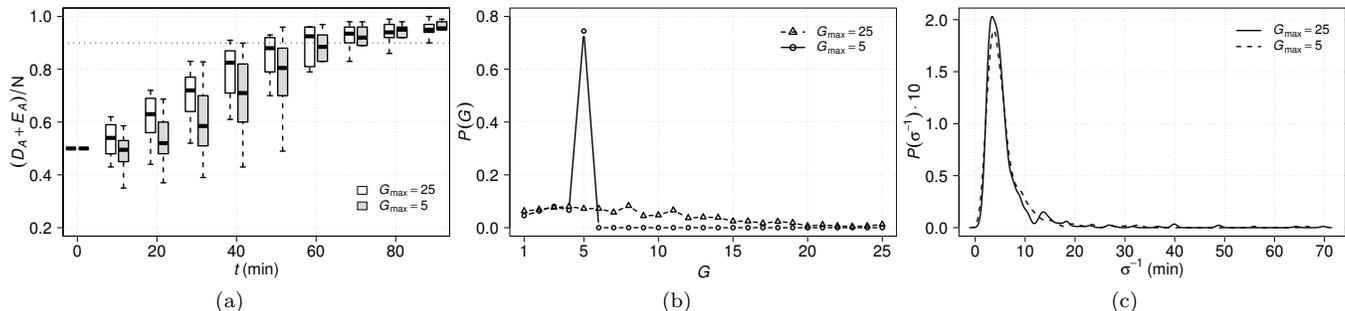


Figure 3: Illustrations of the results of robot experiments: a) proportion of robots with opinion A over time; b) probability distribution of the robot group size G when applying the majority rule; and c) distribution of the time necessary for the robots to complete the exploration state.

a few preliminary experiments we chose to set g to 8.4 min (i.e., about 500 sec). Fig. 2 shows photos of one of the experiments.

We show in Fig. 3a the dynamics of the proportion of robots with opinion A during the decision process $((D_A + E_A)/N)$. When $G_{\max} = 25$, the swarm takes approximately 60 min to exceed a 90% majority of robots in favor of opinion A (white box-plots). When the maximum group size is reduced, $G_{\max} = 5$, the swarm exceeds the 90% majority in around 70 min, hence taking approximately 10 min longer (gray box-plots). Thus, we observe a positive correlation between the speed of the decision process and the average neighborhood size: the bigger the neighborhood the faster the decision process. Additionally, Fig. 3a shows that even though the swarm establishes a large majority of $> 95\%$, the swarm does not reach a 100%-consensus. This is a consequence of limited performance of some robots despite careful calibration and maintenance efforts. Occasionally, robots would require re-calibration during the experiment, hence loose agility, and are less likely to change opinion. At times, robots have battery failures or switch to stand-by due to short-circuits caused by robot-to-robot collisions. Still, the proposed self-organized decision-making strategy proves to be robust by enabling the swarm to establish a large majority for the best option and, therefore, to take a collective decision.

Robot experiments show both the robustness of the self-organized decision-making strategy and the effects of spatial density on the velocity of the decision process. However, the overhead of robot experiments limits the available data and hence our analysis. In order to deepen our understanding concerning the effects of the robot density on the system performance, we collect additional statistics and use this information to define a qualitative mathematical model (see Sec. 4). A first step towards the definition of a proper model is to analyze the spatial distribution of robots during the experiments. We collect data concerning the size G of the group of opinions over which a robot applies the majority rule. Fig. 3b shows the probability mass function $P(G)$ estimated from a single

experimental run for both settings ($\mathcal{G}_{\max} \in \{5, 25\}$). We measure an average group size of 8.57 when robots are allowed to receive up to $\mathcal{G}_{\max} = 25$ messages. The average group size is reduced to 4.4 for $\mathcal{G}_{\max} = 5$.

