Swarm Intell (2011) 5:305–327
DOI 10.1007/s11721-011-0062-z

Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making

Marco A. Montes de Oca · Eliseo Ferrante · Alexander Scheidler · Carlo Pinciroli · Mauro Birattari · Marco Dorigo

Received: 11 November 2010 / Accepted: 9 September 2011 / Published online: 10 November 2011
© Springer Science + Business Media, LLC 2011

Abstract Collective decision-making is a process whereby the members of a group decide on a course of action by consensus. In this paper, we propose a collective decision-making mechanism for robot swarms deployed in scenarios in which robots can choose between two actions that have the same effects but that have different execution times. The proposed mechanism allows a swarm composed of robots with no explicit knowledge about the difference in execution times between the two actions to choose the one with the shorter execution time. We use an opinion formation model that captures important elements of the scenarios in which the proposed mechanism can be used in order to predict the system's behavior. The model predicts that when the two actions have different average execution times, the swarm chooses with high probability the action with the shorter average execution time. We validate the model's predictions through a swarm robotics experiment in which robot teams

Electronic supplementary material The online version of this article (doi:10.1007/s11721-011-0062-z) contains supplementary material, which is available to authorized users.

M.A. Montes de Oca (✉) · E. Ferrante · A. Scheidler · C. Pinciroli · M. Birattari · M. Dorigo
IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
e-mail: mmontes@ulb.ac.be

E. Ferrante
e-mail: eferrant@ulb.ac.be

A. Scheidler
e-mail: ascheidler@iridia.ulb.ac.be

C. Pinciroli
e-mail: cpinciro@ulb.ac.be

M. Birattari
e-mail: mbiro@ulb.ac.be

M. Dorigo
e-mail: mdorigo@ulb.ac.be

Present address:
M.A. Montes de Oca
Dept. of Mathematical Sciences, University of Delaware, Newark, DE, USA
e-mail: mmontes@math.udel.edu

must choose one of two paths of different length that connect two locations. Thanks to the proposed mechanism, a swarm made of robot teams that do not measure time or distance is able to choose the shorter path.

Keywords Opinion dynamics · Differential latency · Collective decision-making · Self-organization · Swarm intelligence · Swarm robotics

1 Introduction

When a person is immersed in a social context, his/her decisions are influenced by those of others. The effects of social influence on the collective-level behavior of groups of people have been studied by economists and sociologists since at least the 1970s (Schelling 1978; Granovetter 1978). More recently, statistical physicists have developed models to quantitatively describe social and economic phenomena that involve large numbers of interacting people (Chakrabarti et al. 2006; Castellano et al. 2009; Helbing 2010). Some of the models that have emerged from these efforts are referred to as *opinion formation models*.

Krapivsky and Redner (2003) proposed a binary opinion formation model in which a population of agents reaches a consensus with high probability on the opinion initially favored by more than 50% of the population. The process that drives the system to consensus is based on the repeated application of the majority rule at a local level on small teams of agents (see Sect. 2). This model is interesting from a swarm intelligence perspective because the resulting opinion dynamics can be seen as a decentralized collective decision-making process in which the opinion favored by more than half of the population spreads to the rest of the population. However, to be of practical use, the majority-rule opinion dynamics needs to make an initially unbiased population (i.e., a population with 50%–50% opinion distribution) reach consensus on the opinion associated with the "best" alternative. In this paper, we demonstrate that it is possible to achieve this goal in a swarm robotics context. In particular, if agents represent robots, and opinions represent actions whose execution temporarily prevents robots from changing opinion or influencing other robots, we demonstrate that an unbiased population can reach consensus on the action that takes less time, on average, to perform.

We begin by introducing a number of modifications to Krapivsky and Redner's model in order to capture important elements of the interaction of real robots with a physical environment (see Sect. 3). The most important of these modifications builds on the concept of latency, which is a period of time of stochastic duration during which an agent cannot be influenced by other agents, and thus cannot change opinion (Lambiotte et al. 2009). We call this modification *differential latency* because in our model the duration of a latency period is different for different opinions. We demonstrate, both analytically and through Monte Carlo simulation, that with the introduced modifications, a population of agents reaches consensus on the opinion associated with the shorter average latency even if that opinion is initially favored by less than half of the population. The predictions of the majority-rule opinion formation model with differential latency are then validated through a swarm robotics experiment that resembles the well-known double bridge experiment designed by Goss et al. (1989) (see Sect. 4). We show that a swarm of robots, which do not measure distance or time, selects the shorter of two paths that connect two locations.

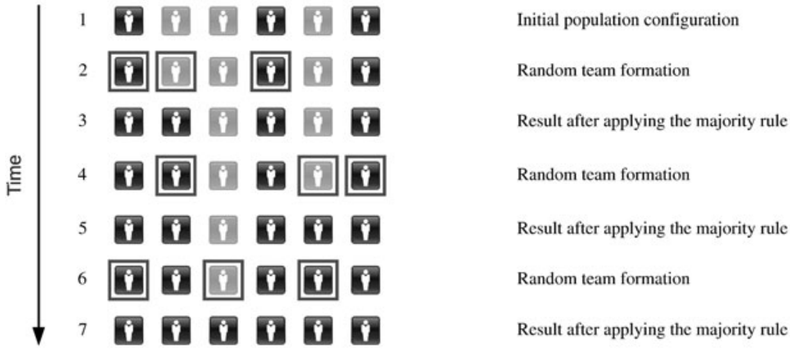

Fig. 1 Majority-rule opinion dynamics. Initially, three agents have opinion A (represented in *black*) and three others have opinion B (represented in *gray*). In this example, after applying three times the majority rule on randomly-formed teams of three agents each (marked with *squares*), the population has reached consensus on one of the two opinions

2 Majority-rule opinion dynamics without and with latency

The majority rule as an element of opinion formation models was first used by Galam (1986) to study voting in hierarchical structures. Krapivsky and Redner (2003) studied the dynamics induced by the majority rule in a well-mixed population case, that is, a situation where everybody can interact with the same probability with everybody else (Nowak 2006). In Krapivsky and Redner's model, a population of agents, each of which can assume one of two states, called opinions (A or B),[1] evolves as follows: First, a team of three randomly chosen agents is formed. Then, the team members adopt the opinion held by the majority within the team. Finally, the team members are put back in the population and the process is repeated (see Fig. 1).

An important aspect of the system's dynamics is the probability of reaching consensus on one opinion, say A, as a function of the initial fraction of the population favoring it (see Fig. 2(a)). In Krapivsky and Redner's model, if the initial fraction of the population in favor of opinion A is greater than a critical initial fraction of 0.5 (i.e., they represent the majority at the population level), then the population reaches consensus on opinion A with a higher probability than on opinion B. If the initial fraction of the population favoring opinion A is 0.5, then the probability of reaching consensus on opinion A is also 0.5. The transition of the probability of reaching consensus on one opinion beyond the initial critical fraction is sharper with larger populations. The average number of team formations required to reach consensus in the majority-rule opinion formation model also depends on the initial fraction of the population favoring one opinion and on the population size. At the critical initial fraction, the system takes the longest to reach consensus. Additionally, the larger the population size, the more team formations are needed to reach consensus (see Fig. 2(b)).

Lambiotte et al. (2009) extended Krapivsky and Redner's model by incorporating *latency*. In the extended model, a team is formed with three randomly picked agents that can be either latent or non-latent. The team's majority opinion is adopted only by the team's non-latent agents. If the team's non-latent agents switch opinion as a result of the majority rule, then they become latent, otherwise they remain non-latent. The team's latent agents

[1] Throughout this paper, we use letters A and B to label the two available opinions.

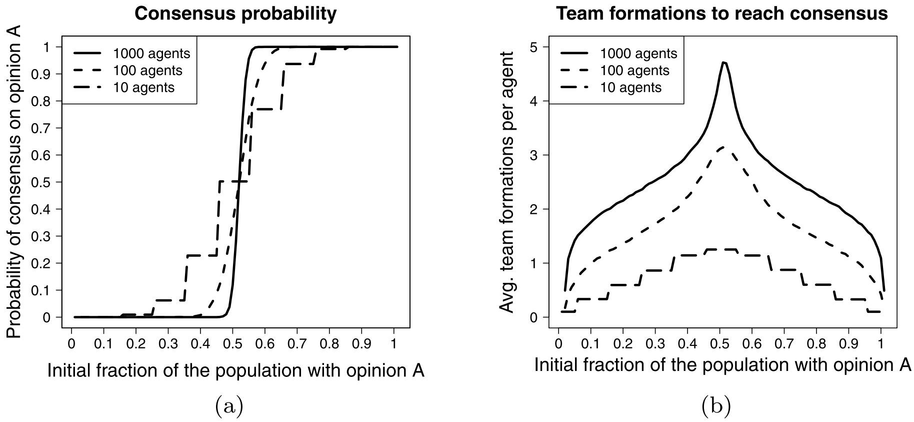

Fig. 2 Dynamics of Krapivsky and Redner's model. (**a**) Estimated probability of reaching consensus on one opinion (labeled A) as a function of the initial fraction of the population in its favor. (**b**) Average number of team formations per agent needed to reach consensus on one opinion. Results are based on averages of 1,000 independent runs of a Monte Carlo simulation

