Towards Dynamic Adaptation of the Majority Rule Scheme

C. Krause
Department of System Analysis and Modeling, Hasso Plattner Institute for Software Systems Engineering

E.P. de Vink*
Department of Mathematics and Computer Science, Technische Universiteit Eindhoven
Centrum Wiskunde en Informatica, Amsterdam

P.J. de Vink
Faculty of Science, University of Amsterdam

Abstract The majority rule scheme has been applied in the setting of robot swarms as a mechanism to reach consensus among a population of robots regarding the optimality of one out of two options. In the context of distributed decision making for agents, we consider two schemes of combining the majority rule scheme with dynamic adaptation for the well-known double bridge problem to cater for a situation where the shortest path changes over time. By modeling the systems as Markov chains, initial results regarding the quality and the trade-off of efficiency and adaptation time can be obtained.

1 Introduction

Distributed decision making by collectives of autonomous agents ultimately relies on the available interaction schemes. For populations of ants it is well-known that by means of pheromones ants can select the shortest of two paths leading from the nest to a place with food. Since pheromones evaporate over time the shortest path is indicated more strongly than the longer one. For robots swarms no satisfactory physical counterpart of pheromones has been agreed upon yet. In [5] Montes de Oca et al. propose the mechanism of the majority rule –as studied in sociology, economics and physics– as a computational alternative. The approach has been developed for swarms of autonomous robots in a static environment where the aim is to reach consensus among all robots on which one out of two paths is the shortest.

Generally, distributed decision making assumes a static environment. However, ants are very well capable to reconsider their preferences in a changing environment. In a variation of the so-called double bridge experiment, it has been shown that starting from a consensus situation where a single path is preferred, the preference of the population will shift from the one to the other path in orders of minutes when dynamically the shortest path has been made the longest and vice versa [1]. Thus, ants are able to deal with spatial dynamicity. For robot swarms, however, it is not obvious how to achieve this.

In a more abstract setting of agent systems, we propose two adaptation schemes that can be combined with the majority rule approach. One is based on a suggestion raised in [5] where there is always a minimum subpopulation of both opinions; another allows teams of agents to switch their opinion with a small probability. The resulting systems can be modeled as discrete-time probabilistic automata. The corresponding Markov chains can be fed to the PRISM model checker [2] for analysis and comparison of the two adaptation schemes.

This paper reports on work-in-progress and has been inspired by [4] where Bio-PEPA is employed for the analysis of a non-adaptive version of the double-bridge problem. To the best of our knowledge this contribution constitutes a first proposal to combine distributed decision making and dynamic adaptation.

*Corresponding author, `evink@win.tue.nl`

Submitted to:
QAPL 2013

2 The *Switch* and *MinPop* schemes

The double bridge problem involves a population of agents having opinion A or opinion B and a number of locations, among which we distinguish the nest N, destination A and destination B. Agents may change their opinion according to the rules described below. Each location has a specific non-negative, integer distance from the nest. The distance from the nest to itself is 0, destinations A and B have distance d_A and d_B from the nest N. We assume that there is a location L_d at distance d from the nest, for each $0 < d \leqslant \max\{d_A, d_B\}$. With specific probability, agents are 'teleported' in teams of three agents from the nest to either of the destinations and return step-by-step via the locations L_d to the nest again. Conceptually, the distance from destination A and from destination B to the nest may vary over time.

Time-steps are discrete. In the base scheme, following [5], at every time point the following happens:

- Each agent at the nest throws a fair coin to decide if it is willing to leave the nest.
- Agents do not leave the nest individually. Instead, teams of exactly 3 agents are randomly formed out of the agents willing to leave. If 1 or 2 agents do not join a team, they will stay at the nest.
- Each team chooses their destination based on the *majority rule*: If a majority of 2 or more agents has opinion A, destination A is chosen; if a majority of 2 or more agents has opinion B, destination B is chosen. Moreover, the minority of 0 or 1 agent changes its opinion to the opinion of the majority. The team will proceed to the chosen destination.
- Agents not at the nest, at some location at distance $d > 0$ say, will proceed towards the nest, i.e. to location L_{d-1}, where L_0 is meant to be the nest N.

So, in the base scheme, the majority rule combines two aspects: (i) the destination that is preferred by the majority of the agents in a team is chosen; (ii) if the opinions of the team are mixed, the agent with the opinion deviant from the majority changes its preference. Note that in our modeling with agents, unlike living ants or physical robots, teams of agents leaving the nest are teleported directly to the preferred destination. Since only the elapse of time while the agents are traveling, the so-called latency [3], is relevant for the scheme, teleportation is a valid, computationally attractive simplification.

