Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Opinion Dynamics for Decentralized Decision-Making in a Robot Swarm

Marco A. Montes de Oca, Eliseo Ferrante, Nithin Mathews, Mauro Birattari, and Marco Dorigo

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2010-004

March 2010
Last revision: May 2010

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle*
Université Libre de Bruxelles
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2010-004

Revision history:

TR/IRIDIA/2010-004.001	March 2010
TR/IRIDIA/2010-004.002	May 2010

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Opinion Dynamics for Decentralized Decision-Making in a Robot Swarm

Marco A. Montes de Oca mmontes@ulb.ac.be
Eliseo Ferrante eferrant@ulb.ac.be
Nithin Mathews nmathews@ulb.ac.be
Mauro Birattari mbiro@ulb.ac.be
Marco Dorigo mdorigo@ulb.ac.be

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium

May 18, 2010

Abstract

In this paper, we study how an opinion dynamics model can be the core of a collective decision-making mechanism for swarm robotics. Our main result is that when opinions represent action choices, the opinion associated with the action that is the fastest to execute spreads in the population. Moreover, the spread of the best choice happens even when only a minority is initially advocating for it. The key elements involved in this process are consensus building and positive feedback. A foraging task that involves collective transport is used to illustrate the potential of the proposed approach.

1 Introduction

Large groups of animals can exhibit behavioral patterns that resemble the ones observed in physical systems that are composed of many simple entities [4, 3]. This observation has led to the development of *opinion dynamics models*, which are used to study large-scale social, economic, and natural phenomena that involve many interacting agents [3].

In this paper, we study how an opinion dynamics model can be the core of a collective decision-making mechanism for swarm robotics. The opinion dynamics model used in our work, originally proposed by Krapivsky and Redner [8], allows a large population of agents to reach consensus on one of two alternatives. We are interested in situations where opinions represent actions that take time to perform. Our goal is to determine if, and under which conditions, the action selected by the swarm is the one that takes less time on average to perform. Such outcome would be useful in swarm robotics applications in which

the number of times a particular action is selected is correlated with the amount of work performed, thus maximizing the productivity of a robot swarm.

In Krapivsky and Redner's model the *majority rule* is repeatedly applied on teams of three agents. With the majority rule, team members adopt the opinion shared by the majority of the team members. In addition to the majority rule, we also use the *expert rule*, through which team members adopt the opinion of a single agent, called the expert, which is selected according to some problem-specific criterion. As these rules are at the opposite ends of the spectrum of all weighted majority rules [1], our study allows us to have an intuition of the results that would be obtained if a weighted majority rule was used.

The main result of our study is that the dynamics of the system makes the swarm select with high probability the action that is the fastest to execute. When using the expert rule, the fastest-to-execute action is selected by the swarm even when only a minority of the robots is initially in favor of it. The potential of the proposed approach as a decentralized decision-making mechanism for swarms of robots is illustrated through a foraging task that involves collective transport. As a result of the application of the proposed approach, the swarm of robots is able to select the shortest path without requiring the robots to measure path-travel times or to rely on pheromone-like information. Other applications where it is desired to select automatically the choice that increases the efficiency of a system composed of many interacting agents could benefit from the use of the mechanism presented in this paper.

The rest of the paper is structured as follows. In Section 2, we describe related work in the area of decentralized decision-making in swarms of robots. In Section 3, we describe the model and the decision rules used. In Section 4, we describe the task and the experimental setup used to evaluate the effectiveness of the proposed approach. Results are presented in Section 5. Conclusions and future work are given in Section 6.

