Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Optimal Collective Decision-Making through Social Influence and Different Action Execution Times

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

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2009-023

August 2009

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/2009-023

Revision history:

TR/IRIDIA/2009-023.001 August 2009

Optimal Collective Decision-Making through Social Influence and Different Action Execution Times

M. A. Montes de Oca, E. Ferrante, N. Mathews, M. Birattari, and M. Dorigo

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
{mmontes,eferrante,nmathews,mbiro,mdorigo}@ulb.ac.be

Abstract. In nature, there are examples of large groups of animals that are capable of making optimal collective-level decisions without the need for global control or information. Understanding the underlying mechanisms of such decentralized decision-making processes may help us to design artificial systems that exhibit some of the desirable properties, like scalability or fault tolerance, that are usually observed in these natural systems. In this paper, we show how a simple social influence mechanism, based on the binary particle swarm optimization algorithm, can make a whole population of agents achieve consensus on one of two possible choices in a completely decentralized way. Furthermore, we show that, if the conditions for achieving consensus are met and each choice is bound to an action that takes time to perform, the population converges to the choice associated with the shortest execution time. We illustrate the applicability of the decision-making mechanism presented in this paper on an example scenario in swarm robotics.

1 Introduction

When an agent is immersed in a social context, its decisions are influenced by the observation of others' decisions and actions [1, 2]. This process makes it possible for an agent to discover information that, otherwise, it may not have access to, or that may be too costly to obtain. Social influence is thus at the root of social learning whereby knowledge is transmitted between agents without using genetic material [3, 4]. The role of social influence in the individual as well as in the collective decision-making process has been studied in human societies in the context of segregation [5], rioting [6], herding in financial markets [7], and other social phenomena [1]. These studies and those performed on animals, for example on social insects, have provided insights that have been used for the design of novel problem-solving techniques [8–12].

In this paper, we study a decentralized decision-making mechanism based on the binary particle swarm optimization algorithm [13, 14], which is inspired by behavioral models of bird flocking (Section 2). We first study whether the mechanism allows a population of agents to achieve consensus on a binary choice problem in a completely decentralized way (Section 3.1). We then study the dynamics of the system when the time needed to perform an action is considered (Section 3.2). We find that when the requirements for reaching consensus are met, the population reaches it on the action with the shortest execution time. We exemplify the utilization of the decision-making mechanism studied in this paper using a swarm robotics scenario where individual robots with limited sensory capabilities must make a binary choice regarding the route to follow to reach a goal location (Section 3.3). Despite their individual inability to measure time or to detect differences between routes, the group of robots, as a whole, is capable of choosing the fastest route in a fully decentralized way. We summarize our results and outline some future works in Section 4.

2 A Decentralized Decision-Making Mechanism

Decentralized decision-making occurs when a group-level decision is made in spite of the fact that individual decisions are based on local information [15]. A flock of birds is a good example of this process as it is capable of choosing a common direction of movement despite the fact that birds continuously change their position, speed and direction within the flock [16].

In this paper, we consider a binary decision mechanism based on the binary particle swarm optimization (PSO) algorithm [13, 14], which in turn is inspired by behavioral models of bird flocking [17, 18]. The system consists of n agents, each of which has a certain predisposition for choosing one of two possible actions, A or B. This predisposition is represented as a continuous variable $p_i^t \in [0, 1]$, $i \in \{1, \ldots, n\}$, which denotes the probability that agent i chooses action A at time step t. The probability that agent i chooses action B at time step t is simply $1 - p_i^t$. Each agent's p_i^t is updated using two pieces of information: the agent's past action and the observation of another agent's action. An agent's own influence leads to the reinforcement of the tendency it already has, whereas the social influence causes the agent to conform to what another agent does.

Each agent i has three variables associated with it: the probability p_i^t of choosing action A at time step t, a state variable v_i^t which governs the dynamics of p_i^t, and a binary variable a_i^t that represents whether action A is taken or not at time step t. The value of a_i^t is a realization of p_i^t. The system evolves using the following rules:

$$v_i^{t+1} = \chi \left[v_i^t + c_1 u_1 \left(a_i^t - p_i^t \right) + c_2 u_2 \left(a_r^t - p_i^t \right) \right] , \quad (1)$$

$$p_i^{t+1} = \frac{1}{1 + e^{-v^{t+1}}} , \quad (2)$$

$$a_i^{t+1} = 1 \text{ with probability } p_i^{t+1}, \text{ and } 0 \text{ with probability } 1 - p_i^{t+1} , \quad (3)$$

where χ is a parameter called *constriction factor* [19], c_1 and c_2 are parameters called *acceleration coefficients*, u_1 and u_2 are uniformly distributed pseudo-random values in the range $[0, 1)$, v_i^{t+1} is clamped such that $v_i^{t+1} \in [v_{min}, v_{max}]$, and r is the index of a randomly selected agent. In Eq. 1, the first term that multiplies χ, that is, v_i^t, acts effectively as a memory term because all updates