The base scheme described above is not adaptive. For any number of agents larger than 2, the population will reach consensus with probability 1, either on opinion A or on opinion B. From there on agents will not change their opinion anymore. In [5], under simplifying assumptions, a continuous-time analysis is made of the base scheme. The probability to eventually have opinion A for all agents or opinion B for all agents depends on the ratio of the path length to destination A and B and on the initial fractions of agents preferring A or B, respectively. Intuitively, if the larger part of agents prefers A, say, they will often have majority in newly formed team, making agents preferring B to change their mind. On the other hand, if there is a substantial difference in path lengths, say the path to B is twice as long as the path to A, then teams choosing B stay away longer from the nest and hence can contribute less to the voting in the teams.

In order to make the scheme adaptive, i.e. that the system can move away from consensus in reaction to a change in path length, we propose two variations. In the so-called *Switch* scheme there is small probability for an unanimous team to switch to the opposite opinion once they have returned to the nest. Thus a team composed of 3 agents preferring A may return as 3 agents with opinion B, and vice versa. The ratio is that in case of consensus, say on A, there is still a non-zero probability for forming teams of majority B. So, there is always a modest drift away from the overall preference for A. When the path to B has become shorter due to a change in the environment, the whole population may advance to B.

The other variation is to appoint 'stubborn' agents. Stubborn agents do not change their mind in a team of opposite majority. So, the overall population of agents of opinion A has as minimum the constant number of stubborn agents preferring A, and analogously for B. We call this the *MinPop* scheme, because there is always a minimal subpopulation of either opinion. These subpopulations are able to seed a shift to a new steady state when the order of the path lengths changes.

Given the choice of parameters, i.e. initial numbers of regular agents, distance for target destinations, switch probability or numbers of stubborn agents, the system can be represented as a discrete-time Markov chain. Currently the dynamicity of the environment is not explicitly modeled; the current path lengths are assumed to be the ones that are newly in place. Based on the Markov chain representation we consider two aspects of the dynamic adaptivity of the two schemes: (i) efficiency, how often are teams sent to the destination with the shortest path; (ii) adaptation time, the expected time needed to let all agents adapt to the new environment.

3 First results

In order to assess the efficiency and adaptivity of the *Switch* and *MinPop* scheme, we use the PRISM model checker [2] to compute corresponding metrics. The results are shown in Figure 1. For simplicity, the path length to destination A is kept constant to 1. The parameter r specifies the path length to destination B. For the *Switch* scheme, we vary the switch probability between 0.002 and 0.04. For *MinPop*, we vary the number of stubborn agents between 2 and 6; these are added to the population for each opinion. Note that in order to allow dynamic adaption, a strictly positive switch probability is required for *Switch*, and at least 2 stubborn agents of either opinion are required for *MinPop*.

For both efficiency experiments, Figure 1ab, the initial configuration consists of 6 A and 6 B agents which are all located at the nest. We calculate the efficiency in two steps. In the first step, we use the number of teams at destination A as a state reward and compute the steady-state reward R_A. In the second step, we use the number of teams at A or at B as a state reward and compute the steady-state reward R_{AB}. We then compute $E = R_A/R_{AB}$ as an efficiency metric, informally the average percentage of teams sent to A. For both schemes, the efficiency increases with higher values for r. The charts show that the efficiency decreases when increasing the switch probability in the *Switch* scheme, or the number of stubborn agents in the *MinPop* scheme. We conjecture that for $r > 1$ the efficiency of the *Switch* scheme can be made arbitrarily high by choosing an appropriately small switch probability.

In the two adaption time experiments, we start from an initial configuration consisting of 12 B-agents in the nest and no A-agents. We define a transition reward of 1 for all transitions. We then compute the expected reward to reach a state with the maximal number of A-agents. This gives us the expected adaptation time T, i.e. the number of steps to turn a B-population into an A-population. Smaller numbers for the adaptation time mean that the scheme is faster in adapting to a change in the environment. The adaptivity of both *Switch* and *MinPop* can be optimized by increasing the switch probability or the number of stubborn agents, respectively.

It is evident that there is trade-off between the efficiency and the adaptation time for both schemes. Increasing the adaptivity is possible only by stimulating the sending of more agents to the currently unfavored path, which in turn decreases the efficiency. This makes it difficult to directly compare *Switch* and *MinPop*. To do compare one may consider the metric $X = E/ln(T)$, which yields values that vary only mildly when changing the switch probability or the number of stubborn agents. For $r = 2$, we obtain averages of $X_{Switch} = 0.100$ and $X_{MinPop} = 0.127$. For the limited range of values considered, we can conclude that *MinPop* performs better than *Switch*, i.e. for comparable efficiency values, *MinPop* has a higher adaptivity. However, especially for small populations, the *Switch* scheme allows a more fine-grained control as the switch probability can be adjusted with arbitrary precision.

