

# Brief Announcement: Distributed Task Allocation in Ant Colonies<sup>\*</sup>

Anna Dornhaus<sup>1</sup>, Nancy Lynch<sup>2</sup>, Tsvetomira Radeva<sup>2</sup>, and Hsin-Hao Su<sup>3</sup>

<sup>1</sup> University of Arizona

<sup>2</sup> Massachusetts Institute of Technology

<sup>3</sup> University of Michigan

A common problem in both distributed computing and insect biology is designing a model that accurately captures the behavior of a given distributed system or an ant colony, respectively. While the challenges involved in modeling computer systems and ant colonies are quite different from each other, a common approach is to explore multiple variations of different models and compare the results in terms of the simplicity of the model and the quality of the results. We consider the task allocation problem as a case study and explore multiple models inspired from both distributed computing and biological experiments. We compare the models with respect to their significance in understanding real ant behavior and also their technical relevance to distributed computing.

**Task Allocation:** In ant colonies, the task allocation problem is a distributed assignment of ants to tasks with the goal of satisfying the demands of all tasks. The first attempt at modeling the task allocation problem from a distributed computing perspective was in [2], where the authors show that the ants can solve the task allocation problem in  $O(|T| \log |A|)$  rounds, where  $A$  is the set of ants and  $T$  is the set of tasks. Biologists have also modeled the ant task allocation process from a distributed perspective by designing various models [5, 6] that try to match the actual ant behavior.

**Summary of results:** We consider two families of models based on the type of input ants receive from the environment. In our first family of models, each ant learns from the environment (1) whether it is successful at the current task it is working on, and (2) a new task it can start working on if it is idle or unsuccessful at its old task. For (1), we consider a function that ensures the number of successful ants working on a given task is no more than the demand for the task. For (2), we consider different options, ranging from a uniformly random task to a task chosen based on the proportion of ants already working on it. We show that, depending on the choice of this function, the running time of the resulting task allocation process ranges from  $O(\log |T|)$  to  $O(|T| \log |A|)$  rounds, also proving a better time bound for the algorithm in [2].

The second family of models we consider captures the individual variation in the work units each ant provides to different tasks. Task allocation with individual variation is NP-hard; we provide a simple mechanism to *approximately* satisfy the demands of each task (assuming this is possible). We show that after

---

<sup>\*</sup> This work is sponsored in part by NSF grants BIO-1455983, CCF-0939370, CCF-1217338, CNS-1318294 and AFOSR grant FA9550-13-1-0042.

$O(|A|^{1-\epsilon})$  rounds, the ants converge to a solution that satisfies the demands with an  $O(W|A|^{1/2+\epsilon})$  additive error, where  $W$  is the ratio between the largest and the smallest number of work units provided by the ants. In each round, each ant switches to the current most promising task with some probability, and that probability diminishes in each subsequent round. The current most promising task for a given ant is the task with the largest deficit (the difference between the demand and the work provided already) weighted by the work units the ant is capable of providing for the task. The main technique in our analysis is derived from the multiplicative weight update method for solving linear programs [1, 7] with modifications to accommodate the limited capabilities of ants. We conjecture that the above technique has potential applications outside the ant world; for example, task allocation among non-communicating agents with individual variation and a global view of task deficits.

**Contribution to biology:** One goal of our analysis is to show that if task allocation is allowed to be approximate, it need not become significantly more difficult with larger colony sizes; this is supported by the fact that, in our results, the convergence time of the task allocation process depends only logarithmically, or does not depend at all, on  $|A|$ . This conclusion is not obvious from prior theory results (e.g. [2]), which contradicted the notion of empiricists that larger colonies perform better at task allocation [4].

Our results also provide a novel hypothesis for the existence of idle ants: we show that the task allocation process is faster when there are extra ants compared to the case where the number of ants is very close to the total sum of demands. While the existence of idle ants is supported by empirical evidence [3], biologists do not have an adequate explanation for this behavior. These general observations make our results broadly relevant for understanding the evolution of division of labor in biological systems.

## References

1. S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
2. A. Cornejo, A. R. Dornhaus, N. A. Lynch, and R. Nagpal. Task allocation in ant colonies. In *Proc. 28th Symposium on Distributed Computing*, pages 46–60, 2014.
3. Anna Dornhaus, Jo-Anne Holley, Victoria G. Pook, Gemma Worswick, and Nigel R. Franks. Why do not all workers work? colony size and workload during emigrations in the ant *temnothorax albipennis*. *Behavioral Ecology and Sociobiology*, 63(1):43–51, 2008.
4. Anna Dornhaus, Scott Powell, and Sarah Bengston. Group size and its effects on collective organization. *Annual review of entomology*, 57:123–141, 2012.
5. Deborah M. Gordon, Brian C. Goodwin, and L.E.H. Trainor. A parallel distributed model of the behaviour of ant colonies. *Journal of Theoretical Biology*, 156(3):293–307, 1992.
6. S. W. Pacala, D. M. Gordon, and H. C. J. Godfray. Effects of social group size on information transfer and task allocation. *Evol Ecol*, 10(2):127–165, 1996.
7. N. E. Young. Randomized rounding without solving the linear program. In *Proc. 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 170–178, 1995.