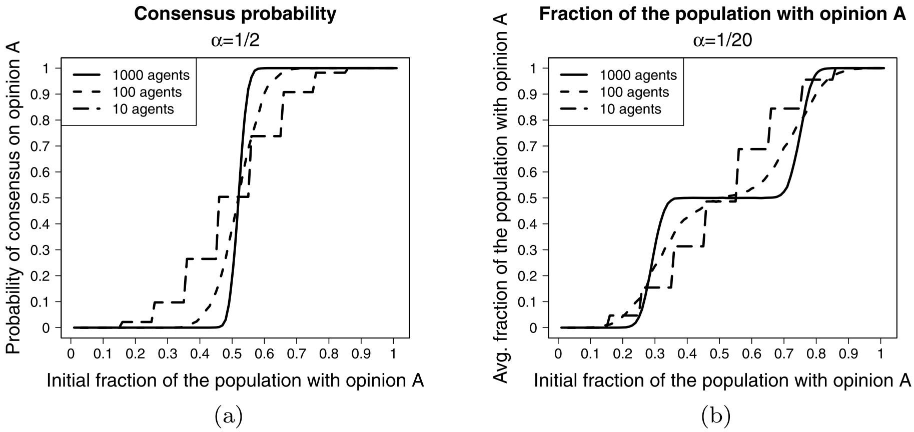

Fig. 3 Dynamics of Lambiotte et al.'s model. Depending on the value of the parameter α, consensus may or may not be the only stable state of the system. (**a**) When $\alpha = 1/2$ consensus is always achieved (we plot the estimated probability of reaching consensus on opinion A). (**b**) When $\alpha = 1/20$ the population does not always achieve consensus (we plot the average fraction of agents with opinion A after 100,000 team formations). Results are based on averages of 1,000 independent runs of a Monte Carlo simulation

become non-latent with probability α, which is a parameter of the model. Lambiotte et al. showed that when $\alpha \geq 1/4$, the system's final state is the same that it would reach without latency (see Fig. 3(a)). However, when $\alpha < 1/4$, consensus is not the only stable state. When $\alpha < 1/4$, a state in which the fraction of the population favoring one opinion fluctuates around 0.5 is also stable (see Fig. 3(b)).

3 Majority-rule opinion dynamics with differential latency

In Lambiotte et al.'s extension of Krapivsky and Redner's model, the duration of latency periods is independent of the opinion adopted by the agents. In this section, we introduce the idea of *differential latency* and apply it to Krapivsky and Redner's model. In a differential latency setting, the duration of latency periods depends on the opinion adopted by the agents.

We interpret agents as robots, opinions as actions that robots have to repeatedly execute while solving a task, and latency as the state of being executing the action associated with the adopted opinion. Under our interpretation, it is reasonable to assume that different robot actions may have different durations. Hence, in our setting, opinions may be associated with latency periods of different duration.

The rest of this section is structured as follows. First, in Sect. 3.1, we describe the main features of the proposed opinion formation model. Then, in Sect. 3.2, we analytically study the proposed system's dynamics under some simplifying assumptions using an ordinary differential equations (ODE) model. Finally, in Sect. 3.3, we relax some of the assumptions made in the formulation of the ODE model and investigate the system's behavior using Monte Carlo simulation.

3.1 Main features of the majority-rule opinion formation model with differential latency

The main features of our opinion formation model are the following:

1. k independent teams of three individuals are formed simultaneously. This feature models the fact that in a real swarm robotics scenario robots can form multiple teams concurrently. The number of teams, k, is a parameter of our model. Our model is different from Krapivsky and Redner's and Lambiotte et al.'s models because in them only one team at a time is formed.
2. Teams are formed with non-latent agents only. As a result, latent agents cannot change opinion and they cannot contribute to other agents' opinion changing process. This feature models situations in which a robot cannot be part of two teams at the same time (e.g., when an action requires some form of physical interaction among team members). Our model is similar to Lambiotte et al.'s model in that latent agents cannot change opinion; however, our model is different from Lambiotte et al.'s model in that latent agents cannot influence other agents.
3. Agents become latent regardless of whether they changed opinion or not after team formation. This feature models situations in which agents immediately execute the action associated with the team's majority opinion. In Lambiotte et al.'s model, agents become latent only if they switch opinions after team formation.
4. Different opinions are associated with latency periods of different duration. This feature is what we call *differential latency*. As explained above, differential latency models situations in which different actions take different amounts of time to perform. As in Lambiotte et al.'s model, the duration of latency periods is stochastic.

A system governed by the proposed model evolves as follows: Initially, agents are non-latent and each agent has a certain opinion. Then, k teams of three randomly chosen agents are formed. The majority rule is used within each team in order to update its members' opinions. Agents that belong to a team enter a latent state whose duration depends on the team's adopted opinion. When a team's latency period finishes, its agents become non-latent and eligible to form a new team. As soon as a team of agents becomes non-latent, a new team is formed. This new team's members are picked randomly from the whole population of non-latent agents. The process is repeated until the population reaches a consensus.

Table 1 Notation used in our ODE model

State	Opinion A	Opinion B	Total
Latent	$l(t)$	$1-\beta-l(t)$	$1-\beta$
Non-latent	$n(t)$	$\beta-n(t)$	β
Total	$n(t)+l(t)$	$1-n(t)-l(t)$	1

3.2 Analytical study

Our analytical study has two goals: (i) to determine the effects of differential latency on the probability of reaching consensus, and (ii) to determine whether consensus is a stable state of the system's dynamics. We investigate the dynamics generated by the proposed model in the continuum limit using a set of ordinary differential equations. For this purpose, we assume the existence of a very large but finite population of constant size. This assumption allows us to focus on the evolution of the latent and non-latent fractions of the population that are in favor of the different opinions. We also assume that an agent's probability of switching from a latent to a non-latent state is constant over time, that is, we assume that the duration of latency periods is exponentially distributed. These assumptions, a large population, and constant probability of switching states, simplify the analysis and are standard practice in the statistical physics literature—see Castellano et al. (2009) and Lambiotte et al. (2009).

3.2.1 Preliminaries

At the microscopic level, agents can be non-latent or latent with opinion A or B, and switch between the resulting four states according to the dynamics of the proposed model. At the macroscopic level, we can aggregate the number of agents in any of the four states and focus on the evolution of the corresponding fractions of the population. Thus, in our analysis, for any $t \geq 0$, $n(t)$ represents the fraction of the population in a non-latent state with opinion A. Likewise, $l(t)$ represents the fraction of the population in a latent state with opinion A. Hence, $n(t)+l(t)$ is the fraction of the population with opinion A and $1-n(t)-l(t)$ is the fraction of the population with opinion B.

In the proposed opinion dynamics model, the number of teams, k, is fixed. In our ODE model, the fixed number of teams is translated into a constant fraction of the population in a latent state (under our robotics-based interpretation of agents and latency, this means that there is always a certain number of robot teams executing an action). Thus, the fraction of the population in a latent state is a parameter of the ODE model, which we denote by $1-\beta$. The constant $\beta \in (0, 1)$ represents the fraction of the population in a non-latent state.

The fractions $n(t)$ and $l(t)$ are sufficient to completely describe the system because the fraction of the population in a non-latent state with opinion B is $\beta - n(t)$ and the fraction of the population in a latent state with opinion B is $1-\beta-l(t)$. In Table 1, we summarize the notation used in our model.

Opinions A and B are associated with exponentially distributed latency periods whose mean duration is denoted by $\mu_A = 1/\lambda_A$ and $\mu_B = 1/\lambda_B$, respectively. λ_A and λ_B are the rate parameters of each distribution. Since we are interested in the relative duration of one latency period with respect to the other, we define the latency period duration ratio as $r = \mu_B/\mu_A = \lambda_A/\lambda_B$. Since for different values of λ_A and λ_B we can obtain the same ratio, we can set $\lambda_A = 1$ and $\lambda_B = \lambda$ with $\lambda \in (0, 1]$ to study the system's behavior without any loss of generality. With this notation, we can model any relation between the two latency periods. For example, if $\lambda = 1/2$ then $r = 2$, which means that the mean duration of the latency

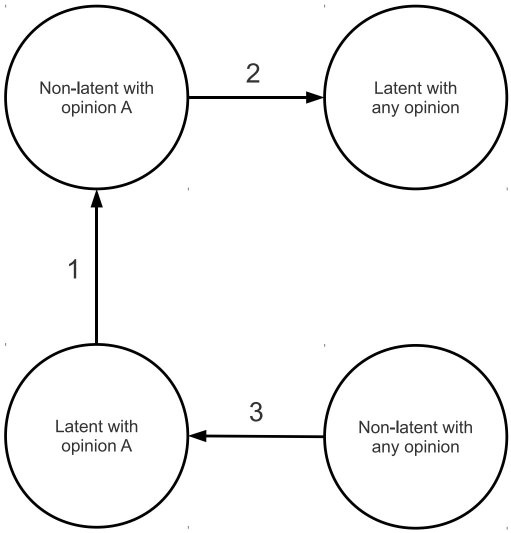

Fig. 4 Schematic structure of the ODE model. 1: Latent to non-latent with opinion A.
2: Non-latent with opinion A to latent with any opinion.
3: Non-latent with any opinion to latent with opinion A

period associated with opinion B is twice as long as the mean duration of the latency period associated with opinion A.

