Adaptable Decentralized Task Allocation of Swarm Agents

Doctoral Consortium

Vera A. Kazakova
University of Central Florida
Orlando, Florida
kazakova@cs.ucf.edu

ABSTRACT

Scalable task allocation in dynamic real-world domains often requires efficient, robust, decentralized, and adaptable approaches. Response threshold reinforcement is a biologically-inspired model of probabilistic action that has been shown to lead to efficient task allocation among swarm agents that do not reason or communicate, making it a highly scalable and low cost solution. The model leads agents to specialize, resulting in reduced costs of interference and task switching, as well as to improved efficiency and adaptability to initially unknown environments. While initial specialization of this and other models is investigated in much of existing literature, subsequent re-adaptation to domain changes is seldom verified. Our goal is to investigate the robustness of response threshold reinforcement to various environmental changes, as well as to compare this model to other decentralized approaches.

KEYWORDS

agent cooperation: biologically-inspired approaches and methods; multi-robot systems; agent societies: self-organization

ACM Reference Format:
Vera A. Kazakova. 2019. Adaptable Decentralized Task Allocation of Swarm Agents. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Scalable task allocation in real-world domains often requires efficient, robust, decentralized, and adaptable approaches. Response threshold reinforcement is well suited for effective and stable decentralized task allocation, but its re-allocation capabilities necessary for dynamic environments have not been thoroughly investigated. Our goal is to analyze how this popular model responds to change and compare it to other decentralized task allocation methods.

Specialization can improve decentralized task allocation by reducing interference and potentially costly task switching, while increasing efficiency and adaptability to changing environments [3, 7, 8]. Scalability and robustness of task allocation can be further improved by using Large-scale Minimalist Multi-robot Systems [4], or swarms. Additionally, approaches that don't rely on inter-agent communication can further improve robustness and general applicability, by allowing for emergent cooperation in environments where communication is unreliable or not feasible [5].

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

Of particular relevance to this work is a biologically-inspired model for probabilistic decentralized task allocation, presented in [9]. To distinguish this model from related approaches, we refer to it as *StimHab*, as agents' probabilistic actions are driven by a combination of task **stim**uli and agents' **hab**ituation thresholds. Its effective specialization-based task allocation has been successfully applied to many real-world problems, but its re-specialization capabilities are not well tested. Additionally, while StimHab is widely used as part of other approaches (e.g. with job queue and dominance contests for factory job assignment [1]), its standalone capabilities, applications, and robustness are not yet well investigated.

We consider for what domains is StimHab uniquely well suited and investigate its task allocation and re-allocation capabilities under a variety of system changes. All work was done in maximally abstracted simulation, to establish a clear baseline without domain-specific influences (e.g. without allowing agents to rely on travel duration to continue gathering stimulus information between consecutive task selection steps). StimHab has demonstrated comparable behavior given 10 to 1000 agents. We will compare StimHab to Ant Colony Optimization (ACO) and Learning Automata (LA) approaches, by measuring proximity to the ideal system-wide task allocation after a change occurs, how quickly equilibrium is reached after each change, and how stable is the agents' assignment at that equilibrium as measured by the amount of task switching.

2 COMPLETED WORK

Two papers from our investigation have been published: (1) addressing the effects of task consideration order when agents employ StimHab to choose from all the tasks available to them [10] and (2) analyzing StimHab's specialization vs. respecialization behavior [6].

