A Theoretical Framework For Self-Organized Task Allocation in Large Swarms

Doctoral Consortium

John Harwell
Department of Computer Science and Engineering, University of Minnesota
harwe006@umn.edu

ABSTRACT

Self-organized task allocation is possible with systems designed using the swarm robotic principles of scalability, flexibility, robustness, and emergence. We summarize (1) our derived quantitative measurements of these principles in 10,000 robot swarms, and (2) our task allocation work using stochastic choice and matroids. We propose extensions to our current task allocation methodology using stochastic processes and graph-theoretic topological invariants to provide a unified algorithmic approach to swarm-robotic foraging and construction.

ACM Reference Format:
John Harwell. 2020. A Theoretical Framework For Self-Organized Task Allocation in Large Swarms. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 2 pages.

1 PROBLEM DESCRIPTION

Swarm Robotics (SR) systems consist of large numbers of robots. The main differentiating factors between SR research and multi-agent robotics research stem from the origins of SR as an offshoot of Swarm Intelligence (SI), which investigates algorithms and problem solving techniques inspired from natural systems such as bees, ants, and termites. The main properties of SR systems are (1) scalability: systems can profitably scale to large numbers of agents, (2) emergence: simple, local robot interactions give rise to complex collective behaviors at the system level, (3) flexibility: systems are capable of exploiting (resisting) beneficial (adverse) environmental changes (4) robustness: systems are extremely tolerant to sensor/actuator noise and changing swarm sizes. As a result, SR systems are able to function effectively in domains where other types of robotic systems cannot.

The biological origins of SR enables effective parallels to be drawn between many naturally occurring problems, such as foraging, and real-world problems including clearing a corridor on a mining operation, hazardous material cleanup, and search and rescue [5]. In the foraging problem, robots are tasked with gathering objects from the environment and bringing them to a single central location under various conditions/parameters, in which the swarm collectively adapts to maximize some performance measure. Once collected, objects can be used to create 3D structures of arbitrary size and complexity, drawing inspiration from the strategies employed by termites [7, 9].

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

2 CONTRIBUTIONS

Our recent works [2–4] have attempted to move beyond the heuristic approaches to problem solving common in SR to more theoretically grounded algorithms. In [2], we investigate task partitioning in the context of a foraging task (Fig. 1). *Generalist* robots choose to not partition the task, and perform the entire foraging task themselves (depth0 decomposition), retrieving an object and bringing it to the nest. Alternatively, robots can choose to partition the task once (depth1 decomposition), and perform one of two subtasks (Fig. 1, bottom gold tasks). *Harvester* robots can bring an acquired object to an intermediate drop site (cache), and *Collector* robots pick up an object from the cache and bring it to the nest. Robots make stochastic task allocation decisions based on local, per-robot estimates of task execution times, and we observe emergent behaviors in which the swarm dynamically changes its task allocation distribution in predictable ways.

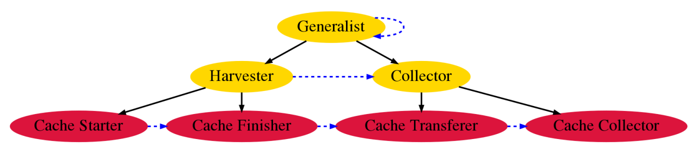

Figure 1: Task decomposition graph for a foraging task. Solid edges indicate task decomposition, dashed blue edges indicate dependent task sequences. Gold tasks were defined in previous works [1, 2, 8], red tasks defined in [4].

In [3], we propose a set of quantitative measurements for swarm scalability, flexibility, and emergence. We apply our measurements to a foraging task using the correctness of predictive hypotheses across scales and operating conditions to demonstrate their utility as part of the iterative SR system design process. We evaluated candidate algorithms at "natural scales" (10,000 robots), as distinct from the very small numbers (≤ 30) of robots common in many related works. Our scalability measure notes the parallels between SR and supercomputing, and frames a swarm's scalability in terms of the fraction of its performance that was due to inter-robot cooperation. To measure emergence, we increase swarm sizes within fixed size operating areas, and measure levels of inter-robot interference across scales, observing that sub-linearities in interference increases indicate self-organization (i.e. in the absence of such organization, doubling the swarm size in a confined space would result in double the interference). Finally, we analyzed swarm performance curves in ideal vs. non-ideal conditions using mathematical measures of curve similarity in order to provide quantitative evaluations of swarm flexibility across experimental scenarios.

In [4], we extend the task partitioning method from [2, 8] recursively (i.e., a large overall task decomposed in different ways into sequences of interdependent subtasks), allowing the swarm to collectively create, utilize, and deplete caches dynamically, rather than using a static predetermined location (Fig. 1, red tasks). We prove that a variant of this task decomposition graph is a matroid [6], and therefore optimal performance is guaranteed via a simple greedy algorithm *iff* the emergent intelligence was rooted in collective learning of vertex weights (task cost estimates) of the original graph. We do not observe optimal performance, suggesting emergent intelligence is instead rooted in collective learning of task decomposition graph connectivity.

