

Brief Announcement: On Site Fidelity and the Price of Ignorance in Swarm Robotic Central Place Foraging Algorithms

Abhinav Aggarwal
University of New Mexico
Albuquerque, NM
abhiag@unm.edu

Diksha Gupta
University of New Mexico
Albuquerque, NM
dgupta@unm.edu

G. Matthew Fricke
University of New Mexico
Albuquerque, NM
mfricke@unm.edu

Melanie E. Moses
University of New Mexico
Albuquerque, NM
melaniem@unm.edu

ABSTRACT

A key factor limiting the performance of central place foraging algorithms is the awareness of the agent(s) about the location of food items around the nest. We study the ratio of how much time an ignorant agent takes relative to an omniscient forager for complete collection of food items in the arena. This effectively quantifies the penalty each algorithm pays for not knowing (or choosing to ignore information gained about) where the resources are located. We model the effect of depletion of food items from the arena on the foraging efficiency over time and analytically verify that returning to the location of the last food item found strongly helps in counteracting this effect. To the best of our knowledge, these results have only been empirically argued so far.

CCS CONCEPTS

• **Computing methodologies** → **Bio-inspired approaches; Distributed algorithms**; Continuous space search;

KEYWORDS

Central Place Foraging, Site Fidelity, Price of Ignorance, Bio-Inspired Computation

ACM Reference Format:

Abhinav Aggarwal, G. Matthew Fricke, Diksha Gupta, and Melanie E. Moses. 2019. Brief Announcement: On Site Fidelity and the Price of Ignorance in Swarm Robotic Central Place Foraging Algorithms. In *2019 ACM Symposium on Principles of Distributed Computing (PODC '19)*, July 29–August 2, 2019, Toronto, ON, Canada. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3293611.3331572>

1 INTRODUCTION

Foraging of food items around a central location is a well motivated problem in understanding biological systems like ant colonies as well as gaining theoretical insight into distributed collaborative search algorithms. To better understand the effects of overlaps in

search trajectories and depletion of food items from the arena, we define the *Price of Ignorance* as a measure of how much time an ignorant agent takes relative to an omniscient forager for complete collection of food items in an arena.

We analyze two algorithms: a deterministic **Distributed Archimedes Spiral Algorithm** (DASA) [4, 7, 12] in which the agents move along interleaved Archimedes spirals [10] around the nest; and a randomized **Ballistic Central Place Foraging Algorithm** (BALCPFA) [15] in which the agents move radially away from the nest until they find a food item or hit the arena boundary. Upon returning to the nest, each agent independently chooses a uniformly random direction for the next ballistic run.

Our analysis shows that preferential foraging around the location of the last found food item (also called *site fidelity*) significantly lowers the foraging time whereas food depletion and large overlaps in search trajectories have an adverse effect. This interplay of exploration vs. exploitation, thus, defines the quality of foraging.

Our Model. We assume a continuous circular arena of radius R around a fixed central location, called the *nest*¹. This arena consists of F individually collectible food items that are arbitrarily clustered into piles. We assume that each pile is placed independently and uniformly at random in the arena². Our system consists of N agents that move synchronously at a constant unit speed and have a vision radius R_v , *i.e.*, the food item can be detected by any agent that is within a distance of R_v from it. We assume that the agents know R and have unique IDs in $\{1, \dots, N\}$. However, the locations of the piles and the number of food items in each pile is unknown to the agents. The agents can navigate back to the nest from any location in the arena but cannot remember more than $\Theta(1)$ previously visited locations. We assume that each agent can pick up only one food item at a time and does not communicate with any other agent during the algorithm. We do not model unexpected environmental effects, agent failures, and external noise in this paper.

Site Fidelity. In BALCPFA, we say that an agent employs site fidelity with parameter $p_s \in [0, 1]$ if after returning to the nest with

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PODC '19, July 29–August 2, 2019, Toronto, ON, Canada

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6217-7/19/07.

<https://doi.org/10.1145/3293611.3331572>

¹While most existing work assumes an infinite discrete grid [6, 14], we note that assuming a continuous arena with Euclidean distances only changes the runtime of our algorithms by a factor of $\sqrt{2}$ compared to the Manhattan distance. The assumption of finite arena is biologically motivated since most ant species do not travel very large distances before giving up their search [8, 9].