Table 1: Parameter settings. The default settings used when exploring the effects of varying individual parameters are shown in boldface. The value $c_i = 2\,, i \in \{1, 2\}$ is chosen so that the expected value of the product $u_i c_i$ is equal to $1\,, i \in \{1, 2\}$. The values chosen for the constriction factor come from the PSO literature [19].

Parameter(s)	Value(s)	Parameter(s)	Value(s)
n	$\{10^1, 10^2, \mathbf{10^3}\}$ agents	v_{max}	$\mathbf{4}$
k	$\{2, 0.5n, 0.75n, \mathbf{n}\}$ agents	μ_A	**10** time steps
p_i^0	$\{0.05, \ldots, \mathbf{0.5}, \ldots, 0.95\}$	μ_B	$\{10, \mathbf{20}, 40\}$ time steps
χ	$\{0.729, 1.0, \mathbf{1.371}\}$	σ_A and σ_B	$\{\mathbf{1}, 4, 8\}$ time steps
c_1 and c_2	$\{0.0, \mathbf{2.0}\}$	Simulation duration (T)	$\mathbf{10^5}$ time steps
v_{min}	$\mathbf{-4}$	Independent Runs (R)	$\mathbf{10^3}$

depend on its previous value. The second term makes the value of the agent's probability change in such a way that it reinforces its previous action: if the action $a_i^t = 1$, p_i^{t+1} will tend to increase; if the action $a_i^t = 0$, p_i^{t+1} will tend to decrease. The third term changes the agent's probability in the same way as the second term, but according to agent r's action. Eq. 2 ensures that v_i^{t+1} is mapped to a value in the range $[0, 1]$ so that p_i^{t+1} can be interpreted as a probability.

We also consider the case in which agents' actions represent tasks that take some time to perform. The execution times associated with actions 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. Once an agent selects an action, its internal variables are updated and the corresponding action execution time distribution is sampled to obtain the number of time steps during which the agent's variables remain unchanged. Furthermore, we also consider the situation in which agents represent embodied robots executing a task in a finite-size environment. In this case, a certain number of agents can be idle at any given time. We let k and $n - k$ be the number of active and idle agents respectively.

3 Experiments and Results

Three series of experiments were carried out. First, we explored under which conditions a population of agents following the rules described in Section 2 achieves consensus. Second, we designed a series of experiments to determine the effects of action execution times on the dynamics of the system. In both cases, the effects of using different parameter settings were investigated. Finally, the decentralized decision-making mechanism was used on a population of simulated mobile robots engaged in a navigation task.

We used Monte Carlo simulation to study the convergence properties of the mechanism described in the previous section using different parameter settings. In Table 1, we list the values used for each parameter of the system and of the simulator.

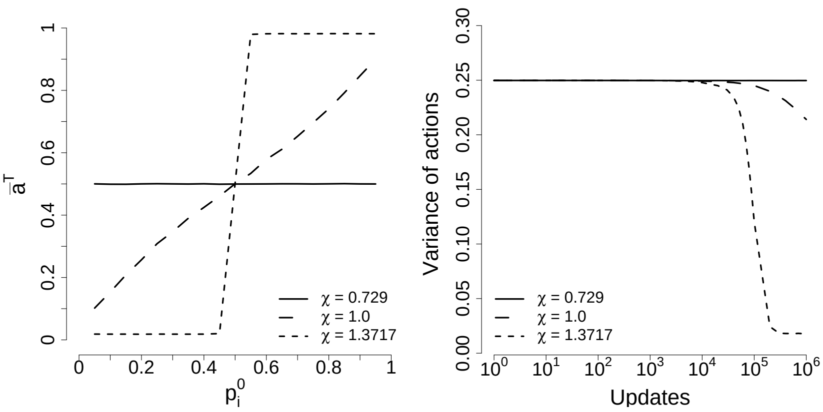

(a) (b)

Fig. 1: Typical behavior of the system without action execution times. (a) Average action $\bar{a}^T$ at the end of the simulation as a function of the agents' initial probabilities p_i^0. (b) Average (over 1000 runs) variance of the agent's actions a_i^t as a function of the number of agent updates ($n = 10^3$ and $T = 10^3$).