A second step is that of analyzing the time required by a robot to complete the procedure of the exploration state. Fig. 3c shows the probability density function of the exploration time σ^{-1} . As expected, the distribution of the exploration time in the two settings is similar: $P(\sigma^{-1})$ is approximately centered around its mean value $\sigma^{-1} = 6.06$ min. We also observe a few samples with very high values ($\gg 6.06$ min), that is, $P(\sigma^{-1})$ has a long tail. This result is related to the above discussion of technical problems. That is, due to motion difficulties, a few robots spend a long time in the exploration state. Additionally, we also collect data on how often a robot applies the majority rule on average. Robots take on average 6.65 and 6.96 decisions, respectively, in the first and second setting. Hence, each robot performs about 7 trips going from the nest to one of the two sites and therefore takes about 7 site quality samples.

4 Mean Field Model

With the results of robot experiments at hand, we define a qualitative mathematical model aimed at understanding the influence of spatial densities. We use tools of dynamical system theory to study the effects of the agents' neighborhood size on the dynamics of the decision process. For each opinion, we model the time evolution of the expected proportion of agents in the dissemination state, d_A and d_B respectively, and the expected proportion of agents in the exploration state, e_A and e_B respectively. We assume that (i) robots have a constant neighborhood size \mathcal{N} and (ii) each robot has already a valid quality estimate associated with its initial opinion at time $t = 0$. Although these simplifying assumptions differ from the actual robotic implementation, they enable the definition of a concise mathematical model.

An essential feature of the proposed decision-making strategy is the modulation of the time agents spend in the dissemination state advertising their own opinion. This modulation biases the system dynamics towards consensus on the best opinion and is achieved by letting the agents weight the unbiased duration g of the dissemination state by the quality ρ_A (ρ_B) of their opinion A (B). We define coefficients $\alpha = (\rho_A g)^{-1}$ and $\beta = (\rho_B g)^{-1}$ as shortcuts to represent the rates at which agents move from the dissemination state to the exploration state for opinions A and B . In the robotic experiments we set the design parameter $g = 8.4$ min. We estimate the mean duration σ^{-1} of the exploration state from the data shown in Fig. 3c to be $\sigma^{-1} = 6.06$ min.

We model the outcome of the majority rule by considering the probability p_{AB} that an agent with opinion A switches to opinion B as the result of applying the majority rule for a neighborhood of size \mathcal{N} (similarly for p_{BA}). In addition, we also consider those cases when the majority rule does not trigger a switch which is modeled by probabilities p_{AA} and p_{BB} . Since agents only advertise their opinion when they are in the dissemination state, these prob-

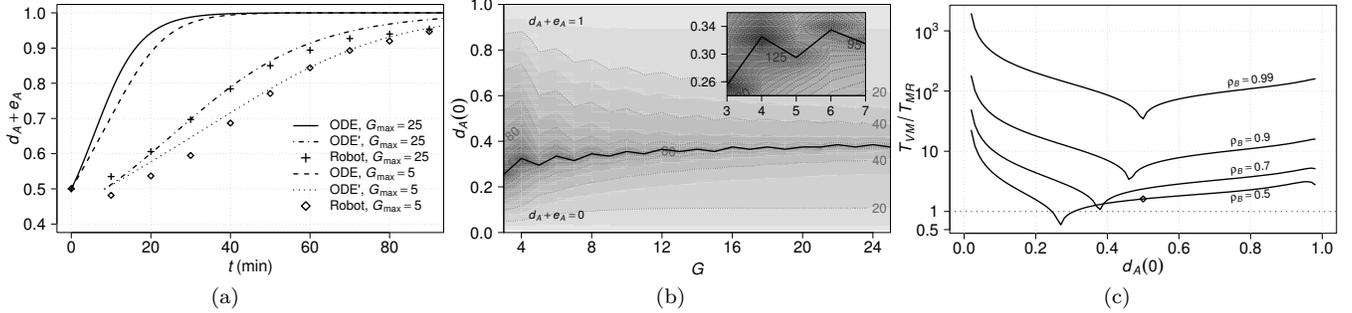


Figure 4: Illustrations of the analysis of the ODE model in Eqs 3: a) predictions of the ODE model against robot experiments; b) heat-map representing the time to consensus and the border between basins of attraction of consensus states (the darker the shades of gray, the longer the decision time); c) decision time ratio between the weighted voter model (T_{VM}) and the majority rule (T_{MR}).

