Influence of Inter-agent Variation on System Redundancy in Multiagent Systems
(Doctoral Consortium)

Ramya Pradhan
Advisor: Dr. Annie S. Wu
Electrical Engineering &
Computer Science
University of Central Florida
Orlando, FL 32816
ramya.pradhan@knights.ucf.edu

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

Keywords

Agent coordination, self-organization, redundancy

1. INTRODUCTION

The goal of our research is to understand and model the effects of inter-agent variation in adaptively building and maintaining a back up pool of agents in multiagent systems (MAS). Redundancy is important in maintaining system stability during pertubations. We consider a decentralized MAS to which tasks are assigned. Specifically, we consider tasks in which agents gain experience by performing the tasks, and where past experience improves future performance. We refer to *redundancy* as the presence of a pool of extra agents with work experience.

The motivation for our research is drawn from how we, as humans, build a back up team of workers with experience. Consider the following example of a football team. The gameday roster consists of 45 players of whom eleven play at any given time. Let it be the case that the team has eleven best players. The team achieves best performance each time it plays with these eleven. If, unfortunately, one or more members of this best team are unable to play, then the team no longer has the best performance. However, loss of a team's performance is seldom seen due the presence of the back up players with experience for the eleven best players. Thus, the presence of back up players with experience contributes to maintaining team performance.

The key to our investigation is balancing performance versus redundancy, that is, balancing the optimization of team performance versus the creation of a back up pool. When the best players always play, the team performance is at its best. The team, however, is unstable without the back up pool during perturbations. When the players in the back up pool play, the back up players gain experience. The team, however, is not capable of best performance. We investigate a method that uses inter-agent variation to balance the team performance and system redundancy.

We assume a response threshold-based MAS [1]. Each agent has a threshold at which it responds to a task's stimulus. Agents with lower thresholds are more likely to respond to a task's stimulus than agents with higher thresholds. If lower threshold agents are the only ones that respond to a task (like the best players always playing), then there is little room for forming a back up pool of agents. In order for the lower threshold agents to not always act, we introduce *response probability* [3], an inter-agent variation, to the response threshold MAS. In this model, agents choose to act or not to act when their response threshold is met based on response probability.

In this research abstract, we discuss the current state of our research with reference to the modeling and formal analysis of the response probability model. The specific areas that we investigate are the use of response probability as a tunable parameter for balancing team performance and system redundancy, the effects of response probability when multiple tasks are assigned to a MAS, and finally, the applications suitable for using the response probability model.

2. CURRENT STATE OF RESEARCH

We test the response probability model on a decentralized task allocation problem. We assume a response threshold-based MAS consisting of $n, n > 1$ agents. We assume a task assigned to the MAS. The task requires $x, 1 \leq x \leq n$, agents to work on it. The agents gain experience when they work on the task. An agent's performance improves with experience. Based on the variations in the agents' response thresholds, there exists an *implicit* ordering of agents (from the lowest to the highest) for the task. We introduce response probability $s, 0.0 \leq s \leq 1.0$, with which agents act when their response threshold is met.

An agent can be in one of three states. In the *inactive* state, the agent does not work on the task. When an inactive agent's response threshold is met, it gains an opportunity to act and transitions to the *candidate* state. When in the candidate state, an agent can choose to act with probability equal to response probability. If an agent chooses to act on a task, then it transitions to the *actor* state. When x actors are obtained, a *response team* is formed.

In the response threshold model, where agents always act when their response threshold is met, the response team consists of only the first x agents (the best players). Over multiple task occurrences, the number of agents with experience is exactly x. With the addition of response probability ($0 < s < 1.0$), where agents may not always act when their response threshold is met, the response team consists of agents beyond the first x agents. Over multiple task occurrences, the number of agents with experience is at least x, since the same agents may not always act [2]. Using Chernoff inequality, we have shown in [4, 5] that a response team is always formed when $s > \frac{x}{n}$ and most surely not formed when $s < \frac{x-1}{en}$.

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.*

Given the implicit ordering of agents for a task, a back up pool of $cx, c > 0$, agents should consist the cx^{th} agent. Using Chernoff inequality, we have shown in [4, 5] that a back up pool will contain the cx^{th} agent with constant probability when $s \leq \frac{1}{ec}$ and will almost certainly not contain the cx^{th} agent when $s > \frac{1}{c}$.