3.1 Results Without Action Execution Times

To evaluate whether consensus on one of the two possible choices is achieved, we performed the following analysis. First, we computed the probability to end with all agents choosing action A after T time steps (i.e., at the end of the simulation) as a function of the agents' initial probabilities p_i^0. This probability was estimated by computing $\bar{a}^T = \sum_R \sum_{i=1}^n a_i^T / nR$, which is the average, over all the R independent runs, of the proportion of agents choosing action A at time step T (the probability of choosing action B is simply $1 - \bar{a}^T$). This measure tells us which of the two choices is preferred by the agents. Second, we recorded the variance of the agents' actions at time step T for each run. We then averaged it over all runs which gives us a measure of the agents' average disagreement. Hence, a small value for this measure corresponds to consensus.

We studied the effects that different values of χ, c_1, c_2 and n have on the ability to reach consensus. The typical result is shown in Fig. 1. The results obtained for all the combination of parameters values shown in Table 1 are available at [20].

After analyzing the results, we found that the parameter with the strongest influence on the results is the constriction factor χ. When $\chi \leq 1$ the group does not achieve consensus: The average variance is ~ 0.25 which is the value that corresponds to a 50% A – 50% B population. When $\chi < 1$, $\bar{a}^T$ is equal to 0.5 regardless of the agents' initial probabilities. When $\chi = 1$, $\bar{a}^T$ depends on the initial probabilities in an almost linear way. However, when $\chi > 1$, agents achieve consensus. Additionally, if $p_i^0 = 0.5$, the group chooses action A with a probability of 0.5. As soon as we increase or decrease the initial bias, the population converges to $a_i = 1$ when $p_i^0 > 0.5$ or to $a_i = 0$ when $p_i^0 < 0.5$.

To summarize, consensus is achieved with $\chi > 1$. This is due to the fact that this setting amplifies the value of $v_i^t \rightarrow v_{min}$ if $p_i^0 < 0.5$ or $v_i^t \rightarrow v_{max}$ if $p_i^0 > 0.5$. When $p_i^0 = 0.5$, v_i^t converges to either limit. From a PSO perspective, this result is not surprising. Indeed, Clerc and Kennedy [19] showed that the constriction factor has a critical effect on the stability of the particle swarm. For the experiments shown in the next section we keep $\chi = 1.3717$.

3.2 Results With Action Execution Times

Fig. 2 shows the effects of varying different parameter values on the behavior of $\bar{a}^T$ when the execution time ratio $\mu_{ratio} = \frac{\mu_B}{\mu_A} > 1$[1], which is the case when action A is, on average, faster than action B.

Fig. 2a shows the effect of changing the value of μ_{ratio}. The results show that when μ_{ratio} is high, the population chooses action A even if the initial probability p_i^0 is small. This means that when the execution time ratio is high, only a minority of the population with a preference for the action with the shortest execution time is necessary to eventually make the whole population choose that action. This happens because the agents that choose the fastest action finish before the others and hence, have higher chances to be observed and to influence the choice of other agents. In turn, these influenced agents will have a higher probability of choosing the fastest action and so on. Social influence thus creates a positive feedback process whereby the best choice (in this case, the action with shortest execution time) is propagated in the population. In Fig. 2b, the effect of varying the population size n is shown. Consensus is not achieved when $n = 10$ as the agents' action variance is not small enough[2]. However, with $n = 100$ and $n = 1000$, the population does achieve consensus. This result may be due to the insufficient number of simultaneous interactions with the environment (that is, the maximum value of k) that a small population can have. Fig. 2c shows the effect of using different settings for c_1 and c_2. The population can only detect the difference in execution time when agents make use of socially acquired information, that is, when $c_2 \neq 0$. Fig. 2d shows the effect of changing the number of agents executing simultaneously an action. The best result is obtained when all the population is active.

Finally, changing the standard deviation of the execution times associated with the agents' actions does not seem to have a significant impact on the results [20].