abilities depend only on the opinion distribution d_A and d_B of agents in the dissemination state. The probability to have a neighbor with opinion A is given by $p_A = \frac{d_A}{d_A + d_B}$. For example, for an agent with opinion A and neighborhood size \mathcal{N} , the probability p_{AB} can be defined by considering all possible majorities that make the agent switch opinion towards B . Say $\mathcal{N} = 2$, then an agent switches from A to B if and only if it encounters a neighborhood BB (i.e., two neighboring agents with opinion B). The agent would keep opinion A in the case the neighborhood is AB , BA , or AA . Under the well-mixed assumption, probability p_{AA} is equal to $p_A^2 + 2p_A(1 - p_A)$ while $p_{AB} = (1 - p_A)^2$. For an agent with opinion A and neighborhood size \mathcal{N} , we have

$$p_{AA} = \sum_{i=\lfloor \mathcal{N}/2 \rfloor}^{\mathcal{N}} \binom{\mathcal{N}}{i} p_A^i (1 - p_A)^{\mathcal{N}-i}, \quad (1)$$

$$p_{AB} = \sum_{i=0}^{\lfloor \mathcal{N}/2 \rfloor - 1} \binom{\mathcal{N}}{i} p_A^i (1 - p_A)^{\mathcal{N}-i}. \quad (2)$$

The summations define a discrete integration of a Binomial distribution. Specifically, p_A is the success probability, \mathcal{N} the number of trials, and i the number of successes. Equations for probabilities p_{BB} and p_{BA} are derived by swapping the power indexes in Eqs (1-2).

Table 1: Summary of study parameters: DP, set by designer; RE, estimated from robot experiments; PP, parameter that defines the problem.

<i>Parameter</i>	<i>Value</i>	<i>Type</i>
Quality of Site <i>A</i>	$\rho_A = 1.0$	PP
Quality of Site <i>B</i>	$\rho_B \in \{0.5, \dots, 0.99\}$	PP
Maximum group size	$\mathcal{G}_{\max} \in \{5, 25\}$	DP
Mean neighborhood size	$\mathcal{N} \in \{2, 4, 8\}$	RE
Exploration time	$\sigma^{-1} = 6.06 \text{ min}$	RE
Dissemination time	$g = 8.4 \text{ min}$	DP

Finally, we define a system of 4 ordinary differential equations

$$\begin{cases} \frac{d}{dt}d_A = \sigma e_A - \alpha d_A, \\ \frac{d}{dt}d_B = \sigma e_B - \beta d_B, \\ \frac{d}{dt}e_A = p_{AA}\alpha d_A + p_{BA}\beta d_B - \sigma e_A, \\ \frac{d}{dt}e_B = p_{AB}\alpha d_A + p_{BB}\beta d_B - \sigma e_B. \end{cases} \quad (3)$$

The first two equations model the change in the proportion of agents in the dissemination state. Note that during the dissemination of opinions, agents do not switch opinion. d_A (respectively d_B) increases at a rate σ due to agents returning from the exploration state e_A (respectively e_B) and decreases at a rate α (respectively β) due to agents leaving the dissemination state. The last two equations model the proportion of agents in the exploration state. For the case of e_A , the proportion increases due to agents that were in the dissemination state before and have switched to the exploration state. This includes both agents that were in favor of opinion *A* before and stay with it and agents that were in favor of opinion *B* but have switched to opinion *A*. The proportion e_A decreases at a rate σ due to agents leaving the exploration state (similarly for e_B).