2 Related Work

Many decentralized decision-making mechanisms have been inspired by the behavior of insects. For example, the pheromone-laying and pheromone-following behavior of some ant species [7] has inspired many works. Most of them have focused on the simulation of pheromones through the use of chemical substances [16], by projecting images on the ground [18], by deploying RFID tags [9], or by using robots as message-relay devices or as beacons so as to form robot chains [19, 14, 11, 12]. Recently, pheromones have been simulated by signals sent by robots that belong to a swarm different from the one engaged in the foraging task [5]. These mechanisms have some important disadvantages. For example, designing sensors for detecting chemicals reliably is a very difficult task. Using robots as beacons requires the development of complex controllers to allow an individual robot to play different roles both within and outside a robot chain. Using RFID tags requires the modification of the environment prior to the deployment of the swarm, which is impossible in some cases. Pro-

jected pheromones can be impractical because a central computer is needed. In contrast to all these works, our system does not require the explicit simulation of pheromones.

Other behaviors of insects have also been a source of inspiration. For example, trophallaxis, which is the insect-to-insect exchange of food, has been the inspiration for a distributed mechanism to create a gradient in the environment to help robots navigate [17]. The aggregation behavior exhibited by cockroaches has been the source of inspiration for a site-selection mechanism with robots [6]. The best-of-N selection mechanism proposed by Parker and Zhang [13] has been inspired by the nest-selection mechanism used by some species of ants. To the best of our knowledge, this last work is the most similar to ours. In both of them, a collective decision is the result of the competition between alternatives. The main difference, however, lies in the consensus building mechanism. In Parker and Zhang's approach, robots need to know whether there is a sufficient number of robots in favor of one alternative before committing to it. Robots do that through a quorum test that depends on a parameter that the designer needs to set before deployment. This is a critical issue because the first alternative that is identified as dominant through the quorum test will be the alternative chosen by the swarm. In our work, the collective decision is the result of self-organization.

Another work related to ours is the one of Wessnitzer and Melhuish [20]. In their work, robots have to chase and immobilize two "prey". Robots capture one prey after the other. To select which prey to immobilize, robots apply the local majority rule to break the symmetry of the decision problem and to make the population agree on one choice only. Our work goes a step further by considering the effects of implicit time costs in the robots' actions.

3 Opinion Dynamics and Decentralized Decision-Making

In this section, we describe the model and the decision rules used in our study. We also explain how this model becomes the core of a decentralized decision-making mechanism when opinions represent actions that take time to perform.

3.1 Opinion Dynamics Model and Decision Rules

We use the opinion dynamics model proposed by Krapivsky and Redner [8]. It operates on a population of N agents, each of which can be in one of two possible states, called opinions. The system evolves as follows: A team of three agents is picked at random without replacement from the population. Then, the individual opinions of the team members are aggregated and transformed into a team opinion by a decision rule. After this, all team members adjust their individual opinions to match the resulting team opinion. The team members are put back in the population and a new team is picked. The process is repeated until all agents share the same opinion.

(a) Majority rule (b) Expert rule

Figure 1: Example application of the majority and the expert rules on a team of three agents with different opinions (represented by different color shades). Figure (a) shows the outcome of the majority rule: the team adopts the opinion of the majority. Figure (b) shows the outcome of the expert rule: the team adopts the opinion of the expert (marked with a × symbol).

We use two decision rules: the majority rule (the one studied in [8]), and the expert rule. When the majority rule is used, all team members will assume the opinion that at least two team members share. With the expert rule, agents will assume the opinion of a single agent, called the expert (See Section 4.2 for information on the criterion used to choose which agent plays this role). Figure 1 shows an example of the application of the majority and of the expert rules.

3.2 Opinion Dynamics, Actions, and Robots

The opinion dynamics model described above can be used as the basis of a decentralized decision-making mechanism in a swarm of robots. Three elements need to be taken into account to do it: (i) opinions need to be interpreted as actions that robots execute, (ii) actions take time to perform, and (iii) robots can operate in parallel. These elements are modeled through a parallel version of the opinion dynamics model described in Section 3.1. Instead of picking at random one team of three agents, we pick k teams without replacement. Then, each team selects an opinion according to a decision rule and its members adopt that opinion. The actions associated with the adopted opinions are then executed. The execution time associated with an action is not necessarily the same from one execution to another due to unexpected events during execution (e.g., an obstacle may have to be avoided, or a robot skids while trying to move). When a team finishes executing an action, its members become available again to form another team. The new team cannot have as members robots that are at that moment executing an action. The process continues until the swarm reaches consensus, the time allocated for the task is over, or the demand for the task ceases to exist. Figure 2 shows an example of the process just described.

