

Using Response Probability to Build System Redundancy in Multiagent Systems

(Extended Abstract)

Annie S. Wu
Elec. Eng. & Comp. Sci.
University of Central Florida
Orlando, FL 32816
aswu@cs.ucf.edu

R. Paul Wiegand
Inst. for Simulation & Training
University of Central Florida
Orlando, FL 32816
wiegand@ist.ucf.edu

Ramya Pradhan
Elec. Eng. & Comp. Sci.
University of Central Florida
Orlando, FL 32816
ramya.pradhan@knights.ucf.edu

ABSTRACT

This work shows that a simple response probability model may be used in decentralized multiagent systems to dynamically build and maintain redundancy in terms of agent experience. We model this system theoretically and use this model to provide guidance on how to choose response probability values that ensure that task needs are met and that ensure the a given level of redundancy is achieved.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems performance measures]

Keywords

Agent coordination, self-organization, redundancy

1. INTRODUCTION

In this paper, we examine the use of probabilistic response in decentralized multiagent systems (MAS) and how such a mechanism may be used to tune the level of redundancy in an MAS. Redundancy is defined as “extra” agents beyond the required number that are able to perform a task. These agents provide the MAS with a back up pool in the event that the primary actor or actors for a task are disabled or lost. We present a simple but effective method to build and maintain a back up pool dynamically over time.

We are specifically interested in tasks where previous experience leads to better performance, and redundancy may be generated by allowing agents to gain experience while acting on a task. Such an approach raises issues of how to balance the tradeoff between optimizing performance and generating experience. This balance is similar to the exploration-exploitation balance that many machine learning algorithms seek as they attempt to dynamically learn new information and optimally use the information they currently have at the same time [?, ?]. Many real systems, such as social insect colonies, face a similar tradeoff [?, ?]. These systems appear to use a simple mechanism, an agent *response probability*, to balance the tradeoff between giving individual ants or bees the opportunity to gain experience on a task such as foraging and sending out their best foragers to optimally harvest food.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Examples of problems where such an approach is useful include applications where agents must collaborate and develop trust-based interactions [?, ?, ?]. Experience is the information that an agent gains from previous interactions with other agents or with the environment. Experience is important because it affects an agent’s decision making in future interactions.

Here, we present a formal analysis of this response probability model, how it affects the ability of an MAS to satisfy a task’s needs, and how it may be used to tune the level of redundancy in an MAS.

2. RESPONSE PROBABILITY IN AN MAS

Our test problem is a distributed task allocation problem with a single task that needs to be addressed. Given a decentralized MAS consisting of n agents, $x : x < n$ agents must respond (form a *response team*) each time the task occurs. Each agent independently decides whether or not to respond to a task and gain experience.

We assume that our MAS is a variable response threshold system in which each agent has a threshold value for each task. These thresholds may be thought of as an agent’s relative willingness or speed in responding to a task in need. Agents with lower thresholds respond faster than agents with higher thresholds. In this *implicit ordering*, lower threshold agents have more chances to respond to the task than the higher threshold agents. Therefore, the response team in this system only consists of first x low threshold agents.

We introduce a response probability, $s : 0.0 \leq s \leq 1.0$, that affects whether or not an agent will act once its threshold has been met. If an agent’s threshold is met before x agents have responded to the task, then that agent is offered an opportunity to act and becomes a *candidate*. A candidate will choose to act with probability s . If the candidate chooses to act, it becomes an *actor* and the number of responders to the task increases by one; if it does not choose to act, then the next agent in the ordering is offered the opportunity to act and becomes a candidate. Thus, the response probability is the probability that a candidate will become an actor. As we move from left to right in the implicit ordering, the addition of the response probability s means that the response team may not always consist of the left most x agents, and that agents beyond the first x can receive an opportunity to act and gain experience. We define a *trial* to be one instance in which a task requires action from the MAS. While only a maximum of x agents gain experience in each trial, over multiple trials, we expect more than the first x agents to gain experience on the task. We define the *pool* to be the redundant agents beyond the first x agents that have gained experience over multiple trials. At high s values, because most of the first x agents have a high probability of choosing to act, agents beyond first x have a low probability of becoming

a candidate, and therefore we expect the pool size to be close to zero. At low s values, because most of the first x agents have a low probability of choosing to act, agents beyond first x have a high probability of becoming a candidate, and therefore we expect the pool size to increase. Thus, we see that s affects both the ability of an MAS to form a response team and the level of redundancy that can be achieved by the MAS over multiple trials.

3. ANALYSIS

Our analysis focuses on two aspects of the response probability, s . First, how does s affect the ability of an MAS to successfully form a response team. Second, how does s affect the ability of an MAS to build and maintain redundancy.

Recall that an MAS is only successful in responding to a task if it can form a response team, and a response team is formed when x agents choose to work on the task. The first question is whether we can determine what s values will allow the MAS to successfully form a response team. Using the Chernoff inequality, we find that, as n grows, a complete team will almost surely be formed when $s_1 > \frac{x}{n}$ and will almost surely not be formed when $s_2 < \frac{x-1}{en}$.