5 Analysis

In this section, we analyze the mean field mathematical model introduced in Sec. 4 with the aim of deepening our understanding of the proposed decision-making strategy. We first validate the mathematical model defined in the system of Eqs (3) against the results of robot experiments. Next, we use the model to study the speed versus accuracy trade-off in reference to the majority rule. Finally, we compare the proposed decision-making strategy with the weighted voter model of Valentini et al. [36]. We summarize all parameters studied in the following analysis in Table 1.

5.1 Model Validation

We validate the predictions of our ODE model by comparing them to our robot experiments. We set the neighborhood size to $\mathcal{N} = 8$ for the case of $\mathcal{G}_{\max} = 25$, and $\mathcal{N} = 4$ for $\mathcal{G}_{\max} = 5$ by approximating the data collected during robot experiments (see Sec. 3.2). We compare transient and asymptotic dynamics of the ODE model to the results from robotic experiments. For both problem settings, Fig. 4a shows that trajectories predicted with the model (labeled ODE, solid and dashed lines) resemble the average robot performance (cross and diamond symbols) but at an increased velocity. The prediction of the transient dynamics improves when scaling the time as $t' = 3(t + g)$ (labeled ODE', dotted and dot-dashed lines). That is, robot experiments are approximately 3 times slower than predictions of the ODE model. In addition, the offset represented by g in t' is due to the fact that, in the experiments, robots do not have an initial estimate of their opinions, which need an average time g to be acquired. The discrepancies between the transient dynamics of the model and those of robot experiments are a consequence of the simplifying assumptions (i) and (ii) of the model (see Sec. 4); of the fact that the model is a continuous approximation and does not account for finite-size effects; and of spatial interferences among robots that, by increasing the spatial correlation of opinions, make the system depart from the well-mixed assumption underlying the model. The asymptotic dynamics are correctly predicted by the model. The stability analysis of the system of Eqs (3) determines three roots: two asymptotically stable solutions correspond to consensus on A ,

$$\{d_A = \frac{\rho_A g}{\rho_A g + \sigma^{-1}}; d_B = 0; e_A = \frac{\sigma^{-1}}{\rho_A g + \sigma^{-1}}; e_B = 0\},$$

and consensus on B ,

$$\{d_A = 0; d_B = \frac{\rho_B g}{\rho_B g + \sigma^{-1}}; e_A = 0; e_B = \frac{\sigma^{-1}}{\rho_B g + \sigma^{-1}}\};$$

the third solution is an unstable fixed point.

5.2 Speed Versus Accuracy Trade-Off

We analyze the system of Eqs (3) with the aim to understand the speed versus accuracy trade-off [6, 24] in our robot system. On the one hand, we want to determine how the initial distribution of opinions among the agents affects the performance of the swarm. On the other hand, we want to quantify, at least qualitatively, how the spatial density of robots influences the speed and accuracy of the collective decision, i.e., how an increase or decrease in the group size \mathcal{G} impacts the performance of the swarm.

Fig. 4b shows the speed versus accuracy trade-off arising from different parameter configurations. The solid line is the border that separates the basins of attraction between the two asymptotically stable solutions, respectively, consensus on A and on B . This border increases roughly logarithmically with the

group size \mathcal{G} . The higher the value of the border, the smaller the basin of attraction of the best option (site A). This result indicates that increasing \mathcal{G} reduces the accuracy and the robustness of the decision-making process (note that fluctuations could move the system towards B even from the unbiased initial condition $d_A = d_B = 0.5$). The time necessary to take a decision increases with the proximity of the initial conditions to the border of the basins of attraction and decreases for increasing values of the group size \mathcal{G} . This is shown in Fig. 4b where the darker the shades of gray, the longer is the time necessary for a collective decision. The inset in Fig. 4b highlights an unexpected feature of the majority rule. When the size \mathcal{G} of the group of opinions used in the majority rule is even, the decision process takes longer to complete and the best opinion’s basin of attraction shrinks. Consequently, the accuracy is reduced for even group sizes. Groups of even sizes have a chance of ties. In case of ties, the agent keeps its current opinion, which eventually favors the current majority at the swarm level [7, 8]. This analysis raises the interesting question of whether tie-breakers could increase accuracy and/or speed. Future empirical and theoretical studies will focus on the design of mechanisms to improve the swarm performance in the case of even group sizes (e.g., random tie-breaker).