3.3 Decentralized Decision-Making in a Swarm of Robots

As an example of an application scenario, we simulated a group of robots that are required to go from one point to another in the fastest possible way (see

[1] We remind the reader that μ_A and μ_B are the means of the distributions from which execution times are sampled.

[2] At [20], the interested reader can find the complete set of results.

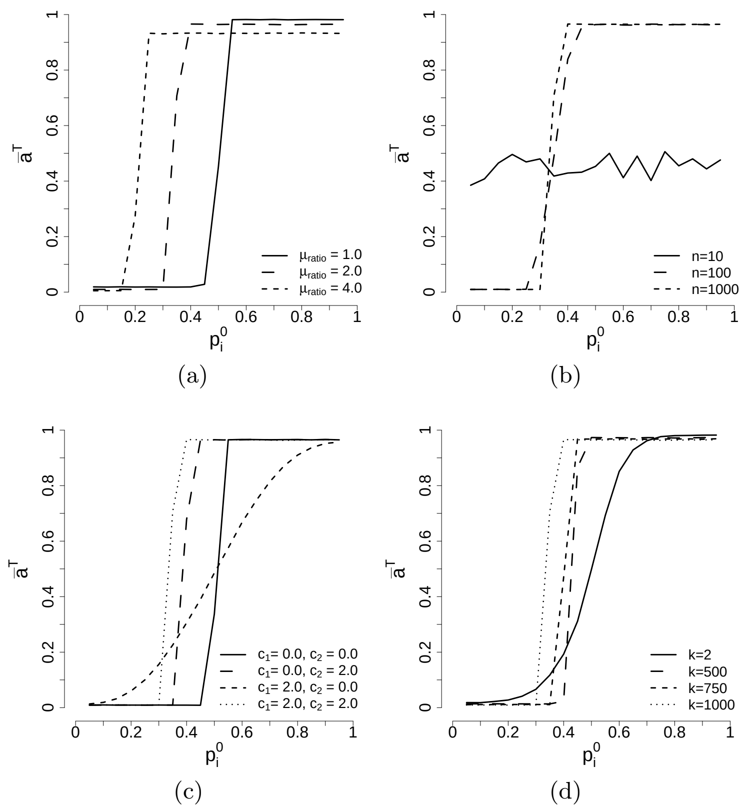

Fig. 2: Typical behavior of the system with action execution times. All figures show $\bar{a}^T$ as a function of the agents' initial probabilities p_i^0. (a) Effect of varying μ_{ratio}. (b) Effect of varying the number of agents n. (c) Effect of varying c_1 and c_2. (d) Effect of varying the number of simultaneously active agents k. The values of the parameters that are not varied are shown in boldface in Table 1.

Fig. 3). Each robot in the group (composed of 20 robots) must choose between going straight to the goal and traverse a muddy area which slows it down, or going around the muddy area through a free, but longer, passage. The best choice can only be determined by trial and error because the robots are not equipped with means to measure time or detect terrain differences. In the example, a maximum of $k = 15$ robots can run in parallel. Initially, each robot chooses a path completely at random ($p_i^0 = 0.5$). To simulate the "observation" of an agent's action, robots communicate their choice with others through an RFID tag that is placed on the ground. Action IDs are read and overwritten each time a robot passes over it. As expected, all robots eventually choose the fastest path even

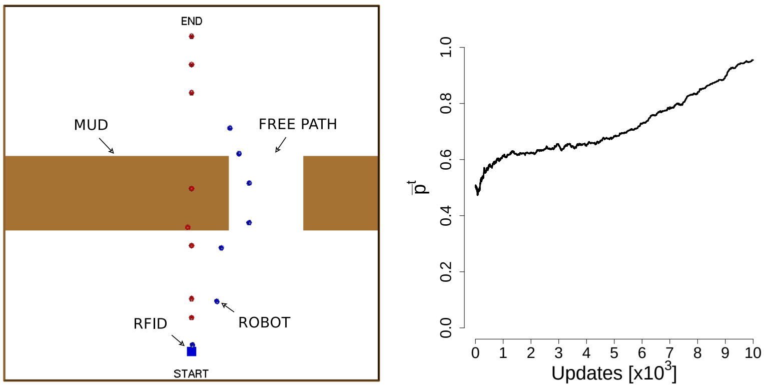

(a) (b)

Fig. 3: (a) Example scenario: Robots need to go from one point to another in the fastest possible way. Each robot can choose between taking the shortest path, which traverses a muddy area, or a longer, but free of mud, path. (b) Development over time of the probability that the group of robots chooses the free path.

though there is no direct intentional communication nor centralized decision-making mechanism. A video showing the system in action can be downloaded from [20].