We have empirically verified our formal predictions with agent-based simulations. We present and analyze a simple application in [2]. In this application, when an agent performs a task, the time taken to complete that task decreases, and therefore, the one with the most experience is the one that can complete the task the fastest. We show that when agents act with low response probability, the size of the back up pool increases, and when agents act with high response probability, the size of the back up pool decreases. We compare the performances of two systems when agents are removed, where one is with response probability and the other is without response probability. We find that at low response probability values, the system without response probability has better performance due to *windowing effect* (see [2]), and at high response probability values, the system with response probability has better performance. We have shown for a MAS that is assigned one task, response probability can be used to achieve both team performance and system redundancy. For further discussions, interested readers are also referred to [5].

3. EXTENSIONS

We are investigating the extension of the response probability model to the case where multiple tasks are assigned to a MAS. The extension presents unique challenges for formally analyzing the response probability model. Though somewhat inter-related, we have categorized the challenges into those that are specific to the tasks and those that are specific to the MAS. We illustrate these challenges in this section.

We first present challenges due to the specifications of the tasks assigned to a MAS. Consider the case where two concurrent tasks are assigned to a MAS consisting of n agents. The first and the second tasks require x_1 and x_2 agents respectively. Since the formal prediction for s values for team formation is a function of x and n in the single task case, the expectation is to extend the formal prediction in terms of x and n for the two tasks case. Given the possibilities of $x_1 + x_2 < n$, $x_1 + x_2 = n$, and $x_1 + x_2 > n$, the bounds for s that ensure team formation with a high probability will vary accordingly.

We next present challenges due to the response threshold-based model of the MAS. Again, we consider the case where two concurrent tasks are assigned to a MAS consisting of n agents. Each agent now has two response threshold values, one for each task. Therefore, there are two implicit orderings of agents, one for each task. Given the two orderings, there exists $\frac{k}{n}$ probability that the same k agents are present (at same or different positions) in the first k positions in both orderings. Given that the tasks are concurrent, an agent acting on one task cannot act on another task at the same time. Therefore, depending on the tasks' requirements, if $x_1 + x_2 \geq n$, then there is a constant probability that one of the teams will not be formed.

Using the two tasks example we see that as the number of tasks increase, the variations in the tasks' requirements will increase exponentially. Also, as the number of tasks increase, the overlaps among the tasks' orderings and the requirements of the task taken together will affect generalizing the formal bounds for the s values that ensure team formation and redundancy.

We are investigating applications that are expected to benefit from the use of the response probability model. We believe that an application with the following characteristics will benefit the most: (1) the application should require only a subset of agents to act and does not fail if agents do not act, (2) the application should occur more than once and provide for agents to act and gain experience based on their response threshold values, (3) the application should provide for measuring individual agent performance and team performance.

4. CONCLUSION

In our research, we study the effects of adding a response probability to the behavior of agents in a response threshold multiagent system and how this mechanism may be used to dynamically adapt the level of redundancy in the MAS. We are interested in problems in which past experience on a task affects (benefits) an agent's future performance on that task. In such problems, redundancy in the form of extra agents with experience can lead to greater team robustness in the face of loss of agents. Response probability then allows us a parameter for adjusting the balance between using the most experienced agents to obtain optimal performance versus building redundancy by giving less or inexperienced agents opportunities to gain experience. We present a formal analysis of the response probability model, discuss the challenges and avenues for application.

In summary, we propose the use of a probabilistic response in agent decision making as a way to balance the optimization of a system's performance versus a system's need for robustness via redundancy. Our method, though simple, provides an elegant and powerful way to tune this balance and build robust and adaptable teams.

5. ACKNOWLEDGEMENTS

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

6. REFERENCES

[1] Eric Bonabeau, Guy Theraulaz, and Jean-Louis Deneubourg. Fixed response thresholds and the regulation of division of labor in insect societies. *Bulletin of Mathematical Biology*, 60:753–807, 1998.

[2] Ramya Pradhan and Annie S. Wu. On the relationship between response probability and redundancy in teams of collaborating agents. In *Proc. of the 7th Intl. Conf. on Collaborative Computing*, 2011.

[3] Anja 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.

[4] Annie S. Wu, R. Paul Wiegand, and Ramya Pradhan. Using response probability to build system redundancy in multiagent systems. In *Proc. of the 12th Intl. Conf. on AAMAS*, 2013.

[5] Annie S. Wu, R. Paul Wiegand, Ramya Pradhan, and Gautham Anil. The effects of inter-agent variation on developing stable and robust teams. In *Proc. of the AAAI Spring Symposium*, 2012.