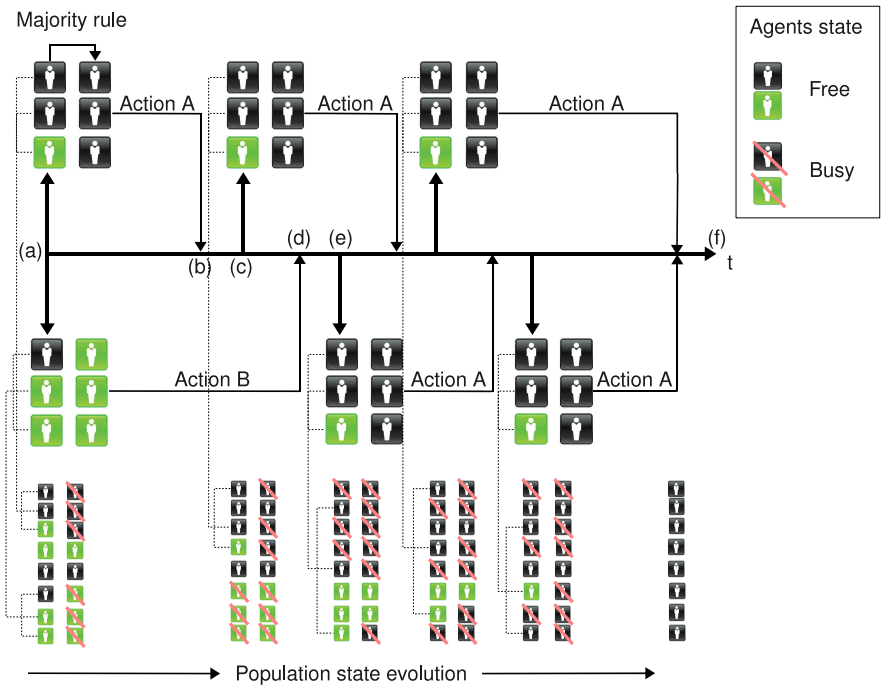

Figure 2: Example of the dynamics induced by the majority rule on a population of 8 agents with 2 teams of 3 agents each. The opinion represented in black is associated with action A, while the opinion represented in light color is associated with action B. Action A is faster to execute than action B, on average. This system evolves as follows. First, in point (a), two teams are formed at random and the majority rule is applied on each one of them. Each team then executes the selected action. In point (b), a team finishes. In point (c) a new team is formed from the set of free agents (busy agents are not considered when the selection occurs). The time it takes to form a team is represented by the distance between points (b) and (c). After the application of the majority rule in (c), the team performs the agreed action (action A, in this case). In point (d), the other team finishes and a new team is formed (point (e)). Again, the majority rule is applied once more to decide which action to perform (action A, again). The process continues until the population reaches consensus (point (f)). In this example, the population changes from a heterogeneous opinion state to a homogeneous one that corresponds to the fastest-to-execute action.

4 Evaluation Scenario and Setup

In this section, we describe the task, the simulation environment, and the experimental setup used to evaluate the effectiveness of the proposed approach.