5.3 Weighted Voter Model Comparison

The proposal of a new approach to a certain problem requires a comparison to existing state-of-the-art solutions. However, as discussed in the introduction, most of the existing algorithms for the best-of- n decision problem rely on the exploitation of particular features of the environment for the modulation of positive feedback. In the scenario tackled in this study, we deliberately avoided such features with the aim of devising a more general decision-making strategy. As a consequence, these algorithms do not directly apply to our scenario with the only exception being the weighted voter model [36]. Although a comparison based on robot experiments would be desirable, the high cost of experiments would result in a limited number of independent runs and a resulting low significance of statistical tests. Instead, we use the ODE model introduced in Sec. 4 and the one discussed in [36]. We compare analytically the performance of these decision-making strategies using the same parameter values (with the exception of the average neighborhood size \mathcal{N}).

We therefore conclude the analysis with a comparison of the proposed majority rule-based decision-making strategy with the weighted voter model. Fig. 4c shows the ratio T_{VM}/T_{MR} between the decision time T_{VM} of the weighted voter model and the decision time T_{MR} of the majority rule for $\mathcal{N} = 2$ and different values of $\rho_B \in \{0.5, \dots, 0.99\}$. The majority rule enables decisions that are generally faster than those of the weighted voter model, reaching in extreme cases a speed-up of up to 10^3 . In a limited range of parameter configurations, we observe a slow-down of the majority rule around the border between different basins of attraction of consensus states (solid line with $T_{VM}/T_{MR} < 1$). The diamond in Fig. 4c indicates the speed-up of 1.61 provided by the majority rule in the investigated robotic scenario. In summary, the majority rule allows for much

faster decisions at the cost of reduced accuracy while the weighted voter model takes much longer to establish a decision but guarantees the optimal solution.

6 Discussion

In this section, we deepen our discussion concerning the primary contributions of this study. We first focus on the the proposed self-organized decision-making strategy and we highlight the differences from previous works. Then, we discuss the implementation of this strategy on real robots and we summarize the relevant design choices.

6.1 Decision-Making Strategy

The generality of the proposed self-organized decision-making strategy is a result of the abstraction of options’ qualities from their particular origins. This abstraction allows us to employ a simple and general decision-making mechanism: the majority rule. The majority rule is invariant to the number of options, which could possibly be unknown. Therefore, the proposed strategy directly applies to the general case of n different options (agents can replicate the dissemination and exploration states for each newly encountered option). Also the reported ODE model is easily generalized to the case of $n > 2$ options. More options are modeled by introducing new equations in the system of Eqs (3) and by adapting Eqs (1–2) to sum over a Multinomial distribution instead of a Binomial distribution.

The proposed decision-making strategy substantially differs from the canonical majority rule model. In the canonical majority rule model [7, 8], all agents are perpetually located in a unique environment and repeatedly apply the majority rule over a set of equally valued alternatives. The focus of the canonical model is therefore on breaking the symmetry between the different alternatives. As a consequence, it lacks a mechanism that allows the agents as a whole to process the information provided by the options’ qualities. In the proposed decision-making strategy, this mechanism is represented by the direct modulation of the duration of options’ promotion performed by the agents—modulation that takes place in the dissemination state. Provided sufficient initial conditions are met (see Fig. 4b), this modulation mechanism gives the swarm the momentum necessary to steer the collective decision towards the best alternative. Note that neither of these two mechanisms alone would solve the considered scenario. On the one hand, using only the majority rule would not favor the best alternative. On the other hand, the sole modulation in the dissemination state does not provide agents with the means to change their opinions and eventually take a collective decision.