Redundancy is achieved when there is a pool of agents beyond the first x agents that have gained experience on the task. Because, in each trial, a maximum of x agents can act and gain experience, redundancy can only be generated over multiple trials. The second question is whether we can determine the appropriate value of s to use in a system when a specific level of redundancy is desired.

We define c to be the desired level of redundancy where cx is the desired number of agents with experience (and the size of the pool is $cx - x$). Our goal is to determine the s values for which the cx^{th} agent is highly likely to have gained experience and be part of the pool. The probability of the i^{th} agent acting and therefore gaining experience (P_i) is no smaller for agents preceding the cx^{th} in the ordering (i.e., $P_i \geq P_{cx}$ if $i \leq cx$) since those will be given the opportunity to act sooner and all agents have the same s . Using the Chernoff inequality, we find that if $s_3 \leq \frac{1}{ec}$, there is a constant probability that the cx^{th} agent will gain experience and if $s_4 > \frac{1}{c}$ then it will almost certainly fail to gain experience.

We present an empirical case containing 100 agents ($n = 100$), a task that requires 20 agents ($x = 20$), and a redundancy factor of 2 ($c = 2$) to show consistency with our formal advice such that the first 40 agents will gain experience on that task given enough trials. Figure 1 shows influence of s on the formation of a response team and a backup pool of agents. The x -axis plots s values and the y -axis measures the average number of actors in 20 simulations. At $s_1 > 0.2$, successful team formation occurs and is shown here by the average number of actors being greater than or equal to 20 (task requirement). The bounds for the 40th agent to gain experience and belong to the backup pool ($s_3 \leq 0.18$, $s_4 > 0.5$) is also satisfied and thus the redundancy requirement is also met.

4. CONCLUSIONS

In this work, we examine the use of an agent response probability as a method for generating redundancy in an MAS. Redundancy in this work is defined as a pool of extra agents with experience on a task beyond the necessary number of agents for that task. Having agents respond probabilistically rather than deterministically to a task demand means that not all of the primary actors for a task may act when a task requires attention, giving other agents an opportunity to act and gain experience on the task.

The inspiration for this work comes from social insect societies which exhibit a similar mechanism for balancing the gain of new information (let agents gain experience and improve performance)

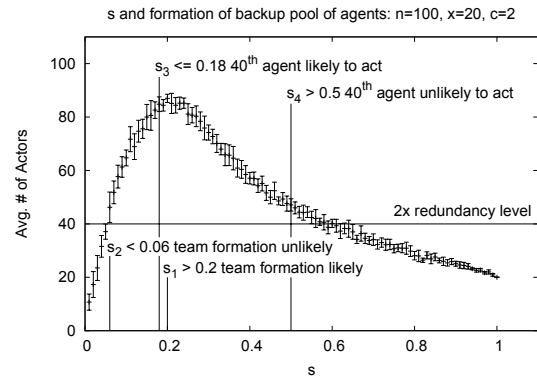


Figure 1: Effect of s on the response team formation and formation of backup pool of agents in a single task.

and the optimal use of known information (send out the best performing agents to do a job). The simple response probability mechanism not only effectively generates redundancy in a decentralized system, but it also allows the system to maintain redundancy dynamically over time despite perturbations to the system such as loss of agents. These characteristics of robustness and adaptability are also desirable in computational decentralized multiagent systems and our goal is to understand how response probability may be effectively used in engineered MAS.

We present a formal analysis of a response threshold based system in which agents act probabilistically. Given the number of agents in the multiagent system and the number of agents needed to satisfy a task, we are able to estimate two factors regarding the system's response probability value. First, we can estimate the range of response probability values that are likely to allow the MAS to satisfy the task demands. Second, we can estimate the range of response probability values that are likely to allow the MAS to achieve a specific level of redundancy.

5. ACKNOWLEDGEMENTS

This work was supported in part by ONR grant #N000140911043.

6. REFERENCES

- [1] J. Atkinson and D. Rojas. On-the-fly generation of multi-robot team formation strategies based on game conditions. *Expert Systems Applications*, 36(3):6082–6090, 2009.
- [2] D. Deangelis and K. S. Barber. Security applications of trust in multi-agent systems. *Computer Security*, 19(1):57–99, 2011.
- [3] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan, 1975.
- [4] F. Ravary, E. Lecoutey, G. Kaminski, N. Châline, and P. Jaisson. Individual experience alone can generate lasting division of labor in ants. *Current Biol.*, 17:1308–1312, 2007.
- [5] T. Sugawara, K. Fukuda, T. Hirotsu, S. Y. Sato, O. Akashi, and S. Kurihara. Fluctuated peer selection policy and its performance in large-scale multi-agent systems. *Web Intelligence and Agent Systems*, 8(3):255–268, 2010.
- [6] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [7] A. Weidenmüller. The control of nest climate in bumblebee (*bombus terrestris*) colonies: interindividual variability and self reinforcement in fanning response. *Behavioral Ecology*, 15(1):120–128, 2004.