Existing task allocation techniques often assume the existence of discrete tasks that can be assigned to one agent at a time (e.g. market-based approaches [1]), or roles in team of agents responsible for each task (e.g. recruitment-based approaches [2]). In real-world domains, however, there are many tasks that are always available to the agents in unlimited quantities (e.g. patrolling, cleanup, gathering, etc.). Such domains preclude the use of market-based approaches and require a strategy for agents to repeatedly choose among multiple ubiquitously available tasks, while ensuring that the overall allocation is commensurate with system needs. In [10] we assess how the order in which tasks are considered by probabilistic agents affects task fulfillment, specialization tendencies, and system robustness. We employ an agent-based simulation where agents decide and act independently, do not communicate, and select tasks using the response threshold reinforcement model StimHab. The tested task ordering schemas are: (1) ascending agent thresholds, (2) descending task stimuli, (3) descending action probability combining thresholds and stimuli [9], and (4) a random ordering as a baseline for comparison. Results indicate that when resources/agents are plentiful, task ordering strategy has no effect on performance: the effort devoted to each task over time is proportional to task demands and an appropriate proportion of agents remain idle when available resources exceed overall demands, thus contributing to StimHab's robustness. Under more restrictive conditions, descending probability ordering leads to more stable task allocations.

Despite StimHab's popularity for decentralized task allocation, its respecialization capabilities have not been thoroughly assessed. Given StimHab's promising ability to handle multi-task environments, we conduct new experiments to quantify how this ability translates to dynamic domains. In [6], we showcase performance differences between specialization and respecialization under identical initial task demands and subsequent demand changes. Results clearly demonstrate that respecialization behavior leads to slower task allocation and a higher amount of unnecessary task switching. To understand the underlying cause of the observed differences, we (1) analyze the behavior driving task thresholds, task stimuli, and action probabilities; (2) observe how these values change during specialization; and (3) how the values change during respecialization. Our analysis shows that agent behavior near low-stimulus-low-habituation values negatively affects respecialization, due to a tug-of-war between "tasks that are no longer as needed, but the agents have grown habituated toward" and "tasks that are now needed more, but the agents have grown habituated against". Results show that respecialization can be problematic even when agents can easily specialize to initial demands, as the same forces that drive agents to specialize, can get in the way of respecialization.

3 WORK IN PROGRESS

The following two investigations are completed but not yet published: (1) we propose an automated detection mechanism to allow agents to respecialize more effectively when their environment changes, (2) we apply StimHab task allocation to a hierarchical domain to demonstrate its wider applicability to real-world problems.

As demands are likely to fluctuate in real-world domains, it is of interest to further investigate the conditions causing agents to fail to re-specialize, as well as to devise potential solutions. We propose a decentralized approach to improve respecialization by allowing agents to autonomously decide when to reset existing specialization thresholds and to respecialize from scratch. Results show that agents are able to accurately recognize changes in their environment by observing changes in task stimuli, detecting that current task specializations have become outdated. Automated specialization resets decrease the number of unnecessary task switches and allow for faster readaptation to changes in the environment.

To further test the applicability of StimHab to real-world problems, we apply it to dynamic area deployment defined as a hierarchical set of areas with changing demands (fig. 1). The new tests also verify StimHab's scalability to a larger number of always available tasks and a group of 1000 agents. We observe: (1) performance for every area over time, measured by how closely agents match the current deployment needs; (2) specialization quality by tracking the amount of area switches over time. Results show that while StimHab allows agents to self allocate proportionately, a 2% task switching is present after agents reach equilibrium and before the next change in demands. Upon closer analysis, this task switching is again a result of the aforementioned tug-of-war. Additional preliminary testing suggests that task allocation stability can be further improved by small changes to StimHab's probability formula.

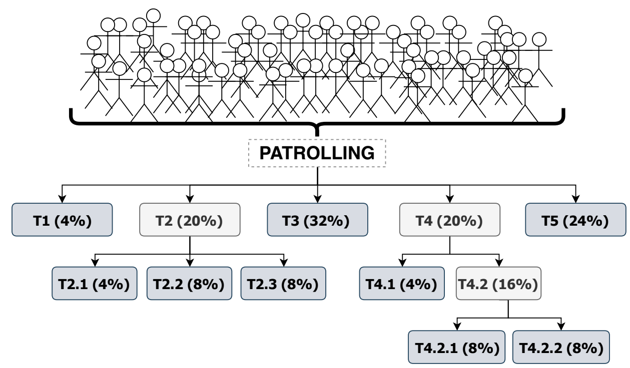

Figure 1: Hierarchical Deployment To Nested Areas
As demands for the areas T1-T5 and their sub-areas change (initial demands shown as % of the entire workforce), agents self-deploy based on performance values and their own evolved thresholds.