²We compare our model and results with those for the adversarial setting considered in the ANTS model by Feinerman and Korman [6] in Section 3.

a food item, with probability p_s , the agent chooses to return to the location of this item for the next ballistic run. Otherwise, the agent moves along a uniformly random direction.

Omniscient Foraging. Our analysis is relative to an omniscient foraging scenario in which each agent is initially equipped with the exact locations of all food items in the arena and given perfect navigation capabilities to forage these items in the smallest possible time. This ideal algorithm assumes information about the location of the next available food item that no other agent is currently scheduled/planning to collect. The expected number of time steps until the last food item is brought to the nest by some agent is $4RF/3N$. This is because each pile is at an expected distance of $2R/3$ from the nest. Thus, it takes $2(2R/3) = 4R/3$ time steps to collect each food item.

The Price of Ignorance. With respect to the ideal foraging described above, we define the price of ignorance, denoted χ , of a given foraging algorithm A as $\chi(A) = 3NT_A/(4RF)$, where T_A is the number of time steps until some agent brings the last food item to the nest.³ This metric helps quantify the cost of being oblivious to the (or choosing to ignore any gained) knowledge about the resources and their location in the arena.⁴

2 OVERVIEW OF OUR RESULTS

We now summarize our main results and present an overview of our approach, wherever possible. We refer the reader to the full version of our paper [1] for more details of our proofs.

Cooperative Foraging Speeds up Deterministic Search. We prove that the total time taken by DASA is $O\left(\frac{R^2}{N} + R\left\lceil\frac{F}{N}\right\rceil\right)$ and that $\chi(\text{DASA}) < \left\lceil\frac{5}{3R_v F} - \frac{N}{2RF}\right\rceil R + \sqrt{2}$. This implies that increasing the number of agents helps reduce the foraging time for DASA. We attribute the improvement in the performance with multiple agents to the fact that the search trajectories are non-overlapping and the location of each food item found helps to continuously eliminate part of the arena for the search that follows (see Figure 1). Additionally, we show that $\chi(\text{DASA}) \geq \left\lceil\frac{\pi}{8R_v F} - \frac{9N}{4RF}\right\rceil R + 1$, implying that the bound above is tight up to constant factors.

BALCPFA has a Higher Foraging Time. BALCPFA can be shown to have a higher foraging time compared to DASA, with high probability, using a result about random arc distances on a unit circle by King and Saia [11]. They prove that when M points are placed uniformly at random on the circumference of a unit circle, then the maximum arc distance between any two neighboring points is $\Theta(\log M/M)$ and the minimum distance is $\Theta(1/M^2)$, with probability at least $1 - O(1/M)$.

³We refer to this as the time to complete collection for algorithm A . Note that we do not address the problem of protocol termination in this paper and measure the time only until the last food item is collected.

⁴We note that this measure is similar to the notion of *competitiveness* for online algorithms [5], which has been used in the context of foraging in [6]. However, in our application, we do not measure competitiveness with respect to the amount of advice available for the value of N (as in [6]). Rather, from a biological and engineering perspective, we measure competitiveness with respect to the knowledge of the distribution of resources around the nest and in the decision of the foraging algorithm to avoid overlaps in the search trajectories. Additionally, we assume a uniformly random pile placement, which is significantly different from the adversarial setting often for online algorithms.

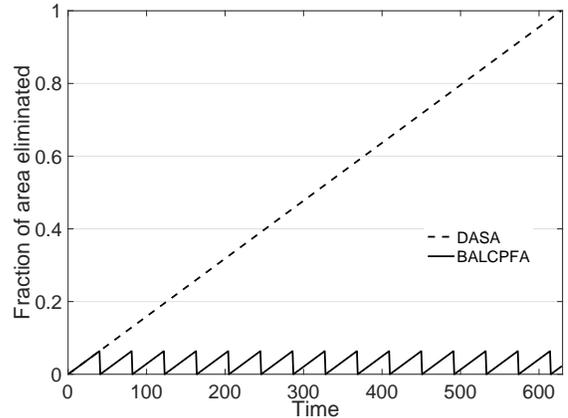


Figure 1: Plot of the fraction of area eliminated by a single agent ($R_v = 1, R = 20$) vs. search time while foraging the arena using DASA and BALCPFA.