In contrast to the weighted voter model proposed by Valentini et al. [36], we implement individual agents’ decisions using the majority rule. In the limit of an infinite number of agents, the weighted voter model guarantees consensus on the best option available to the swarm. However, there is a trade-off—as shown

in Fig. 4c, this extremely high accuracy comes at the cost of much longer convergence times. This feature reduces the overall efficiency of the decision-making strategy and may also prevent designers from using the weighted voter model whenever the chosen robot platform suffers from limited energy autonomy—as was the case for our robot experiments. In contrast, the majority rule allows the swarm to take much faster collective decisions, at the cost of reducing the overall accuracy. Nonetheless, it allows us to implement the decision-making strategy and to perform tests using a swarm of 100 Kilobots. Therefore, the proposed decision-making strategy is preferable in scenarios where a short decision time is the main goal whilst the weighted voter model should be considered whenever the accuracy of the collective decision is of primary importance (provided that the used robot platform has sufficient autonomy). Furthermore, the majority rule gives designers an additional degree of freedom: the size of the group of opinions used for individual agents’ decisions. Given a sufficiently high spatial density, designers may operate on the maximum group size with the aim of calibrating the speed and the accuracy of the collective decision.

6.2 Robot Implementation

The Kilobot is very limited in its sensory capabilities and its performance is extremely sensitive to passive obstacles that cannot be perceived, such as walls or objects scattered in the environment. For example, a Kilobot may get stuck at one of the arena walls. Due to the lack of appropriate sensors, the robot would not be capable to detect this problematic situation, and therefore, to trigger appropriate escape actions. The situation worsens when many Kilobots are placed in a small environment. They may start to form clusters near a wall which might result in the failure of the overall experiment. In addition, the Kilobot is lightweight, and therefore it is sensitive to the level of the surface it is operating on. In order to minimize the undesirable effects resulting from the arena’s walls and surface level, we designed the arena with a slightly concave profile, in contrast to a perfectly flat and level surface. In this way, when moving, the robots have a natural slight tendency to move towards the center of the arena, reducing the chances of forming clusters at walls that are sufficiently large to block many robots.

In the robot implementation of the proposed decision-making strategy, robots spend an exponentially distributed period of time in both the dissemination and exploration states. The exponential distribution is characterized by high variance, and therefore, this choice injects a certain amount of noise into the system. Although other choices are possible, which may result in faster decisions, we found the exponential distribution extremely beneficial in the robot experiments. By injecting noise into the system, we prevent synchronization effects among robots that, with high probability, would result in spatial fragmentation of robots’ opinions and consequently would prevent the system from achieving a collective decision. That is, a certain level of noise helps to maintain a well-mixed distribution of opinions within the nest. As a result of its self-organized nature, the speed and the accuracy of the proposed decision-making strategy is

minimally affected by noise, as highlighted in [36] for the case of the weighted voter model. Additionally, we found this choice to be beneficial for minimizing the formation of clusters of robots near the arena’s walls—a scenario that is more likely to occur when robots are synchronized.