3.2.2 ODE model

Our ODE model lets us determine the fractions of the population that switch from a latent to a non-latent state, and vice versa. In Fig. 4, we show the schematic structure of our ODE model.

The first element of the ODE model is the size of the fraction of the population, already with opinion A, that switches from a latent to a non-latent state. The size of this fraction is $1 \cdot l(t)$, since the rate at which agents with opinion A become non-latent is equal to one. Due to the fact that no opinion change occurs, the loss of the latent fraction is a gain for the non-latent fraction (arrow 1 in Fig. 4).

The second element of the ODE model is the size of the fraction of the population in a non-latent state with opinion A that becomes latent with any opinion (arrow 2 in Fig. 4). As described before, randomly chosen non-latent agents form teams before switching to a latent state. Thus, the probability of choosing an agent with opinion A from the fraction of non-latent agents is $p(t) = n(t)/\beta$. Therefore, the fraction of non-latent agents with opinion A must decrease by a quantity equal to $p(t)f(t) = n(t)f(t)/\beta$, where $f(t)$ is the total fraction of the population that changes from a non-latent to a latent state. To compute $f(t)$, we need to recall that the fraction of the population in a non-latent state has a constant size β. Thus, the total fraction of the population that changes from a non-latent to a latent state must be equal in size to the fraction of the population that changes from a latent to a non-latent state. Such a fraction is given by

$$f(t) = 1 \cdot l(t) + \lambda \cdot \big(1 - \beta - l(t)\big), \tag{1}$$

where the first term is the fraction of the population that is in a latent state with opinion A multiplied by the rate at which agents with opinion A become non-latent. Similarly, the second term is the fraction of the population that is in a latent state with opinion B multiplied by λ, which is the rate at which agents with opinion B become non-latent.

The third element of our model is the size of the fraction of the population that changes from a non-latent state with any opinion to a latent state with opinion A (arrow 3 in Fig. 4). Becoming latent with opinion A can occur if agents are part of teams in which at least two agents favor opinion A. Such an event occurs only in the following four configurations of opinions in a randomly formed team: $A\ A\ B$, $A\ B\ A$, $B\ A\ A$, and $A\ A\ A$. Thus, the probability that an agent becomes latent with opinion A is given by $3p(t)^2(1-p(t)) + p(t)^3 = 3p(t)^2 - 2p(t)^3 = 3n(t)^2/\beta^2 - 2n(t)^3/\beta^3$. Therefore, the rate of change of the fraction of latent agents with opinion A must increase by a quantity equal to $[3n(t)^2/\beta^2 - 2n(t)^3/\beta^3]f(t)$.

Putting these three elements together, we can model the dynamics of the system with the following equations:

$$\begin{aligned} \frac{dn(t)}{dt} &= l(t) - \frac{n(t)}{\beta} f(t), \\ \frac{dl(t)}{dt} &= -l(t) + \left[3\frac{n(t)^2}{\beta^2} - 2\frac{n(t)^3}{\beta^3}\right] f(t). \end{aligned} \tag{2}$$

Figure 5(a) shows six example trajectories of this ODE model. The curves give the value of $n(t) + l(t)$, that is, the fraction of the population with opinion A. The values of $n(0)$ and $l(0)$ are determined as $n(0) = \beta F$ and $l(0) = (1-\beta)F$, where F is the initial fraction of the population with opinion A.

There are two important cases to be examined: (i) when $\lambda = 1$ and (ii) when $0 < \lambda < 1$. If $\lambda = 1$, or equivalently, when the latency periods have equal mean duration (dashed lines), and $F = 0.5$ (middle line) no dynamics takes place. By contrast, if $\lambda = 1$ and $F \neq 0.5$ the population reaches consensus on the opinion favored by the initial majority. For example, if the system starts with the majority favoring opinion A ($F = 0.52$, top line), the population reaches consensus on opinion A. On the contrary, if opinion A is favored by the minority ($F = 0.48$, bottom line), consensus is reached on opinion B. This result is similar to the one obtained by Krapivsky and Redner with their original model (see Sect. 2).

If $0 < \lambda < 1$, the system can reach consensus on an opinion that is initially favored by a minority of the agents. For instance, if $\lambda = 0.5$ (solid lines of Fig. 5(a)), that is, agents with opinion B remain latent for twice as long as agents with opinion A, it is sufficient that an initial fraction of only $F = 0.4$ favors opinion A to lead the system to consensus on this opinion (top line). However, there is a critical initial fraction that determines the opinion on which the system finally reaches consensus. If the initial fraction is lower than the critical initial fraction, the system converges to the opinion associated with the latency period with longer mean duration. For example, if $\lambda = 0.5$ and $F = 0.36$ (bottom line), the system reaches consensus on opinion B.

In Fig. 5(b), we compare the dynamics predicted by the ODE model (labeled as "ODE") with results obtained through Monte Carlo simulations (labeled as "MC Sim."). This figure shows the fraction of the population that favors opinion A as a function of the elapsed time t. The simulation results are obtained with 1000 agents over 100 independent runs. The bars indicate the interquartile range of the obtained simulation values. In the Monte Carlo simulation, all agents start non-latent with a fraction F in favor of opinion A. In the very first step, the desired number of teams is formed (agents are randomly selected from the non-latent agents, the majority rule is applied and the selected agents become latent). In the ODE model, this initial team formation step determines the model's initial conditions. Thus, for a given initial bias F, the values of $n(0)$ and $l(0)$ are calculated as $n(0) = \beta F$,

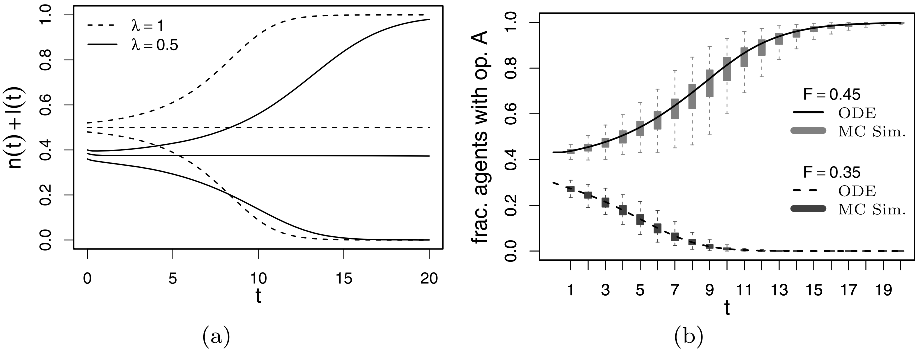

Fig. 5 Example trajectories of the model defined in (2). In all cases, $\beta = 0.25$. (**a**) shows the fraction of agents with opinion A over time as predicted by the ODE model. For $\lambda = 1$ (*dashed lines*) the system is initialized with $F = 0.52$ (*top*), $F = 0.5$ (*middle*), $F = 0.48$ (*bottom*). For $\lambda = 0.5$ (*solid lines*) the system is initialized with $F = 0.40$ (*top*), $F = 0.3845$ (*middle*), $F = 0.36$ (*bottom*). (**b**) shows a comparison of the ODE model with results of Monte Carlo simulations with 1000 agents over 100 independent runs. The *bars* indicate the interquartile range of the obtained simulation values. As predicted by the ODE model, when $F = 0.45$, the population reaches consensus on opinion A. By contrast, when $F = 0.35$, the population reaches consensus on opinion B

Fig. 6 Direction field of the system defined by (2) with isoclines (*dotted*) and example trajectories (*solid*). The *thick line* represents an unbiased system

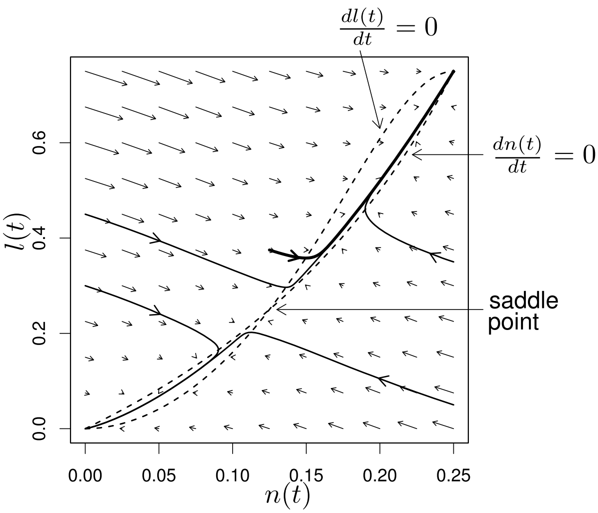

and $l(0) = (1 - \beta)(3F^2 - 2F^3)$, respectively. Our results show that the predictions of the ODE model fit the simulation results well. For $F = 0.35$ the fraction of the population with opinion A converges to zero. By contrast, if $F = 0.45$ the system reaches consensus on opinion A.

A more detailed analytical study of the proposed opinion formation model is presented by Scheidler (2011) who relaxes our assumption of a very large population and determines, using the Fokker–Planck equation, the probability of reaching consensus on the opinion associated with the shorter latency period as well as the time needed to reach consensus.