For a single adversarially placed pile of diameter Δ inside the arena, the number of ballistic runs required to locate this pile is $\Omega(\Delta \log \Delta)$, with high probability. Thus, a total of $\Omega\left(R\left\lceil\frac{F\Delta \log \Delta}{R_v N}\right\rceil\right)$ time steps are needed for complete collection of food items from this pile. Clearly, this is higher than that of DASA.

When the pile is placed uniformly at random instead, we show that when agents do not deploy site fidelity (*i.e.* when $p_s = 0$), it holds that $\chi(\text{BALCPFA}) = \Omega(\lceil R/R_v \rceil)$, which, asymptotically, is at least F times higher than the upper bound for DASA. Intuitively, this is because BALCPFA chooses to completely ignore all the information obtained while collecting the first food item while it forages for the second one. Thus, each food item is essentially searched as if it were independently and uniquely placed in the arena, causing large overlaps in the search trajectory (see Figure 1).

Food Depletion has Adverse Effects. If the arena contains multiple piles, the time to complete collection is affected by the rate of depletion of these piles. This is because as these piles become smaller (due to collection of food items in them) and their number decreases, it becomes more challenging for the agents to find the next available food item. To the best of our knowledge, we make the first attempt to model the effect of such a behavior in collaborative foraging and prove that when food depletion is taken into account, one can expect a multiplicative blowup up to $\Omega(MN)$ in the price of ignorance, where M is the number of piles. We model this behavior using a probabilistic recurrence relation and compute the expected foraging time for BALCPFA with depleting piles to obtain our result [1]. Intuitively, while a larger number of agents on one hand collect food items faster (by working in parallel), they also deplete the arena at a faster rate, thus, significantly slowing down locating the remaining food items as time progresses.

Site Fidelity Helps Keep Randomized Foraging Efficient. Our final result proves that deploying site fidelity strongly counteracts the effect of food depletion and helps keep BALCPFA (and thus, other randomized foraging schemes) efficient. This has mostly been empirically argued so far [13] and our result provides one of the first

theoretical evidence into this phenomenon. Intuitively, site fidelity is the behavior of the agent in which it chooses to return to the location of the last food item it found. We show that when the agents deploy perfect site fidelity (*i.e.* with $p_s = 1$), the expected price of ignorance of BALCPFA is at most $\left\lceil \frac{\pi}{R_v} - \frac{N}{2FR} \right\rceil R$, which is a significant improvement over the bound obtained above. Note that this does not imply that the agent only moves along one direction indefinitely. When the pile is exhausted, the ballistic run in this direction will go all the way to the arena boundary and return to the nest without any food item in sight. This will cause the agent to move along a random direction for the next ballistic run since site fidelity only applies when a food item is found.

Communication Enhances Foraging Efficiency. We further argue that communication among the agents, and not only working in parallel, can help collect the food items faster. Similar to the approach by Steels [16] using *crumbs* and that by Hecker and Moses [9] using virtual waypoints, we envision a pedagogical scenario with BALCPFA in which agents have a special marker which they can place in the arena at any time. This marker is only visible to the nearby agents and once detected, indicates a desired direction to them. Assuming only a single pile in the arena, when some agent finds a food item in this pile, it places the marker near the nest in the correct direction and all other agents then move along this direction to collect the remaining food items. To further optimize the algorithm, all agents can search the first food item in parallel and the one that finds it first places the marker. The expected distance covered in this marker-based algorithm is $O\left(\frac{R^2}{N} + R\left\lceil \frac{F}{N} \right\rceil\right)$, which closes the gap between the asymptotic performance of BALCPFA and DASA.

Note that this form of communication is different from deploying site fidelity since only one agent returns to the last visited location in the latter. In biological systems, such a communication has been conjectured to happen by laying pheromones that the other agents can pick up and follow [9].

3 DISCUSSION

In this section, we briefly compare our approach to the work by Feinerman and Korman [6] for the ANTS problem. Taking motivation from the distribution resources in the wild [8], we assume a random food placement in the arena in contrast to the adversarial setting considered in [6]. Moreover, in the adversarial case, if the adversary is allowed to place the food items in an online manner, a lower bound of $\Omega\left(R\left\lceil F + \frac{RF}{N} \right\rceil\right)$ can be proved for complete collection of food items in the arena using Yao’s minimax principle [17] and the Feinerman-Korman bound in [6]. In our case, however, we show that DASA takes only $O\left(\frac{R^2}{N} + R\left\lceil \frac{F}{N} \right\rceil\right)$ time in expectation for complete collection, which is significantly better than the bound above. Additionally, while only the time to locate the first food item is considered in [6], we study the effect of clustering of food items into piles and the time for complete collection of resources. In fact, it can be shown that when the food items are clustered into a pile of diameter Δ , then even under adversarial placement, there exists a single-agent deterministic foraging strategy which can locate this pile in $\Theta(R^2/\Delta)$ time steps [2, 3].