Finally, we investigated the cause of the relatively long time that is necessary to take a decision. Is it a consequence of the implemented decision-making strategy itself or does it depend considerably on the robots’ limited speed? From data collected during the robot experiments we know that, during the 90 min of execution, each robot performs only about 7 trips going from the nest to one of the two sites and back. Thus, a robot takes about 7 samples of site qualities. Moreover, from Fig. 3a we know that a collective decision is already taken after about 60 min on average. Hence, we can expect the number of necessary samples to be approximately two-thirds of what was measured for the full 90 min. Given that the number of site visits is small, we conclude that the long execution time is a result of the limited speed of the Kilobots. A seemingly more efficient alternative could be to have each robot visit all sites first and then let it compare its estimate to determine the best option. In this case, only 2 site visits would be required. However, such a decision-making strategy would not be self-organized, would not utilize cooperation among the robots, and would consequently suffer from several drawbacks. On the one hand, the direct comparisons of option qualities is more sensitive to noisy estimates [23]. On the other hand, it would reduce the efficiency of the swarm with increasing number of alternative options. Indeed, when the information from the environment is processed in a self-organized manner, agents are not required to estimate the quality of all available options individually, as shown by studies of house-hunting honeybees [33, 34], where the large majority of bees visit only one or two of the available alternatives.

7 Conclusions

We described a self-organized collective decision strategy to solve the best-of- n decision-making problem in a swarm of robots. This strategy couples a time-modulation mechanism, that biases the system dynamics towards consensus on the best option, with individual robots’ decisions based on the majority rule. We have shown that our strategy can be implemented on a swarm of 100 robots with minimal actuation, perception, and computational capabilities. The robot experiments prove the feasibility of our strategy in large swarms and its robustness to robot failures. The consensus states are the only asymptotically stable solutions, as shown using the mean-field model. Using this model, we have (i) explored the speed versus accuracy trade-off that arises in the majority rule from different robot densities and (ii) compared the proposed strategy to the weighted voter model [36]. The comparison shows that our strategy speeds up the decision-making process considerably. We have shown that the speed of the decision-making process increases with the neighborhood size while its accuracy decreases.

Future empirical work will include the comparison of the proposed strategy with the weighted voter model through robot experiments as well as the investigation of direct comparisons of opinion qualities. As discussed in Sec. 5.2, room for improvement exists for individual agents' decisions in the case of even group sizes. We will investigate the effect of tie-breakers in future research. Further theoretical studies will be focused on the general case of n alternative options by extending the mean field model of Eqs 3 following our discussion in Sec. 6.1. Additional empirical and theoretical work will consider the network dynamics arising from the robots' spatial interactions. Building on the interaction network abstraction, we will focus on the definition of quantitative stochastic models for the accurate prediction of decision dynamics.

8 Acknowledgments

This work has been partially supported by the European Research Council through the ERC Advanced Grant “E-SWARM: Engineering Swarm Intelligence Systems” (contract 246939). Marco Dorigo acknowledges support from the Belgian F.R.S.–FNRS.

References

- [1] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
- [2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- [3] A. Brutschy, A. Scheidler, E. Ferrante, M. Dorigo, and M. Birattari. “Can ants inspire robots?” Self-organized decision making in robotic swarms. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 4272–4273. IEEE Press, 2012.
- [4] A. Campo, S. Garnier, O. Dédriche, M. Zekkri, and M. Dorigo. Self-organized discrimination of resources. *PLoS ONE*, 6(5):e19888, 2010.
- [5] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(01n04):87–98, 2000.
- [6] N. R. Franks, A. Dornhaus, J. P. Fitzsimmons, and M. Stevens. Speed versus accuracy in collective decision making. *Proc. R. Soc. Lond. B*, 270:2457–2463, 2003.
- [7] S. Galam. Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. *Journal of Mathematical Psychology*, 30(4):426–434, 1986.