3.2.3 Stability

We investigate the stability of the equilibrium points of our ODE model. The stationary solutions of the system ($\frac{dn(t)}{dt} = 0$, $\frac{dl(t)}{dt} = 0$) are the states $[n(t) = \beta, l(t) = 1 - \beta]$, $[n(t) = 0, l(t) = 0]$, and $[n(t) = \beta/2, l(t) = \lambda(1 - \beta)/(1 + \lambda)]$. In Fig. 6, we show a visual example of the state space of the system when $\beta = 0.25$ and $\lambda = 0.5$. The arrows indicate how the system evolves at the specific point in the state space. The two dashed lines mark the two isoclines. These are the points at which $\frac{dn(t)}{dt} = 0$ and $\frac{dl(t)}{dt} = 0$, respectively. The two isoclines meet at the three stationary solutions $[n(t) = 0.25, l(t) = 0.75]$, $[n(t) = 0, l(t) = 0]$, and $[n(t) = 0.125, l(t) = 0.25]$. Four example trajectories for different starting conditions are depicted. The trajectories of the system that start at $[n(t) = 0.25, l(t) = 0.05]$ and $[n(t) = 0, l(t) = 0.3]$ both end up in $[n(t) = 0, l(t) = 0]$ (consensus on B), whereas for the starting conditions $[n(t) = 0.25, l(t) = 0.35]$ and $[n(t) = 0, l(t) = 0.45]$, the system converges to $[n(t) = 0.25, l(t) = 0.75]$ (consensus on A). The trajectory plotted with a thick line starts at $[n(t) = 0.125, l(t) = 0.375]$. This case represents an unbiased system, that is, when exactly half of the agents start with opinion A.

The solutions $[n(t) = \beta, l(t) = 1 - \beta]$, $[n(t) = 0, l(t) = 0]$ correspond to consensus on A and B, respectively. The Jacobian matrix evaluated at these two points results in $\begin{bmatrix} \frac{\beta-1}{\beta} & \lambda \\ 0 & -\lambda \end{bmatrix}$ and $\begin{bmatrix} \frac{\lambda(1-\beta)}{\beta} & 1 \\ 0 & -1 \end{bmatrix}$, respectively. The eigenvalues of these matrices are $(\beta - 1)/\beta$ and $-\lambda$ for the first matrix and $(\lambda\beta - \lambda)/\beta$ and -1 for the second matrix. In the range of variability of the values β and λ, $0 < \beta, \lambda < 1$, these eigenvalues are real and negative. Thus, both consensus states are asymptotically stable. The linearization of the system near the third equilibrium point has one positive and one negative eigenvalue (details are omitted because of excessive lengthiness). Thus, this third point is a saddle point and, hence, it is not stable. Consequently, the only stable equilibrium points of the system are $[n(t) = \beta, l(t) = 1 - \beta]$ and $[n(t) = 0, l(t) = 0]$. Thus, we can conclude that with large populations and exponentially distributed latency periods the opinion dynamics induced by the majority-rule opinion formation model with differential latency always lead to consensus.

3.3 Simulation study

In the ODE model, we made some simplifying assumptions that are not likely to hold in a real world robotics scenario. Specifically, in a real swarm robotics scenario, the number of robots is not very large and the duration of the latency periods is not likely to be exponentially distributed. Thus, we now relax these assumptions and study the system's dynamics using Monte Carlo simulation. In our simulation study, we assume that the duration of latency periods is normally distributed. Thus, we model situations in which robot actions have a mean duration with a symmetrical deviation. In Sect. 4, we further relax this assumption and test the system in a scenario in which the duration of latency periods is the result of robots interacting with other robots and with their environment.

Our simulation study is performed in three steps. First, we explore the effects of different parameters on the system's dynamics. In particular, we focus on the effects of different durations of the latency periods associated with each opinion and of the number of teams. Next, we study the system's dynamics when the number of teams is equal, or very close to the limit $N/3$, where N is the total number of agents. Such a case is of interest because in a continuous time system the probability of two teams becoming non-latent at exactly the same time is zero (in our discrete time simulation, this effect is implemented by not allowing teams that happen to finish at the same time to exchange team members). As a result, new team formations do not change the opinions of the agents involved. Thus, we expect that when $k = N/3$ the system will behave differently than when $k < N/3$. Finally, we study the system's dynamics when the distributions of the latency periods have various degrees of overlap.

3.3.1 Setup

The latency periods associated with opinions A and B are modeled as two normally distributed random variables with means μ_A and μ_B, and standard deviations σ_A and σ_B, respectively. The latency period duration ratio is defined as before: $r = \mu_B/\mu_A$. In our simulations, we used populations of $N \in \{9, 90, 900\}$ agents.[2] For each population size, we vary the number of teams: $k \in \{1, 2, 3\}$, when $N = 9$, $k \in \{1, 10, 20, 30\}$, when $N = 90$, and $k \in \{1, 100, 200, 300\}$, when $N = 900$. We also vary r by changing the value of μ_B. The explored values of r are 1, 2, 3, and 4. The reference mean, μ_A, is fixed to a value of 100 time steps. We set $\sigma_A = \sigma_B = 20$ time steps. With these settings, the two distributions do not significantly overlap. In Sect. 3.3.3, we study the system's dynamics when the distributions of the latency periods significantly overlap.

3.3.2 Dynamics

Figure 7 shows the dynamics of the proposed model with a population of 900 agents. The relation between the initial configuration of the population and the probability of reaching consensus on one of the alternative opinions follows the same nonlinear pattern observed in Fig. 2(a). However, when latency periods have a different mean duration, it is more likely that the system achieves consensus on the opinion associated with the shorter latency period. This fact is reflected by a lower critical initial fraction. In Fig. 7(a), for example, the critical initial fraction is approximately equal to 0.35 when $r = 4$, while it is approximately equal to 0.42 when $r = 2$. In every case, the peak on the number of team formations needed to reach consensus occurs at the critical initial fraction (see Fig. 7(b)). Additionally, at this critical point, the larger the latency period duration ratio, the more team formations are needed to reach consensus.

A second aspect that we study in this experiment is the effect of the number of teams on the system's dynamics. An example of the obtained results is shown in Figs. 7(c) and 7(d). For a latency period duration ratio greater than one, increasing the number of teams reduces the critical initial fraction. In terms of the number of team formations to achieve consensus, the results are similar to the ones observed in Fig. 2, that is, the maximum number of team formations occurs at the critical initial fraction. As expected, when k approaches $N/3$, the system exhibits different dynamics and stops following the aforementioned tendencies. Except for cases in which consensus is reached after the first team formations (e.g., with very small populations and very low or large initial densities), when $N = 3k$ the system does not reach consensus (curve plotted in gray in Fig. 7(c)). The reason for this result is the following. When $N = 3k$, every time a team is destroyed and formed anew, it is composed

[2] Due to space constraints, we present only the results obtained with 900 agents. Where appropriate, the results obtained with smaller populations are mentioned. The complete set of results as well as other supplemental material can be found in Montes de Oca et al. (2011).

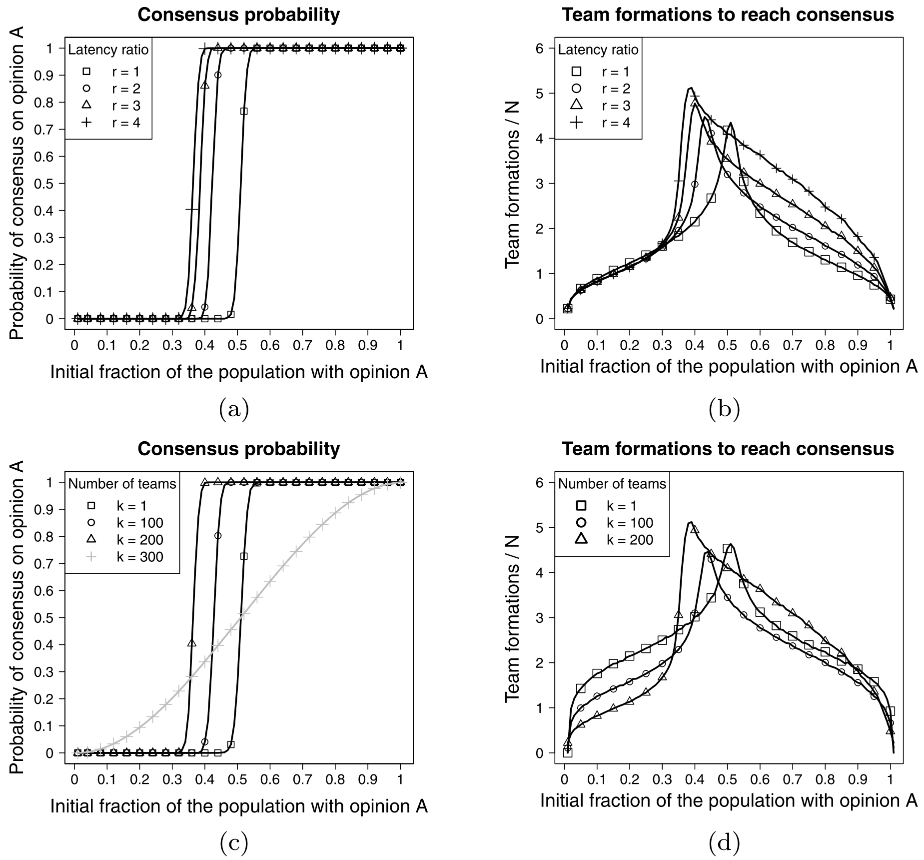

Fig. 7 Dynamics of the majority-rule opinion formation model with normally distributed latency periods on a population of 900 agents. (**a**) and (**b**) respectively show the probability of reaching consensus on opinion A and the number of team formations per agent necessary to reach consensus for different latency period duration ratios and for a fixed number of teams ($k = 200$). (**c**) and (**d**) respectively show the probability of reaching consensus on opinion A and the number of team formations to reach consensus for different number of teams and for a fixed latency period duration ratio ($r = 4$). The *gray curve* in (**c**) shows the case $k = N/3$ in which the system does not reach consensus (in this case, the data plotted is the fraction of the population with opinion A after 100,000 time steps). Results obtained through 1,000 independent runs of a Monte Carlo simulation

of exactly the same members because the probability of two teams leaving the latent state at exactly the same time is zero. This means that when $N = 3k$ there is no change in the number of agents with one or another opinion after the initial team formations. We have also tested the case in which $N = 3k + 1$ and, as in the previous case, consensus is not reached. In this case, this phenomenon occurs because three of the four non-latent agents available at the moment of forming a new team have the same opinion. Thus, while there may be a different agent in a new team, the team's opinion does not change, eliminating the possibility of an eventual consensus. When $N = 3k + 2$ the population always reaches consensus. Two non-latent agents are enough for possibly changing the opinion of one agent that just switched from a latent to a non-latent state. Thus, a non-latent population of at least two non-latent agents guarantees consensus.