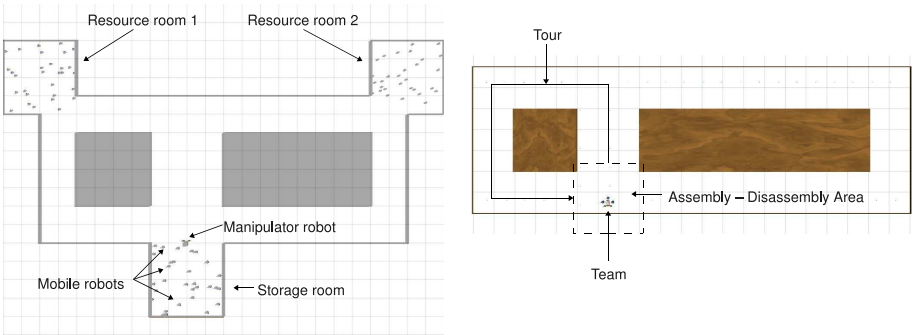

(a) Complete scenario (b) Simplified scenario

Figure 3: Test task. Figure (a) shows a complete view of the scenario. A swarm of robots must collect as many objects as possible within some time limit from the two resource rooms. Figure (b) shows a simplified version of the scenario. Robots must choose between making the left or right tour. The simplified scenario allows us to concentrate on the decision-making aspect of the task. See text for more information.

4.1 The Task

We chose a foraging task that involves collective transport as an example of the kind of applications the proposed approach could be used for. The environment consists of a storage room and two resource rooms in which there are objects of interest. There are two kinds of robots: robots that can manipulate the objects of interest but that cannot move by themselves, and robots that cannot manipulate the objects of interest but that can move autonomously and carry the manipulator robots. The task is to collect as many objects of interest as possible within some time limit from the two resource rooms, and deposit them in the storage room. To accomplish the task, the robots that can move autonomously must form small teams to carry the manipulator robots. These teams of robots go back and forth between the storage and resource rooms. It is in general desirable to go to the closest resource room in order to maximize the number of collected objects in a given amount of time. The path that leads to the closest resource room is a priori unknown to the robots. Figure 3(a) shows a picture of the complete envisioned scenario.

In the complete scenario, the choice that robots face is whether to turn left or right to go to a resource room. To focus only on the decision-making aspect of the task and not on other technicalities, such as object manipulation, collective obstacle avoidance, or self-assembly, we use a simplified version of the scenario described above, which is shown in Figure 3(b). Instead of a storage room there is an assembly-disassembly area where the mobile robots attach to and detach from the manipulator robots. Navigation is possible thanks to visual aids similar to the ones used in [12]. In this simplified scenario, the choice faced by the robots is simply to either turn left or right. These choices represent the

robots' opinions on which the decision-making mechanism will operate. The execution of a chosen action starts with the formation of a team of robots, it continues with the navigation around the big obstacle in the chosen direction, and finishes once a complete lap is made.

4.2 Setup

All the experiments reported in this paper were performed in simulation. The simulator that we used [15] was developed for the *SWARMANOID* project.[1] This simulator uses the Open Dynamics Engine library[2] to simulate accurately physical interactions with the environment and between robots. The robot models are based on the physical and electronic design of the actual *SWARMANOID* robots (currently under development). The mobile robot model is based on the *Foot-bot*, and the manipulator robot model is based on the *Hand-bot* [2].

We ran simulations with different swarm sizes ($N \in \{8, 16, 32, 64\}$), different initial opinion biases (according to a parameter $p \in \{0.05, 0.1, \ldots, 0.95\}$, which is the probability of a robot's initial opinion to advocate for the left path), different numbers of teams ($k \in \{2, 4, 8, 16\}$), and different action-execution time ratios ($r \in \{1, 2, 4\}$). The action-execution time ratio is defined as $r = l_{right}/l_{left}$, where l_{left}, and l_{right} are the length of the left and right paths, respectively. The reference length, l_{left}, was adjusted so that no collisions between teams of robots occurred. This was done by making the length of the left path much greater than the total length that results from lining up the k teams used in a given experiment. Teams had 3 robots each. A simulation was run until the swarm achieved consensus. 100 trials were run for each combination of parameters.