4 Conclusions

In this paper, we studied a decentralized decision-making mechanism in which agents consider their own history and the actions of others to bias their own choice. We found that when agents are subject to social influence, the population can achieve consensus on one of two alternative choices. If agents' actions are linked to tasks that need time to be performed, the decentralized decision-making mechanism described in this paper allows the population to choose the fastest action, even if the individual agents cannot detect any difference between the alternatives. Furthermore, the population can better identify the action with the shortest execution time when the number of agents and the ratio between the alternatives' execution time grow. The mechanism and examples presented in this paper show that social behavior can lead to collective cooperative behavior.

Future work includes the investigation of simpler mechanisms, multi-choice problems, continuous ranges of decisions and dynamic environments.

Acknowledgments. This work was supported by the *SWARMANOID* project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission (grant IST-022888). M. Birattari and M. Dorigo acknowledge support from the F.R.S.-FNRS of the French Community of Belgium.

References

1. Chamley, C.P.: Rational Herds. Economic Models of Social Learning. Cambridge University Press, New York (2004)
2. López-Pintado, D., Watts, D.J.: Social influence, binary decisions and collective dynamics. Rationality and Society **20**(4) (2008) 399–443
3. Nehaniv, C.L., Dautenhahn, K., eds.: Imitation and Social Learning in Robots, Humans and Animals: Behavioral, Social and Communicative Dimensions. Cambridge University Press, Cambridge, UK (2007)
4. Curran, D., O'Riordan, C.: Increasing population diversity through cultural learning. Adaptive Behavior **14**(4) (2006) 315–338
5. Schelling, T.C.: Dynamic models of segregation. Journal of Mathematical Sociology **1** (1971) 143–186
6. Granovetter, M.: Threshold models of collective behavior. American Journal of Sociology **83**(6) (1978) 1420–1443
7. Schiller, R.J.: Irrational Exuberance: Second Edition. Princeton University Press, Princeton, NJ (2005)
8. Bonabeau, E., Dorigo, M., Theraulaz, G.: Swarm Intelligence: From Natural to Artificial Systems. Oxford University Press, New York (1999)
9. Bonabeau, E., Dorigo, M., Theraulaz, G.: Inspiration for optimization from social insect behaviour. Nature **406**(6791) (2000) 39–42
10. Kennedy, J., Eberhart, R.: Swarm Intelligence. Morgan Kaufmann, San Francisco, CA (2001)
11. Dorigo, M., Stützle, T.: Ant Colony Optimization. MIT Press, Cambridge, MA (2004)
12. Engelbrecht, A.P.: Fundamentals of Computational Swarm Intelligence. John Wiley & Sons, Chichester, UK (2005)
13. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of IEEE International Conference on Neural Networks, Piscataway, NJ, IEEE Press (1995) 1942–1948
14. Kennedy, J., Eberhart, R.: A discrete binary version of the particle swarm algorithm. In: Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, Piscataway, NJ, IEEE Press (1997) 4104–4108
15. Marshall, J.A.R., Dornhaus, A., Franks, N.R., Kovacs, T.: Noise, cost and speed-accuracy trade-offs: decision-making in a decentralized system. J. R. Soc. Interface **3**(7) (2006) 243–254
16. Couzin, I.D., Krause, J.: Self-organization and collective behavior in vertebrates. Advances in the Study of Behavior **32** (2003) 1–75
17. Reynolds, C.W.: Flocks, herds, and schools: A distributed behavioral model. ACM Computer Graphics **21**(4) (1987) 25–34
18. Heppner, F., Grenander, U.: A stochastic nonlinear model for coordinated bird flocks. In: The Ubiquity of Chaos. AAAS Publications, Washington, D. C. (1990)
19. Clerc, M., Kennedy, J.: The particle swarm - explosion, stability, and convergence in a multidimensional complex space. IEEE Transactions on Evolutionary Computation **6**(1) (2002) 58–73
20. Montes de Oca, M.A., Ferrante, E., Mathews, N., Birattari, M., Dorigo, M.: Optimal collective decision-making through social influence and different action execution times: Complete data (2009) Supplementary information page at `http://iridia.ulb.ac.be/supp/IridiaSupp2009-005/`.