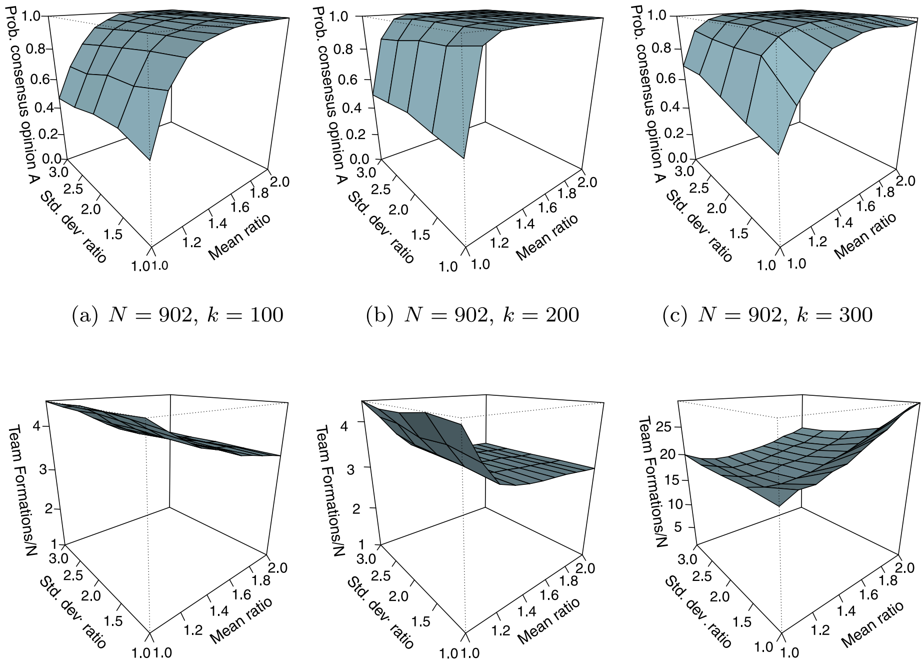

(a) $N = 902$, $k = 100$ (b) $N = 902$, $k = 200$ (c) $N = 902$, $k = 300$

(d) $N = 902$, $k = 100$ (e) $N = 902$, $k = 200$ (f) $N = 902$, $k = 300$

Fig. 8 The probability of reaching consensus on the opinion associated with the shorter latency period and the average number of team formations needed to do it as a function of different levels of overlap between latency period duration distributions. Results obtained through 1,000 independent runs of a Monte Carlo simulation

3.3.3 Dynamics of the system when latency distributions significantly overlap

If the distributions of the duration of latency periods significantly overlap, we expect that the population of agents will not be able to consistently reach consensus on the opinion associated with the shorter latency period. Thus, it is important to assess the ability of the system to discriminate between the two distributions if the system's dynamics are to be used as a decision-making mechanism.

The following experiment is aimed at measuring the extent to which the population can still reach consensus on the opinion associated with the shorter latency period when the two latency duration distributions overlap. We assume that there is no a priori information about which opinion is associated with the shorter latency period. Thus, the initial fraction of agents in favor of one opinion or the other is equal to 0.5. We fix the parameters of the distribution associated with the shorter latency period (μ_A, σ_A). We vary both the mean and the standard deviation of the distribution associated with the longer latency period (μ_B, σ_B). The explored ranges are: $\mu_B = r\mu_A$ with $r \in [1.0, 2.0]$ in increments of 0.1, and $\sigma_B = s\sigma_A$ with $s \in [1.0, 3.0]$ in increments of 0.5. The parameters used for the distribution associated with the shorter latency period are $\mu_A = 100$, and $\sigma_A = 10$. Other values were explored, but the system does not exhibit different dynamics as long as the relations between the distributions' coefficients of variation remain the same. As discussed in Sect. 3.3.2, two extra non-latent agents are needed to ensure consensus. Thus, in these experiments, we increase

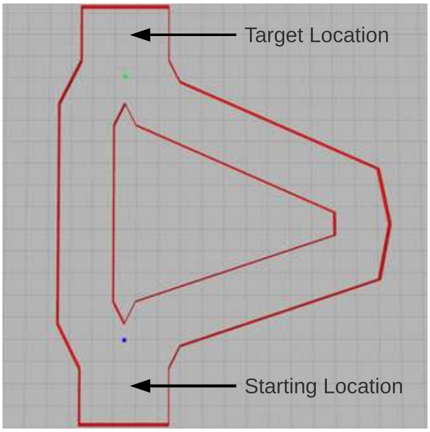

(a) Environment

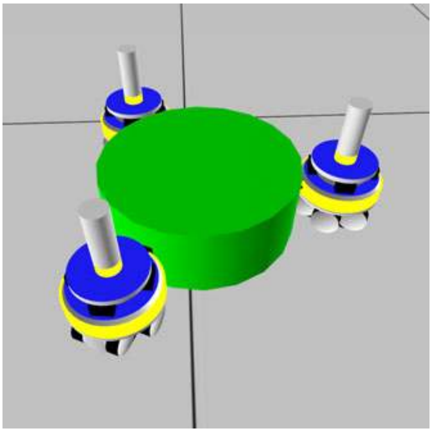

(b) Team

Fig. 9 Task Scenario. The environment, shown in (**a**), is a bridge-like arena with two branches of different lengths. In (**b**), we show a team of robots attached to an object. The task of the robots is to transport objects from the starting to the target location of the arena. The choice robots must make is to take either the longer or the shorter path

the population size with respect to the previous experiments. The results obtained with 902 agents are shown in Fig. 8.

The probability of reaching consensus on the opinion associated with the shorter latency period grows more rapidly when a large number of teams and, consequently, a large population is used. For example, with 11 agents the system has great difficulties in detecting the opinion associated with the shorter latency period (results shown in Montes de Oca et al. 2011). With 11 agents, the maximum probability of reaching consensus on the opinion associated with the shorter latency period is approximately 0.8. In contrast, in the example shown in Fig. 8, the system is able to discriminate latency periods under a wide range of combinations of means and standard deviation ratios. With 100 and 200 teams (Figs. 8(a) and 8(b)), the system is mostly affected by the ratio between means. When using 200 teams, the system reaches a probability of 1 for achieving consensus on the opinion associated with the shorter latency period already from $r \geq 1.3$. At the same time, the number of team formations needed to reach consensus decreases as r increases (Figs. 8(d) and 8(e)). With 300 teams (Fig. 8(c)), the system exhibits a good discrimination ability (although not as good as with 200 teams) but at a much higher cost in terms of team formations (Fig. 8(f)).

Irrespective of the size of the population, the standard deviation ratio does not have a significant impact on the probability of the system discriminating between the two distributions. We believe that this is the result of an "averaging" effect due to the large number of team formations needed to reach a consensus. The effects of short-term fluctuations due to the high variability of one of the distributions become negligible in the long run.

4 Self-organized collective decision-making in swarms of robots

In this section, we use a physics-based simulator to validate the predictions of the opinion formation model described in Sect. 3. We use a scenario that resembles the well-known double bridge experiment designed by Goss et al. (1989) (see Fig. 9(a)). The task of the robots is to transport objects from a starting location (at the bottom of the figure) to a target location (at the top of the figure). The objects that need to be transported weigh more than what a single robot can carry. Thus, robots need to team up in order to move the objects. An assembled team ready to transport an object is shown in Fig. 9(b). While performing this task, robots must choose to take either the short or the long path to reach the target location. These two options represent the robots' "opinions." The time needed by robots to go from the starting location to the target location and back is the duration of the latency period associated with the chosen path, and is called traversal time. Hence, in this example, latency periods arise naturally as a result of the robots' embodiment and the interactions of the robots with their environment. Moreover, the physical dimensions of the environment determine the maximum number of teams that can be used to perform the task. It is important to keep in mind that in this experiment active robots are those that traverse a path, and thus are latent with respect to the decision-making process because they can neither change opinion nor influence other robots to do so. Similarly, inactive robots are those that wait in the starting location, and thus are non-latent because they can potentially form new teams, and change or spread their opinion. Like the ants in Goss et al.'s experiment, robots do not have any knowledge about the length of the paths and do not measure distances or travel times.