From a biological perspective, we provide one of the first analytical insights into naturally observed phenomena of site fidelity [13] and argue in favor of determinism for collaborative foraging. The formal analysis shows that deterministic spirals are best in theory, but simulations and our prior work show that stochastic algorithms can work as well in practice, but with significantly higher variance. Moreover, from a theoretical perspective, our analysis helps quantify the importance of keeping overlaps in search trajectories to a minimum, which is an argument in favor of deterministic search for limited-memory systems (which was also argued empirically by [7, 12, 15]). For stochastic search, our results favor walks with site fidelity as a solution to balancing the effects of spatial sparsity that arises over time (due to depletion).

ACKNOWLEDGMENTS

This work is supported by the NASA MUREP, NNX15AM14A Swarathon Cooperative Agreement and the James S. McDonnell Foundation, Complex Systems Scholar Award. We also thank Jared Saia, Department of Computer Science, William F. Vining, Moses Biological Computation Lab, University of New Mexico, and David Peleg, Weizmann Institute of Science for their comments and discussions.

REFERENCES

- [1] A Aggarwal, D Gupta, W F Vining, G M Fricke, and M E Moses. 2019. Ignorance is Not Bliss: An Analysis of Central-Place Foraging Algorithms. *Manuscript Under Submission*. Available at <https://bit.ly/2QcPaX8> (2019).
- [2] A Aggarwal and J Saia. 2019. Robustly Finding a Cluster Using the Golden Ratio. *Manuscript Under Submission*. (2019).
- [3] A Aggarwal, W F Vining, D Gupta, J Saia, and M E Moses. 2019. A Most Irrational Foraging Algorithm. *Manuscript Under Submission*. (2019).
- [4] Ricardo Baeza-Yates and René Schott. 1995. Parallel searching in the plane. *Computational Geometry* 5, 3 (1995), 143–154.
- [5] Allan Borodin and Ran El-Yaniv. 2005. *Online computation and competitive analysis*. Cambridge University Press (2005).
- [6] Ofer Feinerman and Amos Korman. 2017. The ANTS problem. *Distributed Computing* 30, 3 (2017), 149–168.
- [7] G Matthew Fricke, Joshua P Hecker, Antonio D Griego, Linh T Tran, and Melanie E Moses. 2016. A distributed deterministic spiral search algorithm for swarms. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 4430–4436.
- [8] Deborah M Gordon and Alan W Kulig. 1996. Founding, foraging, and fighting: colony size and the spatial distribution of harvester ant nests. *Ecology* 77, 8 (1996), 2393–2409.
- [9] 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.
- [10] HC Holland. 1957. The archimedes spiral. *Nature* 179, 4556 (1957), 432.
- [11] Valerie King and Jared Saia. 2004. Choosing a random peer. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*. ACM, 125–130.
- [12] Malika Meghjani, Sandeep Manjanna, and Gregory Dudek. 2016. Multi-target rendezvous search. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2596–2603.
- [13] Danielle P Mersch, Alessandro Crespi, and Laurent Keller. 2013. Tracking individuals shows spatial fidelity is a key regulator of ant social organization. *Science* 340, 6136 (2013), 1090–1093.
- [14] Cameron Musco, Hsin-Hao Su, and Nancy Lynch. 2016. Ant-inspired density estimation via random walks. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*. ACM, 469–478.
- [15] Lu Qi, Antonio D. Griego, G. Matthew Fricke, and Melanie E Moses. 2019. Comparing Physical and Simulated Performance of a Deterministic and a Bio-inspired Stochastic Foraging Strategy for Robot Swarms. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*. IEEE.
- [16] Luc Steels. 1990. Cooperation between distributed agents through self-organisation. In *Intelligent Robots and Systems ’90. Towards a New Frontier of Applications, Proceedings. IROS’90. IEEE International Workshop on*. IEEE, 8–14.
- [17] Andrew Chi-Chin Yao. 1977. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. IEEE, 222–227.