An extra parameter of the system is the criterion used to select the expert when applying the local expert rule. In our experiments, that criterion is the absolute number of completed laps (no distinction between paths is made), that is, the most "experienced" robot is selected as the expert. The study of other criteria to select the expert is left for future work.

5 Results

We are interested in two aspects of the system: (i) the probability with which the swarm selects the fastest-to-execute action as a function of the initial opinion bias, and (ii) the number of team formations needed to reach consensus on one of the alternative choices as a function of the initial opinion biases and of the swarm size. For both aspects, we evaluate the effects of different numbers of teams, k, and of different action-execution time ratios, r.

[1] http://www.swarmanoid.org/
[2] http://www.ode.org/

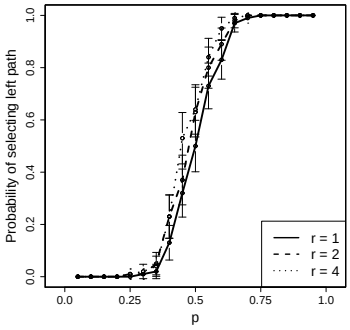

(a) Majority rule: $N = 64$, $k = 4$

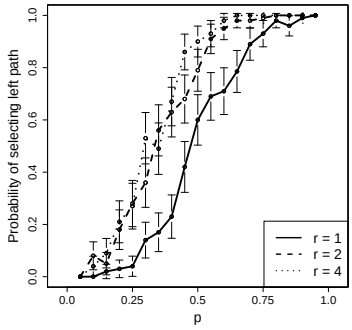

(b) Expert rule: $N = 64$, $k = 4$

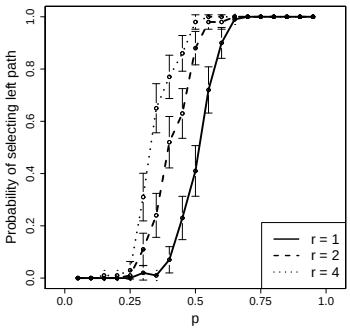

(c) Majority rule: $N = 64$, $k = 16$

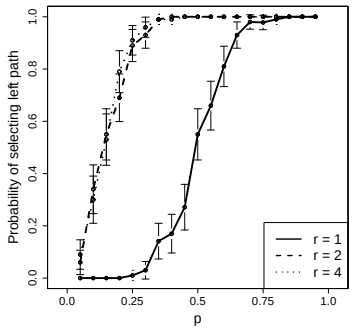

(d) Expert rule: $N = 64$, $k = 16$

Figure 4: Probability of selecting the fastest-to-execute action as a function of the initial opinion bias in a swarm of 64 robots with 4 and 16 teams. Figures (a) and (c) show the results obtained with the majority rule. Figures (b) and (d) show the results obtained with the expert rule. Error bars indicate the confidence interval at a 95% level.

5.1 Probability of Selecting the Fastest-to-execute Action

Figure 4 shows the estimated probability of selecting the left path, that is, the fastest-to-execute action, as a function of the initial opinion bias in a swarm of $N = 64$ robots.[3] We estimate this probability by dividing the number trials the system reached consensus on the left path by the total number of trials.

There is a nonlinear relationship between the initial opinion bias and the final probability with which the swarm chooses an alternative. In all cases, there is a critical bias p_c such that if $p < p_c$ the swarm will choose one opinion, and if $p > p_c$ the swarm will choose the other opinion. In the case where $r = 1$, that is, when there is no difference between the alternative choices, the critical initial bias is $p = 0.5$. When $r > 1$, $p_c < 0.5$ for both decision rules. With the

[3] We refer the reader to [10] for the complete set of results.

majority rule, the higher the action-execution time ratio, the lower the critical bias. Furthermore, the critical bias decreases as the number of teams active in the environment increases (the critical bias is lower when $k = 16$ than when $k = 4$). With the expert rule, the critical bias is, in general, lower than with the majority rule; however, the actual value depends more strongly on the number of teams than on the action-execution time ratio.