3 FUTURE WORK

Extending [3], we will first develop a metric for robustness using mathematical curve similarity measures. We will also perform a correlation analysis between common swarm convergence measures (when the collective search process to solve a problem has stabilized) and observed performance, in order to determine which types of measures are best suited to different application domains. Such convergence measures include (1) positional entropy (i.e., is the swarm's spatial distribution generally stable over time?), (2) robot nearest neighbors (i.e., have robots generally grouped themselves into stable clusters over time?), (3) task allocation distribution (i.e., is the fraction of the swarm engaged in each task stable over time?).

We will refine the stochastic task partitioning algorithm presented in [4], investigating task distribution load balancing within the task decomposition graph, now represented as a Markov Chain. Each robot will compute the projected steady state distribution of the Markov Chain from its own localized information about task cost estimates, and choose its next task based on (1) what task it has most recently executed, (2) which task of the parents and/or children of its most recently executed task has a projected steady-state distribution that is most dissimilar from its predecessors/descendants. The goal of this algorithm will be to minimize unmet dependencies between tasks in task sequences (dashed blue edges in Fig. 1) which are being worked on by (potentially) different robots, thus reducing wait times and improving collective efficiency.

Similar to previous work in foraging, recent works on swarm based construction have largely been heuristic in nature [7, 9], and therefore few mathematical performance or convergence guarantees of algorithms exist. We begin to address this gap by defining a new task decomposition graph which breaks down the overall *Construction* task (i.e. locating a block and placing it on an in-progress structure) into the *Forager* graph shown in Fig. 1, and the *Builder* decomposition graph, which breaks down the task of taking a block from a specific cache and placing it onto the in-progress structure. We will demonstrate the generality and extensibility of our task allocation methodology by applying the derived theoretical tools from our previous works to this much richer and more complex task decomposition graph. This stochastic foraging-construction system will be evaluated at truly real-world scales of 1,000,000 robots, and will be capable of efficiently building simple but useful structures such as walls and ramps. If successful, the system will introduce the beginning of a theoretical basis for automation in building basic infrastructure in real-world applications.

To create such a system, we will extend the stochastic task allocation algorithms of the *Forager* task from [4] by modeling caches as queues, and applying queueing theory to influence task allocation decisions to maintain a minimum number of objects in caches in steady state. These stable piles of heterogeneous objects will then be used by robots engaged in the *Builder* task. Stochastic processes will be used to model the arrival of blocks within the swarm's operating area, as well as track their progress from their original location to one of the existing caches in the arena (dynamically created by the swarm), so that overall construction rate can be accurately predicted as the swarm's task distribution converges to a steady state.

The *Builder* task will be designed to maintain topological invariants of the graphical representation of the structure to build (e.g., no vertical holes in the graph, no concavities). The invariants will be chosen such that the derived algorithm is provably correct for all possible sub-graphs, and that construction is guaranteed regardless of the number of robots or the current state of progress; that is, the subgraph containing (possibly) heterogeneous blocks comprising the in-progress structure also maintains the invariants. Furthermore, the *Builder* algorithm will be massively parallel, allowing multiple block attachment sites to be active simultaneously (as distinct from [7], which although it allowed multiple robots to traverse the in-progress structure in parallel, only had a single active block attachment site). This will be achieved through recursive decomposition of the overall graph and *Builder* task into multiple subtasks and construction lanes, each with their own ingress and egress points.

ACKNOWLEDGEMENTS

Support is gratefully acknowledged from Amazon Robotics, MnDRIVE RSAM at the University of Minnesota, and the Minnesota Supercomputing Institute.

REFERENCES

[1] Arne Brutschy, Giovanni Pini, Carlo Pinciroli, Mauro Birattari, and Marco Dorigo. 2014. Self-organized task allocation to sequentially interdependent tasks in swarm robotics. *Autonomous Agents and Multi-Agent Systems* 28, 1 (2014), 101–125.

[2] John Harwell and Maria Gini. 2018. Broadening applicability of swarm-robotic foraging through constraint relaxation. *IEEE Int'l Conf. on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)* (May 2018), 116–122.

[3] John Harwell and Maria Gini. 2019. Swarm Engineering Through Quantitative Measurement of Swarm Robotic Principles in a 10,000 Robot Swarm. *International Joint Conference on Artificial Intelligence (IJCAI)* (August 2019).

[4] John Harwell, London Lowmanstone, and Maria Gini. 2020. Demystifying Emergent Intelligence And Its Effect On Performance In Large Robot Swarms. *Autonomous Agents and Multi-Agent Systems (AAMAS)* (May 2020).

[5] Joshua P. Hecker and Melanie E. Moses. 2015. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence* 9, 1 (2015), 43–70.

[6] James G. Oxley. 2006. *Matroid Theory (Oxford Graduate Texts in Mathematics).* Oxford University Press, Inc., New York, NY, USA.

[7] Kirstin Petersen, Radhika Nagpal, and Justin Werfel. 2011. TERMES: An Autonomous Robotic System for Three-Dimensional Collective Construction. In *Robotics: Science and Systems (RSS).*

[8] Giovanni Pini, Arne Brutschy, Mauro Birattari, and Marco Dorigo. 2011. Task Partitioning in Swarms of Robots: Reducing Performance Losses Due to Interference at Shared Resources. In *Informatics in Control Automation and Robotics.* Springer, 217–228.

[9] Justin Werfel, Kirstin Petersen, and Radhika Nagpal. 2011. Distributed multi-robot algorithms for the TERMES 3D collective construction system. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems.* 1–6.