- [8] S. Galam. Real space renormalization group and totalitarian paradox of majority rule voting. *Physica A: Statistical Mechanics and its Applications*, 285(1-2):66–76, 2000.
- [9] 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.
- [10] S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.
- [11] H. Hamann, T. Schmickl, H. Wörn, and K. Crailsheim. Analysis of emergent symmetry breaking in collective decision making. *Neural Computing and Applications*, 21(2):207–218, 2012.
- [12] Y. Hatano and M. Mesbahi. Agreement over random networks. *IEEE Transactions on Automatic Control*, 50(11):1867–1872, 2005.
- [13] C. Huepe, G. Zschaler, A.-L. Do, and T. Gross. Adaptive-network models of swarm dynamics. *New Journal of Physics*, 13(7):073022, 2011.
- [14] N. E. Humphries, H. Weimerskirch, N. Queiroz, E. J. Southall, and D. W. Sims. Foraging success of biological Lévy flights recorded in situ. *Proceedings of the National Academy of Sciences*, 109(19):7169–7174, 2012.
- [15] 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.
- [16] K. Lerman and A. Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2):127–141, 2002.
- [17] J. A. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, and N. R. Franks. On optimal decision-making in brains and social insect colonies. *Journal of The Royal Society Interface*, 2009.
- [18] M. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5:305–327, 2011.
- [19] A. L. Nevai, K. M. Passino, and P. Srinivasan. Stability of choice in the honey bee nest-site selection process. *Journal of Theoretical Biology*, 263(1):93–107, 2010.
- [20] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006.
- [21] C. A. C. Parker and H. Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251, 2009.

- [22] C. A. C. Parker and H. Zhang. Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4:199–220, 2010.
- [23] C. A. C. Parker and H. Zhang. Biologically inspired collective comparisons by robotic swarms. *The International Journal of Robotics Research*, 30(5):524–535, 2011.
- [24] K. M. Passino and T. D. Seeley. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behavioral Ecology and Sociobiology*, 59(3):427–442, 2006.
- [25] A. Reina, M. Dorigo, and V. Trianni. Towards a cognitive design pattern for collective decision-making. In M. Dorigo, M. Birattari, S. Garnier, H. Hamann, M. Montes de Oca, C. Solnon, and T. Stützle, editors, *Swarm Intelligence*, volume 8667 of *LNCS*, pages 194–205. Springer, 2014.
- [26] W. Ren, R. Beard, and E. Atkins. A survey of consensus problems in multi-agent coordination. In *Proceedings of the 2005 American Control Conference*, pages 1859–1864 vol. 3. IEEE Press, 2005.
- [27] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal. Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems*, 62(7):966–975, 2014.
- [28] M. Rubenstein, A. Cabrera, J. Werfel, G. Habibi, J. McLurkin, and R. Nagpal. Collective transport of complex objects by simple robots: Theory and experiments. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '13*, pages 47–54. IFAAMAS, 2013.
- [29] M. Rubenstein, A. Cornejo, and R. Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799, 2014.
- [30] G. Sartoretti, M.-O. Hongler, M. de Oliveira, and F. Mondada. Decentralized self-selection of swarm trajectories: from dynamical systems theory to robotic implementation. *Swarm Intelligence*, 8(4):329–351, 2014.
- [31] A. Scheidler. Dynamics of majority rule with differential latencies. *Phys. Rev. E*, 83:031116, 2011.
- [32] A. Scheidler, A. Brutschy, E. Ferrante, and M. Dorigo. The k-unanimity rule for self-organized decision making in swarms of robots. Technical Report TR/IRIDIA/2011-023, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium, 2011.
- [33] T. D. Seeley. *Honeybee Democracy*. Princeton University Press, 2010.
- [34] T. D. Seeley and S. C. Buhrman. Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology*, 45(1):19–31, 1999.

- [35] G. Valentini, M. Birattari, and M. Dorigo. Majority rule with differential latency: An absorbing Markov chain to model consensus. In *Proceedings of the European Conference on Complex Systems 2012*, Springer Proceedings in Complexity, pages 651–658. Springer, 2013.

- [36] G. Valentini, H. Hamann, and M. Dorigo. Self-organized collective decision making: The weighted voter model. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '14, pages 45–52. IFAAMAS, 2014.
- [37] G. Valentini, H. Hamann, and M. Dorigo. Self-organized collective decisions in a robot swarm. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence, AI Video Competition*. AAAI Press, 2015. http://youtu.be/5lz_HnOLBW4.