4 PROPOSED WORK

Our goal is to conduct in-depth testing of the robustness, benefits, and shortcomings of StimHab task allocation under changing environmental conditions. To this end, we will compare StimHab task allocation quality and stability to that of ACO and LA approaches. All three methods will be tested in simulated multiagent environments experiencing a variety of dynamic changes to the tasks and agents within the system. Below we list some of these experiments.

We propose a **systematic review of the effects of weights and exponents in StimHab probabilities** on the adaptability and stability of decentralized task allocation. Exponents have generally been used with a default value of '2' in existing literature, while weighting appears to have been generally omitted. Our preliminary investigation suggests that these variables may allow for improved stability of the task assignment at equilibrium.

To further assess the robustness of StimHab, we will test its behavior given a set of **environmental changes** and compare the results to those of ACO and LA. We propose testing a large quantity of tasks (100-1000) with 1000 or more agents, to demonstrate scalability. We will also test variations of the task set, where some tasks will be removed or added over time. We also propose a series of tests to assess robustness to changes of the agent team itself, such as agent removal or replacement with unspecialized units, or replacement with units limited to only performing a subset of the tasks. Additionally, prior testing demonstrates that StimHab can find appropriate allocation when supplied with exactly sufficient agents to fulfill all tasks. When given excess agents, the appropriate amount of agents idle, conserving resources until they are needed on new tasks or as replacements [10]. Preliminary testing suggests that when provided insufficient agents to fulfill task demands, StimHab leads to an unequal levels of deficiency across tasks. Future testing will be conducted to analyze ans possibly improve this behavior.

REFERENCES

[1] Vincent A. Cicirello and Stephen F. Smith. 2004. Wasp-like Agents for Distributed Factory Coordination. *AAMAS* 8, 3 (2004), 237–266.

[2] Fernando dos Santos and Ana LC Bazzan. 2009. An ant based algorithm for task allocation in large-scale and dynamic multiagent scenarios. In *GECCO*. 73–80.

[3] M Ani Hsieh, Ádám Halász, Ekin Dogus Cubuk, Samuel Schoenholz, and Alcherio Martinoli. 2009. Specialization as an optimal strategy under varying external conditions. In *ICRA*. 1941–1946.

[4] C Jones and MJ Mataric. 2003. Adaptive division of labor in large-scale minimalist multi-robot systems. In *IROS 2003*, Vol. 2. IEEE, 1969–1974.

[5] Anshul Kanakia, Behrouz Touri, and Nikolaus Correll. 2016. Modeling multi-robot task allocation with limited information as global game. *Swarm Intelligence* 10, 2 (2016), 147–160.

[6] Vera A. Kazakova and Annie S. Wu. 2018. Specialization vs. Re-Specialization: Effects of Hebbian Learning in a Dynamic Environment. In *FLAIRS-31*.

[7] Ling Li, Alcherio Martinoli, and Yaser S Abu-Mostafa. 2002. Emergent specialization in swarm systems. In *IDEAL*. Springer, 261–266.

[8] Jesús A Román, Sara Rodríguez, and Juan M Corchado. 2014. Improving intelligent systems: Specialization. In *PAAMS*. Springer, 378–385.

[9] Guy Theraulaz, Eric Bonabeau, and Jean-Louis Deneubourg. 1998. Response threshold reinforcement and division of labour in insect societies. *Proceedings of the Royal Society of London B* 265 (1998), 327–332.

[10] Annie S. Wu and Vera A. Kazakova. 2017. Effects of Task Consideration Order on Decentralized Task Allocation Using Time-Variant Response Thresholds. In *FLAIRS-30*. 466–471.