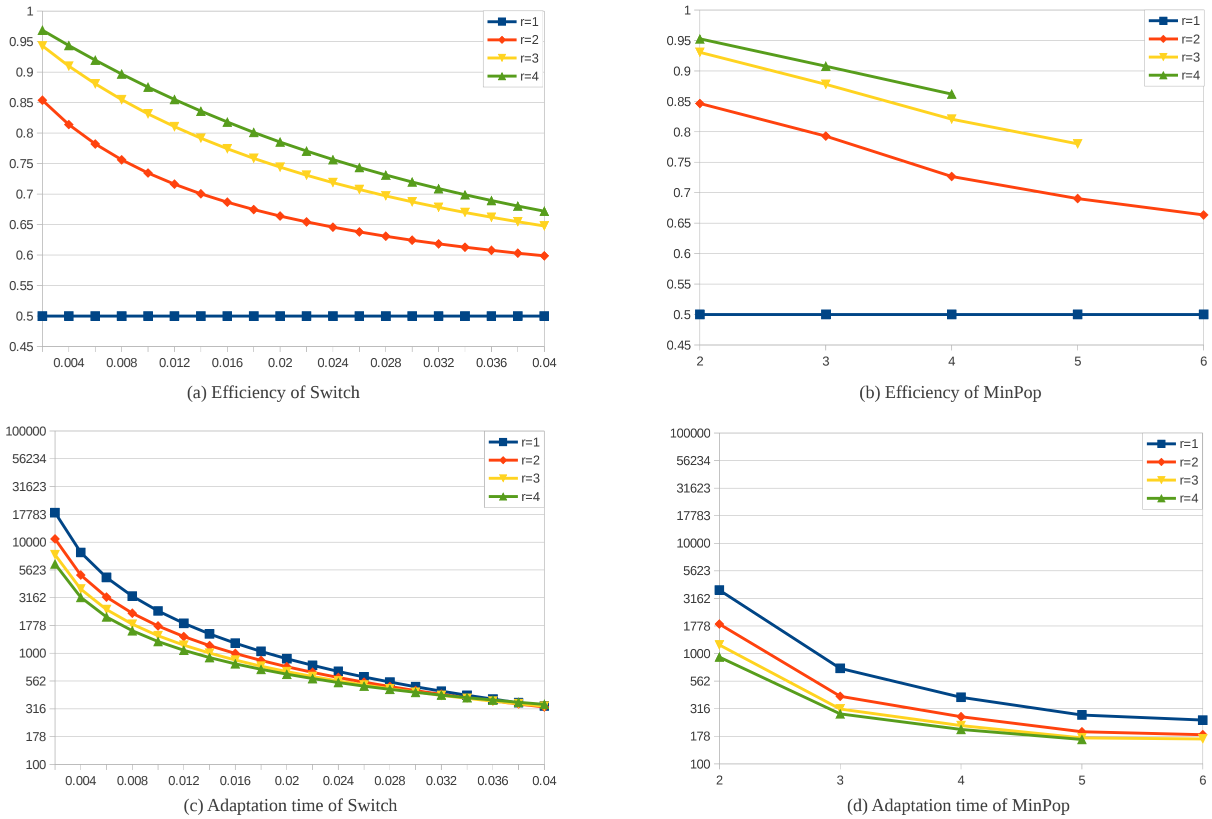

(a) Efficiency of Switch (b) Efficiency of MinPop (c) Adaptation time of Switch (d) Adaptation time of MinPop

Figure 1: (a–b) Efficiency of Switch and MinPop; (c–d) Adaption time of Switch and MinPop (log scale).

4 Conclusion

We reported on first results towards a dynamically adaptive variant of the majority rule scheme. We presented the two adaptive schemes *Switch* and *MinPop* and compared their efficiency and adaptivity using probabilistic model checking. More work is needed to deal with larger population sizes and to study refined but presumably better adaptation schemes.

References

[1] S. Goss, S. Aron, J.L. Deneubourg & J.M. Pasteels (1989): *Self-organized Shortcuts in the Argentine Ant*. *Naturwissenschaften* 76, pp. 579–581.

[2] M. Kwiatkowska, G. Norman & D. Parker (2011): PRISM *4.0: Verification of Probabilistic Real-time Systems*. In: *CAV'11*, LNCS 6806, pp. 585–591.

[3] R. Lambiotte, J. Saramäki & V.D. Blondel (2009): *Dynamics of latent voters*. *Physical Review E* 79(4), pp. 046107/1–6.

[4] M. Massink, M. Brambilla, D. Latella, M. Dorigo & M. Birattari (2012): *Analysing Robot Swarm Decision-Making with Bio-PEPA*. In: *ANTS 2012*, LNCS 7461, pp. 25–36.

[5] M.A. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari & M. Dorigo (2011): *Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making*. *Swarm Intelligence* 5(3–4), pp. 305–327.