From a practical point of view, a small critical bias is desirable because it means that the action that is fastest to execute will be selected by the whole swarm even when only a minority of the agents is initially in favor of it. In this sense, the expert rule is the best suited for this purpose because it can spread more easily the opinion of the minority. For example, see Figure 4(d), when $r > 1$. In this case, the swarm of 64 robots chooses the left path with probability 0.8 when just 13 robots (20% of the swarm) initially choose it.

Small swarms ($N < 32$ with $k < 8$) have greater difficulties than large swarms in detecting differences in the average action-execution times (results shown in [10]). This may be due to the rapid opinion fluctuations that are amplified by the system. For instance, in an 8-robot swarm, one robot represents the 12.5% support for one or another opinion. Consensus is reached, but the chosen alternative is random.

5.2 Number of Team Formations Needed to Reach Consensus

The number of team formations needed to reach consensus (NTFC) depends on the initial opinion bias (see Figure 5). In all cases, the maximum NTFC is reached when the initial opinion bias is equal to the critical bias (see Section 5.1).

The NTFC using either the majority rule, or the expert rule, depends on the number of teams that are active in the environment (k) and the action-execution time ratio (r). With the majority rule, the NTFC increases as k and r increase. With the expert rule, the NTFC decreases as k and r increase. In conclusion, with the expert rule the opinion associated with the fastest action is spread more rapidly than with the majority rule, especially in the presence of large action-execution time ratios.

In case there is no a priori information about the quality of the alternatives the robots must choose from, the most reasonable strategy to initialize the system is to have a balanced initial opinion bias, that is, $p = 0.5$. Under that circumstance, one may ask what would be the NTFC as a function of the swarm size. The answer is shown in Figure 6. Not surprisingly, the NTFC increases with the swarm size in all cases. With the majority rule, there appears to be no significant difference in the NTFC if the number of active teams in the environment changes. With the expert rule, however, if more teams are deployed and the swarm size increases, the NTFC decreases.

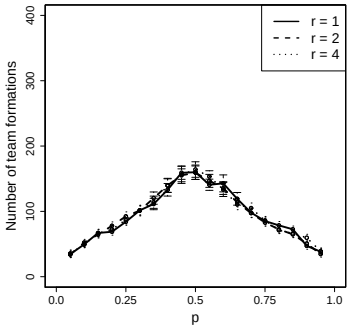

(a) Majority rule: $N = 64$, $k = 4$

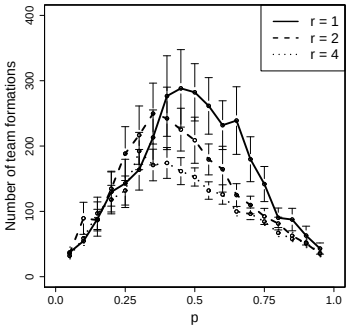

(b) Expert rule: $N = 64$, $k = 4$

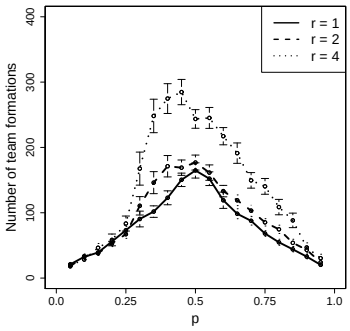

(c) Majority rule: $N = 64$, $k = 16$

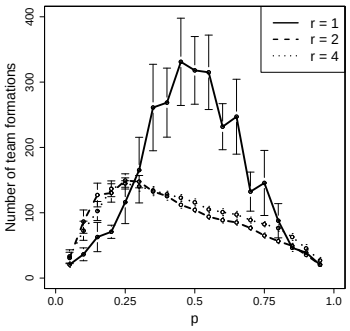

(d) Expert rule: $N = 64$, $k = 16$

Figure 5: Number of team formations needed to reach consensus as a function of the initial opinion bias. These results correspond to the case where $N = 64$. Figures (a) and (c) show the results obtained with the majority rule. Figures (b) and (d) show the results obtained with the expert rule. Error bars indicate the confidence interval at a 95% level.

