

A Unified Mathematical Approach for Foraging and Construction Systems in a 1,000,000 Robot Swarm

John Harwell

Department of Computer Science, University of Minnesota
harwe006@umn.edu

Abstract

Automation in construction is possible with systems designed using the swarm robotic principles of scalability, flexibility, robustness, and emergence. We derive quantitative measurements of these principles in 10,000 robot swarms as a first step in achieving this goal. We summarize our recent task allocation work in the context of an object gathering task and demonstrate its feasibility in the context of automated construction tasks. We present a trajectory to extend our current task allocation methodology using stochastic processes in order to present a unified approach to task allocation in swarm-robotic construction.

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 are (1) scalability: systems can profitably scale to large numbers of agents, (2) emergence: simple, local robot interactions give rise to complex collective behaviors, (3) flexibility: systems are capable of exploiting (resisting) beneficial (adverse) environmental changes (4) robustness: systems are extremely tolerant to sensor/actuator noise and loss of robots. 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 [Hecker and Moses, 2015]. 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.

2 Contributions

Our recent works [Harwell and Gini, 2018; Harwell and Gini, 2019] have attempted to move beyond the historical heuristic

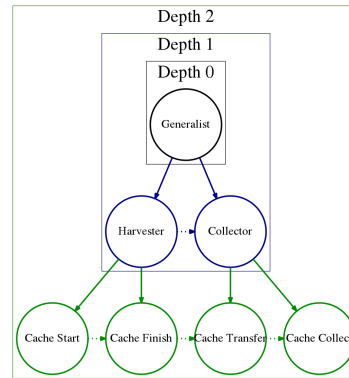


Figure 1: Depth2 (recursive) task decomposition graph for a foraging task. Sequential execution of the tasks from left to right at a given depth from the root is equivalent to execution of the tasks at all higher levels.

approaches to problem solving common in SR to more theoretically grounded algorithms. In [Harwell and Gini, 2018], 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 sub-tasks (Fig. 1, blue 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.

In [Harwell and Gini, 2019], we further address this gap by proposing a set of quantitative measurements of swarm scalability, flexibility, and emergence, and demonstrate their applicability to the SR system design process on an example foraging task. 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 double the interference as well). Finally, we analyzed similarity between swarm performance curves in ideal vs. non-ideal conditions using mathematical measures of curve similarity in order to provide mathematically grounded, rather than qualitative, evaluations of swarm flexibility across experimental scenarios. We evaluated candidate solutions at “natural scales” (10,000 robots), as distinct from the very small numbers (≤ 30) of robots common in many related works. We demonstrated the predictive insights possible using our proposed methodology during an iterative system design process as algorithms are scaled up/modified to meet target application scale and operating conditions.

3 Future Work

Extending [Harwell and Gini, 2019], we will develop a metric for robustness using mathematical curve similarity measures in order to provide a rich set of design tools for swarm engineering. We will further analyze swarm behaviors in order to derive similar measurements of swarm convergence (when the collective search process to solve a problem has stabilized), along the axes of (1) measurement of 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 continue to explore applications directly applicable to real-world problems such as those faced by Amazon Robotics in their fulfillment warehouses, performing an analysis of the following aspects of foraging scenarios: (1) object spatial distributions, (2) object priorities, (3) obstacle filled environments, drawing on [Harwell and Gini, 2019] to define measures of similarity for obstacle conditions, (4) the effect swarm density (ratio of robots to operating area) has on various aspects of performance. Study of these aspects will provide a platform for other researchers to develop more theoretically grounded approaches to foraging tasks by enabling greater insight into *why* a given method performs well/poorly under specified conditions.

We will extend our task partitioning method [Harwell and Gini, 2018] recursively, allowing the swarm to perform multiple partitionings of the overall foraging task and collectively create caches at any location in the arena, rather than using a static predetermined location (Fig. 1, blue+green tasks). Individual robots will solve an optimization problem maximizing the projected utility of a potential cache site in relation to other nearby caches and known objects. Overall task selection will be performed according to swarm task distribution load balancing, in which each robot computes the projected steady state distribution of the Markov Chain which connects the different tasks in the swarm’s task decomposition graph together, and chooses its next task based on (1) what task it most recently executed, (2) which task of its last task’s parents and/or children has a steady-state distribution that is most

dissimilar from its predecessors/descendants.

Beyond foraging, we will demonstrate the generality of our task allocation methodology by extending it to swarm-based construction. Similar to previous work in foraging, recent works on swarm based construction has largely been heuristic in nature [Petersen *et al.*, 2011], and therefore few mathematical performance or convergence guarantees of algorithms exist. We begin to address this gap with a stochastic foraging-construction SR system consisting 1,000,000 robots, which will be capable of building simple but useful structures such as walls and ramps. If successful, this work will introduce the beginning of a theoretical basis for automation in building basic infrastructure in real-world applications.

We will apply queuing theory to design stochastic algorithms for swarms to maintain a specified level of objects in caches, as a precondition for providing stable piles of heterogeneous objects for use in construction tasks. We will then define 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 somewhere), into the 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 an in-progress structure. Builders will use stochastic means to determine where to place a block they have obtained from a cache near the construction site.

Mathematical predictions of construction progress will be computed from (1) localized steady-state task distributions of the swarm’s task decomposition graph, (2) results of queuing theory analysis of caches near the construction site. The resulting system will then be evaluated at natural scales on the similarity between the mathematical predictions of construction progress on a series of simple walls and ramps.

Acknowledgements

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

References

- [Harwell and Gini, 2018] John Harwell and Maria Gini. Broadening applicability of swarm-robotic foraging through constraint relaxation. *IEEE Int’l Conf. on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN)*, pages 116–122, May 2018.
- [Harwell and Gini, 2019] John Harwell and Maria Gini. Swarm engineering through quantitative measurement of swarm robotic principles in a 10,000 robot swarm. arXiv:1907.03880, 2019.
- [Hecker and Moses, 2015] Joshua P. Hecker and Melanie E. Moses. Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intelligence*, 9(1):43–70, 2015.
- [Petersen *et al.*, 2011] Kirstin Petersen, Radhika Nagpal, and Justin Werfel. TERMES: An autonomous robotic system for three-dimensional collective construction. In *Robotics: Science and Systems (RSS)*, 2011.