4.1 Experimental setup

To carry out our experiments, we use ARGoS (Pinciroli et al. 2011), a simulator developed as part of the *SWARMANOID* project.[3] ARGoS accurately simulates physical interactions between robots and their environment. The robot models are based on the physical and electronic designs of the actual *SWARMANOID* foot-bots (Bonani et al. 2010).

In our simulations, only active robots are placed in the environment. The dimensions of the environment do not allow a parallel deployment of teams. Thus, a sequential deployment strategy is adopted. Every time a new team needs to be formed, three robots are chosen at random from the set of inactive robots. Note that since only active robots are in the environment, the process that ensures that team formation occurs with randomly picked robots is not physically simulated. The three chosen robots are placed in the starting location together with the object to be carried. These robots attach to the object using their gripper actuators. Next, the robots determine the team's majority opinion by exchanging messages using their range and bearing communication devices which allow robots to communicate locally with other robots (Roberts et al. 2009). Only robots that are located within a short range and that are in line of sight receive messages. Each robot sends its own opinion to the other two robots of the team, and once a robot receives the opinions of the others, it applies the local majority rule to determine the opinion to adopt. Upon agreement on the path to follow, robots start moving. Two LEDs are placed at the bifurcations to let robot teams know when and in which direction they should turn. Robots detect LEDs using their omnidirectional camera. When robots reach the goal area, they detach from the object they were transporting and go back, as single robots, through the same path they used when they were part of a team. On their way to the target location, robots use the collective transport controller designed by Ferrante et al. (2011). This controller allows robots to transport the object to the target location while avoiding obstacles (walls and single robots that are on their way back to the starting location). Obstacles are detected using a rotating distance scanner. The target location is indicated by a light source located above it, which the robots perceive through their light sensors. To go back to the starting location, robots use the light

[3] http://www.swarmanoid.org/

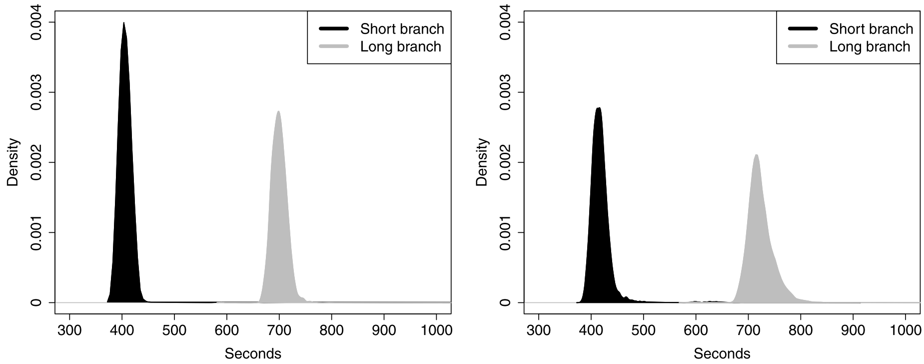

(a) Two teams (b) Ten teams

Fig. 10 Estimated traversal time distributions for the two paths when there are two (**a**) and ten (**b**) teams in the environment. Each density plot is based on 10,000 round trips (100 runs of 100 trips each) of a robot between the starting and goal locations in the environment shown in Fig. 9(a)

source that identifies the target location as a landmark and move away from it. To coordinate the heading direction, robots again use the range and bearing device as described in Ferrante et al. (2011). New teams are deployed every 40 simulated seconds until a specific number of teams is reached, or the environment reaches its maximum capacity, which is determined by the maximum number of teams that the shorter branch can hold. In our experiments, the shorter branch can hold up to ten teams.

4.2 Estimation of the traversal time distributions

The traversal time distributions of robots moving in the environment shown in Fig. 9 are related to the latency duration distributions discussed in Sect. 3. Although these distributions are related, there are some important differences. In the analysis presented in Sect. 3, we assumed that the distributions of the latency periods are independent of the number of agents with a particular opinion. However, in a swarm robotics scenario, this assumption does not generally hold because interference between robots is likely to dynamically change the latency duration distributions and their ratio. In our environment, for instance, a branch could become congested if many robots choose it. This increased congestion translates into longer and more variable path traversal times. To measure the effects of interference in our environment, we deploy from two to ten robot teams and make them traverse several times the environment using only one of the two branches. The estimated traversal time distributions when there are two and ten teams in the environment are shown in Fig. 10.

The results of this experiment show that both the mean and the standard deviation of the traversal time distributions change as a result of the number of teams that choose each branch. When there are only two teams in the environment, the average time needed to traverse the short and long branches of the environment is 408.5 and 699.8 seconds, respectively. Similarly, the standard deviation is 59.3 seconds for the short path and 15.7 seconds for the long path. When there are ten teams, the average time needed to traverse the short and long branches of the environment becomes 419.2 and 724.0 seconds, respectively. The standard deviation in this case becomes 29.0 seconds for the short path and 35.9 seconds for the long path. In our simulations with two agents choosing the short path, there were

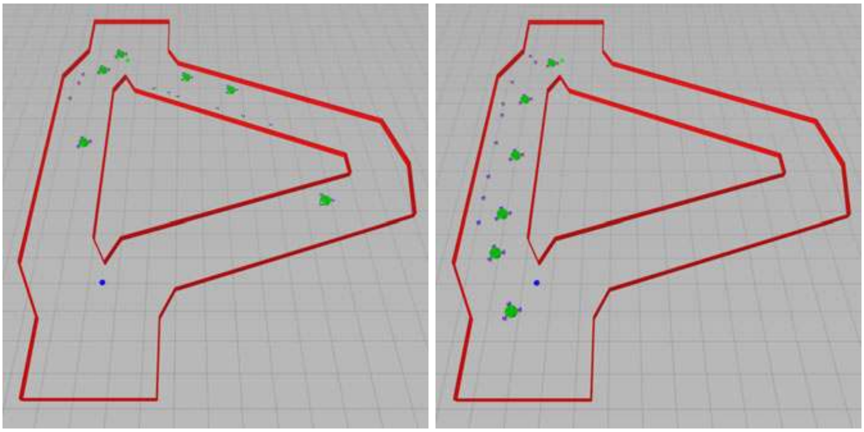

(a) In progress (b) Consensus

Fig. 11 Path selection process. (**a**) shows a swarm of robots in the process of transporting objects from the starting location to the target location. Note that the robots use both branches of the environment. (**b**) shows the state of the environment when the swarm of robots has reached consensus. The shorter path is selected by the swarm of robots. A video visualization of the simulation is available as online supplementary material

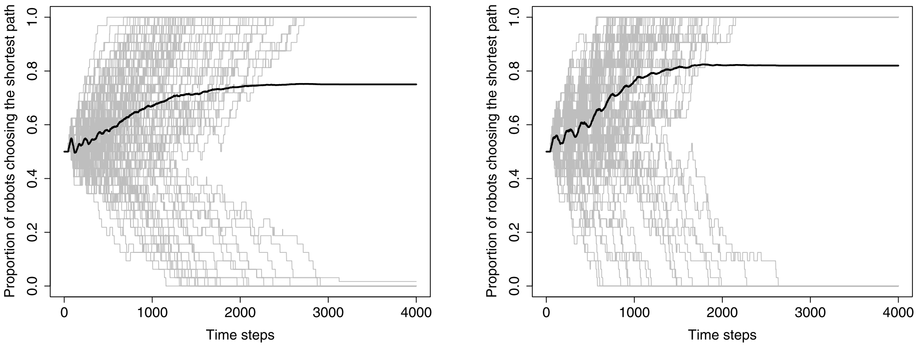

(a) Five teams (b) Ten teams

Fig. 12 Fraction of the population of robots with the opinion associated with the shorter path. The evolution of the system with five and ten teams is shown in (**a**) and (**b**), respectively. Single run results are shown in *gray lines*. The average over 100 runs is marked with a *thick black line*

a few rare cases in which the time needed by a robot to perform a round trip between the starting and target locations was very long. These outliers explain the high standard deviation observed. We also observe that the traversal time distributions are right-skewed. This is because robots that have reached the target location have to avoid collisions with incoming teams. This phenomenon occurs more frequently when the number of teams in the environment increases.

4.3 Collective decision-making

We now test the ability of a swarm of robots to choose the shorter of the two paths that connect the starting location with the target location in the environment shown in Fig. 9(a). In

Table 2 Probability of choosing the shorter branch of the environment as a function of the number of teams k. The population size N is equal to 32 robots. The highest probability is highlighted in boldface. These results are based on statistics taken from 100 independent simulations

k	Physics-Based Simulation		Monte Carlo Simulation	
	Probability	Avg. Team Formations	Probability	Avg. Team Formations
1	0.48	74.29	0.54	70.66
2	0.52	72.67	0.62	74.62
3	0.69	72.75	0.58	74.39
4	0.71	70.28	0.68	71.87
5	0.75	71.60	0.74	70.17
6	0.74	75.22	0.72	71.18
7	0.79	76.20	0.83	80.84
8	**0.86**	77.73	0.82	85.58
9	0.83	81.29	**0.86**	98.43
10	0.81	109.95	0.69	248.25

this experiment, the robots' decisions are governed by the dynamics of the model described in Sect. 3. We use a total of 32 robots (30 of which are executing the task at the same time plus two extra robots that are used in order to ensure consensus). Initially, 16 robots favor the shorter path and 16 favor the longer path. In Fig. 11, we show two snapshots of a simulation that finishes with the swarm selecting the shorter path. The reader can find a video that shows the system in action in the online supplementary material.