6 Conclusions and Future Work

We presented a decentralized decision-making mechanism for swarm robotics systems whose main elements are: (i) a consensus-building mechanism based on an opinion dynamics model, and (ii) actions that take different average time to execute. The consensus-building mechanism can be used with any weighted majority rule. In this paper, we explored the two extremes: the simple majority rule and the expert rule. We showed how the dynamics of the system allows the swarm to choose with a high probability the action that is the fastest to execute. With the expert rule, the fastest-to-execute action is selected by the whole swarm even when only a minority of the agents is initially in favor of it.

The difference in the execution time of the alternative actions induces a positive feedback process that is ultimately the responsible for the swarm's final

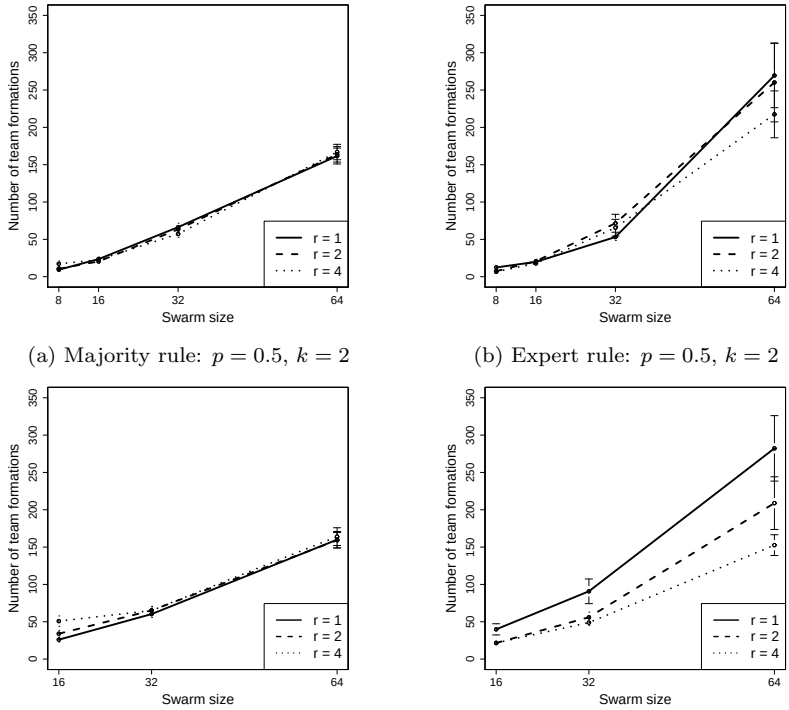

(a) Majority rule: $p = 0.5$, $k = 2$

(b) Expert rule: $p = 0.5$, $k = 2$

(c) Majority rule: $p = 0.5$, $k = 4$

(d) Expert rule: $p = 0.5$, $k = 4$

Figure 6: Number of team formations needed to reach consensus as a function of the swarm size. These results correspond to the case where $p = 0.5$. Figures (a) and (c) show the results obtained with the majority rule. Figure (b) and (d) show the results obtained with the expert rule. Error bars indicate the confidence interval at a 95% level.

choice. This phenomenon has parallels with the decision-making mechanism exploited by ants [7]. However, in our system agents influence each other by forming teams instead of doing it via pheromones.

From a practical point of view, forming small teams allows the number of messages between robots to be kept to a minimum. As a consequence, the system is able to scale up to many robots (in fact, the more robots, the better the system performs) without having to deal with interference or bandwidth problems.

Besides porting the system to real robots, future work includes extending the mechanism to sequential, adaptive and multi-choice decision-making. A first step in this direction will be to switch to a multidimensional probabilistic representation of a robot's opinion.

Acknowledgments.

This work was partially supported by the SWARMANOID project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-022888, and by the VIRTUAL SWARMANOID project funded by the Fund for Scientific Research F.R.S.-FNRS of Belgium's French Community. 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. M. Dorigo and M. Birattari acknowledge support from the F.R.S.-FNRS of Belgium's French Community, of which they are a research director and a research associate, respectively. We thank A. Brutschy and C. Pinciroli for their useful advice.

References

[1] Daniel Berend and Yuri Chernyavsky. Ranking of decision rules with random power distribution. *Mathematical and Computer Modelling*, 48(9–10):1326–1334, 2008.

[2] M. Bonani, S. Magnenat, P. Rétornaz, and F. Mondada. The hand-bot, a robot design for simultaneous climbing and manipulation. In *Proceedings of the International Conference on Intelligent Robotics and Applications (ICIRA 2009) LNCS 5928.*, pages 11 –22. Springer, Berlin, Germany, 2009.

[3] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009.

[4] B.K. Chakrabarti, A. Chakraborti, and A Chatterjee, editors. *Econophysics and sociophysics: Trends and perspectives.* Wiley-VCH, Wienheim, Germany, 2006.

[5] F. Ducatelle, G. Di Caro, and L. M. Gambardella. Cooperative self-organization in a heterogeneous swarm robotic system. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2010).* ACM Press, New York, 2010. To appear.

[6] Simon Garnier, Jacques Gautrais, Masoud Asadpour, Christian Jost, and Guy Theraulaz. Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2):109–133, 2009.

[7] S. Goss, S. Aron, J.-L. Deneubourg, and J.M̃. Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76(12):579–581, 1989.

[8] P. L. Krapivsky and S. Redner. Dynamics of majority rule in two-state interacting spin systems. *Physical Review Letters*, 90(23):238701.1–238701.4, 2003.

[9] M. Mamei and F. Zambonelli. Physical deployment of digital pheromones through RFID technology. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 1353–1354. ACM Press, New York, 2005.

[10] Marco A. Montes de Oca, Eliseo Ferrante, Nithin Mathews, Mauro Birattari, and Marco Dorigo. Opinion dynamics for decentralized decision-making in a robot swarm: Complete data, 2010. Supplementary information page at `http://iridia.ulb.ac.be/supp/IridiaSupp2010-004/`.

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

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

[13] Chris A.C̃. Parker and Hong Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-N problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251, 2009.

[14] David Payton, Mike Daily, Regina Estowski, Mike Howard, and Craig Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, 2001.

[15] Carlo Pinciroli. Object retrieval by a swarm of ground based robots driven by aerial robots. Mémoire de DEA, Université Libre de Bruxelles, Bruxelles, Belgium, 2007.

[16] R. Andrew Russell. Ant trails – An example for robots to follow? In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 1999)*, pages 2968–2703. IEEE Press, Piscataway, NJ, 1999.

[17] Thomas Schmickl and Karl Crailsheim. Trophallaxis within a robotic swarm: Bio-inspired communication among robots in a swarm. *Autonomous Robots*, 17(2):109–133, 2009.

[18] K. Sugawara, T. Kazama, and T. Watanabe. Foraging behavior of interacting robots with virtual pheromone. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2004)*, pages 3074–3079. IEEE Press, Piscataway, NJ, 2004.

[19] B. Werger and M. Matarić. 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)*, pages 625–634. MIT Press, Cambridge, MA, 1996.

[20] Jan Wessnitzer and Chris Melhuish. Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In *Proceedings of the European Conference on Artificial*

Life (ECAL 2003) LNCS 2801., pages 893–902. Springer, Berlin, Germany, 2003.