In Fig. 12, we show two examples of the development over time of the fraction of robots with the opinion associated with the shorter path. Consensus is the final state of all individual runs; however, notice that the swarm does not reach consensus on the shorter path in all runs. The probability of reaching consensus on the shorter path depends on the number of teams deployed. In Table 2, we list the estimated probabilities of reaching consensus on the shorter path. We also include results obtained with the Monte Carlo simulator used in Sect. 3.3 for validation purposes. The simulation setup uses the data gathered in the experiment described in Sect. 4.2. Specifically, we set the mean and standard deviation of the latency period associated with the shorter path to 100 and 20 time steps, respectively. The mean of the latency period associated with the longer path is set to $\frac{724}{408.5} \times 100 = 1.72 \times 100 = 172$ time steps, and its standard deviation is set to $\lceil \frac{35.9}{29} \times 20 \rceil = \lceil 1.23 \times 20 \rceil = 25$ time steps.

The probability of choosing the shorter path increases with the number of teams and reaches its maximum value with eight teams with the physics-based simulator and with nine teams with Monte Carlo simulations. In both cases, the maximum probability is 0.86. The average team formations needed to reach consensus oscillates within the range $[70, 75]$ for most cases and grows when the number of teams approaches the limit $N/3$, where N is the number of robots. Based on these results, we can make the following observations. First, small swarms (a 32-robot swarm can be considered small) have difficulties in discriminating latency distributions whose ratio is lower than two (see also the results in Montes de Oca et al. 2011). Second, as the number of teams approaches the limit $N/3$, the size of the non-latent subpopulation starts playing a role in both the quality of the decision eventually made by the swarm (lowering its quality) and the time it takes to reach consensus (increasing the number of needed team formations).

5 Related work

5.1 Models

In the biological sciences, self-organization models have been proposed to explain the coordination of large groups of animals (Camazine et al. 2001; Couzin and Krause 2003). Self-organization is itself the result of the interaction of several elements that include multiple direct or indirect interactions among the system's components, positive and negative feedback, and random fluctuations (Camazine et al. 2001). These models are particularly relevant for our proposal because the dynamics of the model described in Sect. 3 can be seen as an example of self-organization. In fact, the double-bridge experiment proposed by Goss et al. (1989) is reproduced here with the goal of pinpointing the self-organized nature of the collective decision-making mechanism introduced in this paper. Positive feedback, for example, occurs in experiments with real ants through the reinforcement of pheromone trails. In our experiment, positive feedback occurs because robots that use the shorter path return to the starting location before the robots that use the longer path. Thus, the probability that a new team has a majority in favor of the shorter path increases. Pheromone evaporation is a negative feedback process in path finding with real ants. In our experiment, negative feedback is the result of the increased difficulty with which teams that adopt the opinion associated with the longer path are formed. In both cases, real ants and robots, randomness plays an important role to break symmetries and produces the fluctuations that are amplified by the processes described above.

5.2 Collective decision-making in artificial swarms

Many collective decision-making mechanisms in swarm robotics are based on the simulation of pheromones. Approaches range from the use of real chemicals (Russell 1999; Fujisawa et al. 2008a, 2008b), to message passing between robots (Campo et al. 2010b), to the use of digital video projectors to cast images of pheromone trails on the ground (Sugawara et al. 2004; Garnier et al. 2007; Hamman et al. 2007). There are also works in which the environment is enhanced in order to let it store information. For example, Mamei and Zambonelli (2005) and Herianto and Kurabayashi (2009) deploy RFID tags in the environment so that robots can read from or write in them. Mayet et al. (2010) use an environment whose floor is covered with a paint that glows if robots activate ultraviolet LEDs. Another variant of the pheromone-inspired approach is to use actual robots as markers to form trails. Some works that use this approach are the ones by Werger and Matarić (1996); Payton et al. (2001); Nouyan et al. (2008 2009) and Ducatelle et al. (2010). Simulating pheromones, at least in the way it has been done so far, has important limitations. For example, dealing with chemicals is problematic because very specialized sensors are needed. The level of sophistication is such that some authors have even used real insect antennae (Kuwana et al. 1995; Nagasawa et al. 1999). Using video projectors is an approach that can be adopted only indoors and under controlled conditions. Furthermore, the use of video projectors implies the use of tracking cameras and a central computer to generate the images to be projected. The existence of such a central information processing unit gives the approach a single point of failure. Modifying the environment with special floors or with RFID tags is a cheap and interesting approach. However, its applicability is limited to situations in which it is possible to design and build an environment where it is known that robots are going to be deployed. Finally, using robots as markers allows a swarm to operate in unknown environments and no central control is required. However, complex robot controllers are needed in order to allow individual robots to play different roles in the swarm. While a promising approach, the development of complex robot control software for swarms is in its infancy as we are still trying to understand the connection between individual-level and collective-level behaviors.

Other insect behaviors have also served as inspiration sources. For example, trophallaxis, the exchange of liquid food between insects, was first used in swarm robotics by Schmickl and Crailsheim (2008) to generate gradients through robot-to-robot communication to allow robots to find the shortest path between two locations. Gutiérrez et al. (2010) also used trophallaxis as an inspiration source for a method through which a swarm of robots can locate and navigate to the closest location of interest from a particular origin. Both of these methods need robots to implicitly know that the goal is to find the shortest path between two locations. In Schmickl and Crailsheim's work, robots decrease a numerical value at a certain rate as they move. This value is communicated when there are encounters with other robots. Thus, the exchanged information gives a rough indication of the distance traveled. In Gutiérrez et al.'s work, robots actually measure the distance they have traveled and communicate this information to other robots in order to reduce the uncertainty of each robot's estimate of the location of a target. In our work, robots measure neither travel times nor distances. Nevertheless, the swarm finds the shortest path between the two locations. The aggregation behavior of cockroaches has been the source of inspiration for a site-selection mechanism with robots (Garnier et al. 2009; Campo et al. 2010a). The nest-selection mechanism used by ants, which is based on detecting a quorum in favor of one option, has inspired the work of Parker and Zhang (2009, 2010). In these works, robots need to know whether there are enough committed robots to one of the competing options. In both cases, the more robots are committed to one of the options, the more likely it is for a single robot to commit to that option too. In Garnier et al.'s work, the decision is probabilistic; and in Parker and Zhang's work, the decision depends on whether the number of committed robots is larger than a threshold. Choosing the value of this threshold or the rate at which the commitment probability increases is a critical issue because the first alternative that is identified as dominant will be the alternative chosen by the swarm. In our work, there are no thresholds or probabilities that depend on the number of robots with one opinion or the other. Thus, decision-making is a continuous process that ends when the whole population reaches a consensus.

Finally, robots in the work of Wessnitzer and Melhuish (2003) use the majority rule to decide which of two "prey" to chase and immobilize. Robots capture one prey after the other. Although the decision is collective, the majority rule is used simply to break the symmetry of the decision problem.

6 Conclusions and future work

In this paper, we have introduced a self-organized collective decision-making mechanism for swarm robotics. The settings for which the proposed mechanism can be useful are those in which robots face a binary choice on actions that produce the same result but that may take different amounts of time to perform. The proposed mechanism makes a swarm of robots choose with high probability the action with the shorter execution time in a completely decentralized way.

An opinion formation model, based on the majority-rule opinion formation model originally proposed by Krapivsky and Redner (2003), was used to determine on which opinion consensus was reached and whether consensus was a stable state. The main modification that we introduced to Krapivsky and Redner's model is called *differential latency*, which models the fact that the different actions that the robots can perform may take different amounts of time to be completed. Our analysis demonstrated that the population of agents reaches with high probability a consensus on the opinion associated with the shorter latency period. We showed that this is the case when the duration of latency periods are exponentially and normally distributed. We also showed that if the probability of two teams of agents becoming non-latent at exactly the same time is zero, it is possible to guarantee consensus if there are always at least two non-latent agents in the population.

We used a collective transport task in an environment that resembles the well-known double bridge experiment proposed by Goss et al. (1989) to validate the predictions of the model. In our simulations, robots form teams to transport objects and the actions that robots need to choose from to perform the task are to take the shorter or the longer path from a starting location to a target location. We observed that the swarm of robots indeed chooses the shorter path with high probability. In our simulations, only active robots are placed in the environment. For simplicity, we did not simulate the physics involved in the process that ensures that team formation occurs with randomly picked robots. However, in an experiment with real robots any mechanism that allows robots to be well-mixed in the location where teams are formed (e.g., a random walk) could be used to obtain results similar to those reported in this paper.

We believe future work should focus on the implementation of the proposed mechanism with real robots as well as on its combination with other approaches. Future work should aim to:

- Devise collective decision-making mechanisms that work well with large as well as with small populations. The achievement of this goal would enable experimentation with real robots.
- Relax some of the assumptions made in this work, such as robots knowing the number of alternative actions, the decision points in the environment, and the association between opinions and actions. Also of interest would be relaxing the assumption of a static environment.
- Integrate opinion dynamics with task allocation methods in order to tackle problems for which consensus is a suboptimal solution.
- Extend the proposed mechanism to situations in which the desired collective decision is not based on action execution time but on other aspects. Translating these aspects into latencies of different duration would be a first approach toward a more general collective decision-making mechanism; however, other approaches should be explored.
- Extend the proposed mechanism to situations in which robots need to solve tasks individually and not as part of a team. Our own ongoing work is following this direction. Roughly speaking, the idea is to simulate team formations by making robots memorize observed actions and apply the majority rule on these observations.

To conclude, we believe that collective decision-making in swarms based on opinion formation models is a new and exciting research direction with the potential of cross-pollinating the fields of swarm robotics and statistical physics. On the one hand, the field of swarm robotics may greatly benefit from ideas and tools developed in the statistical physics literature. On the other hand, physicists may regard swarm robotics as a rich source of interesting problems waiting to be modeled and solved.

Acknowledgements This work was supported by the European Commission through the ERC Advanced Grant "E-SWARM: Engineering Swarm Intelligence Systems" (contract 246939). The information provided is the sole responsibility of the authors and does not reflect the European Commission's opinion. The European Commission is not responsible for any use that might be made of data appearing in this publication. Alexander Scheidler is supported by the Postdoc Programme of the German Academic Exchange Service (DAAD). Mauro Birattari and Marco Dorigo acknowledge support from the F.R.S.-FNRS of Belgium's French Community, of which they are a Research Associate and a Research Director, respectively.

References

Bonani, M., Longchamp, V., Magnenat, S., Rétornaz, P., Burnier, D., Roulet, G., Vaussard, F., Bleuler, H., & Mondada, F. (2010). The MarXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2010)* (pp. 4187–4193). Piscataway: IEEE Press.

Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-organization in biological systems*. Princeton: Princeton University Press.

Campo, A., Garnier, S., Dédriche, O., Zekkri, M., & Dorigo, M. (2010a). Self-organized discrimination of resources. *PLoS ONE*, *6*(5), e19888.

Campo, A., Gutiérrez, Á., Nouyan, S., Pinciroli, C., Longchamp, V., Garnier, S., & Dorigo, M. (2010b). Artificial pheromone for path selection by a foraging swarm of robots. *Biological Cybernetics*, *103*(5), 339–352.

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, *81*(2), 591–646.

Chakrabarti, B., Chakraborti, A., & Chatterjee, A. (Eds.) (2006). *Econophysics and sociophysics: trends and perspectives*. Weinheim: Wiley-VCH.

Couzin, I. D., & Krause, J. (2003). Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, *32*, 1–75.

Ducatelle, F., Di Caro, G., & Gambardella, L. M. (2010). Cooperative self-organization in a heterogeneous swarm robotic system. In *Proceedings of the genetic and evolutionary computation conference (GECCO 2010)* (pp. 87–94). New York: ACM.

Ferrante, E., Brambilla, M., Birattari, M., & Dorigo, M. Socially-mediated negotiation for obstacle avoidance in collective transport. In *Proceedings of the 10th international symposium on distributed autonomous robotic systems (DARS 2010)*. Berlin: Springer (2011, to appear).

Fujisawa, R., Dobata, S., Kubota, D., Imamura, H., & Matsuno, F. (2008a). Dependency by concentration of pheromone trail for multiple robots. In *LNCS*: *Vol. 5217. Proceedings of the international conference on ant colony optimization and swarm intelligence (ANTS 2008)* (pp. 283–290). Berlin: Springer.

Fujisawa, R., Imamura, H., Hashimoto, T., & Matsuno, F. (2008b). Communication using pheromone field for multiple robots. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2008)* (pp. 1391–1396). Piscataway: IEEE Press.

Galam, S. (1986). Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. *Journal of Mathematical Psychology*, *30*(4), 426–434.

Garnier, S., Tache, F., Combe, M., Grimal, A., & Theraulaz, G. (2007). Alice in pheromone land: An experimental setup for the study of ant-like robots. In *Proceedings of the IEEE swarm intelligence symposium (SIS 2007)* (pp. 37–44). Piscataway: IEEE Press.

Garnier, S., Gautrais, J., Asadpour, M., Jost, C., & Theraulaz, G. (2009). Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, *17*(2), 109–133.

Goss, S., Aron, S., Deneubourg, J.-L., & Pasteels, J. M. (1989). Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, *76*(12), 579–581.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, *83*(6), 1420–1443.

Gutiérrez, Á., Campo, A., Monasterio-Huelin, F., Magdalena, L., & Dorigo, M. (2010). Collective decision-making based on social odometry. *Neural Computing & Applications*, *19*(6), 807–823.

Hamman, H., Szymanski, M., & Worn, H. (2007). Orientation in a trail network by exploiting its geometry for swarm robotics. In *Proceedings of the IEEE swarm intelligence symposium (SIS 2007)* (pp. 310–315). Piscataway: IEEE Press.

Helbing, D. (2010). *Quantitative sociodynamics. Stochastic methods and models of social interaction processes* (2nd ed.) Berlin: Springer.

Herianto, H., & Kurabayashi, D. (2009). Realization of an artificial pheromone system in random data carriers using RFID tags for autonomous navigation. In *IEEE international conference on robotics and automation (ICRA 2009)* (pp. 2288–2293). Piscataway: IEEE Press.

Krapivsky, P. L., & Redner, S. (2003). Dynamics of majority rule in two-state interacting spin systems. *Physical Review Letters*, *90*(23), 238701.

Kuwana, Y., Shimoyama, I., & Miura, H. (1995). Steering control of a mobile robot using insect antennae. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 1995)* (Vol. 2, pp. 530–535), Piscataway: IEEE Press.

Lambiotte, R., Saramäki, J., & Blondel, V. D. (2009). Dynamics of latent voters. *Physical Review. E, Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, *79*(4), 046107.

Mamei, M., & Zambonelli, F. (2005). Physical deployment of digital pheromones through RFID technology. In *Proceedings of the international joint conference on autonomous agents and multiagent systems (AAMAS 2005)* (pp. 1353–1354). New York: ACM.

Mayet, R., Roberz, J., Schmickl, T., & Crailsheim, K. (2010). Antbots: A feasible visual emulation of pheromone trails for swarm robots. In *LNCS*: *Vol. 6234*. *Proceedings of the international conference on swarm intelligence (ANTS 2010)* (pp. 84–94). Berlin: Springer.

Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: Supplementary information webpage. http://iridia.ulb.ac.be/supp/IridiaSupp2010-014/.

Nagasawa, S., Kanzaki, R., & Shimoyama, I. (1999). Study of a small mobile robot that uses living insect antennae as pheromone sensors. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 1999)* (pp. 555–560). Piscataway: IEEE Press.

Nouyan, S., Campo, A., & Dorigo, M. (2008). Path formation in a robot swarm: self-organized strategies to find your way home. *Swarm Intelligence*, *2*(1), 1–23.

Nouyan, S., Gross, R., Bonani, M., Mondada, F., & Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Transactions on Evolutionary Computation*, *13*(4), 695–711.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, *314*(5805), 1560–1563.

Parker, C. A. C., & Zhang, H. (2009). Cooperative decision-making in decentralized multiple-robot systems: The best-of-N problem. *IEEE/ASME Transactions on Mechatronics*, *14*(2), 240–251.

Parker, C. A. C., & Zhang, H. (2010). Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, *4*(3), 199–220.

Payton, D., Daily, M., Estowski, R., Howard, M., & Lee, C. (2001). Pheromone robotics. *Autonomous Robots*, *11*(3), 319–324.

Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Stirling, T., Gutiérrez, Á., Gambardella, L. M., & Dorigo, M. (2011, in press). ARGoS: a modular, multi-engine simulator for heterogeneous swarm robotics. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2011)*. Piscataway: IEEE Press.

Roberts, J., Stirling, T., Zufferey, J., & Floreano, D. (2009). 2.5D infrared range and bearing system for collective robotics. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2009)* (pp. 3659–3664). Piscataway: IEEE Press.

Russell, R. A. (1999). Ant trails—an example for robots to follow? In *Proceedings of the IEEE international conference on robotics and automation (ICRA 1999)* (pp. 2698–2703). Piscataway: IEEE Press.

Scheidler, A. (2011). Dynamics of majority rule with differential latencies. *Physical Review. E, Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, *83*(3), 031116.

Schelling, T. (1978). *Micromotives and macrobehavior*. New York: Norton.

Schmickl, T., & Crailsheim, K. (2008). Trophallaxis within a robotic swarm: bio-inspired communication among robots in a swarm. *Autonomous Robots*, *25*(1–2), 171–188.

Sugawara, K., Kazama, T., & Watanabe, T. (2004). Foraging behavior of interacting robots with virtual pheromone. In *Proceedings of the IEEE/RSJ international conference on intelligent robots and systems (IROS 2004)* (pp. 3074–3079). Piscataway: IEEE Press.

Werger, B., & Matarić, M. (1996). Robotic "food" chains: Externalization of state and program for minimal-agent foraging. In *Proceedings of the international conference on simulation of adaptive behavior: from animals to animats (SAB 1996)* (pp. 625–634). Cambridge: MIT Press.

Wessnitzer, J., & Melhuish, C. (2003). Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In *LNCS*: *Vol. 2801*. *Proceedings of the European conference on artificial life (ECAL 2003)* (pp. 893–902). Berlin: Springer.