Ignorance is Not Bliss: An Analysis of Central-Place Foraging Algorithms

Abhinav Aggarwal,[1] Diksha Gupta,[1] William F. Vining,[1] G. Matthew Fricke,[*,1,2] and Melanie E. Moses[1,3,4]

***Abstract*— Central-place foraging (CPF) is a canonical task in collective robotics with applications to planetary exploration, automated mining, warehousing, and search and rescue operations. We compare the performance of three Central Place Foraging Algorithms (CPFAs), variants of which have been shown to work well in real robots: spiral-based, rotating-spoke, and random-ballistic. To understand the difference in performance between these CPFAs, we define the *Price of Ignorance* and show how this metric explains our previously published empirical results. We obtain upper-bounds for expected complete collection times for each algorithm and evaluate their performance in simulation. We show that site-fidelity (i.e. returning to the location of the last found target) and avoiding search redundancy are key-factors that determine the efficiency of CPFAs. Our formal analysis suggests the following efficiency ranking from best to worst: spiral, spoke, and the stochastic ballistic algorithm.**

I. Introduction

Autonomous central-place foraging (CPF) is a fundamental task in collective robotics that involves the discovery, collection, and transportation of targets to a collection zone [1]. Central Place Foraging Algorithms (CPFAs) have recently received increased attention as resource collection on other planets, moons, and asteroids by robots is planned by space agencies to enable human exploration. Mining by autonomous vehicles and inventory collection in automated warehouses are essentially CPF tasks in that they require efficient collection and transportation of targets distributed within an area. Search and rescue, collection of bomb fragments for analysis and robotic agriculture also motivate the study of distributed search. Immune systems searching for pathogens and ant-colonies searching for food can also be understood by analyzing CPFA.

Empirical work in real robots leads us to investigate 3 simple algorithms: the Distributed Archimedes Spiral Algorithm (SpiralCPFA) [2], [3], Spoke Central Place Foraging Algorithm (SpokeCPFA) [4], and Random Ballistic Central Place Foraging Algorithm (RandCPFA) [5]. Variants of these three algorithms performed well in the NASA Swarmathon swarm robot foraging competition [4].

Formal analysis allows us to predict the performance of these CPFAs for large areas and swarms of robots for which experiments are currently impractical. We formalise two principles observed in our empirical work: the importance of site-fidelity (i.e. returning to the location of the last found target) and oversampling, i.e. repeatedly searching the same area.

To aid our analysis, we introduce the *Price of Ignorance* metric. This metric is the ratio of the performance of a given algorithm to that of omniscient algorithm. This effectively quantifies the penalty each algorithm pays for not knowing where the targets are, which is one of the main factors determining CPFA performance.

For each of the CPFAs, the proofs presented here provide relative upper-bounds on the expected performance of teams of robots. These upper-bounds suggest the ordering of algorithm performance in idealized scenarios. To test whether the ranking of upper-bounds holds, we run Autonomous Robots Go Swarming (ARGoS) [6] simulations for each of the CPFAs. In combination, the asymptotic analysis, ARGoS simulations, and experiments with real robots give us insight into how CPFAs perform in theory and in practice.

The technical details of our model along with our formal analysis is presented in Section II. We describe our empirical methods in Section III and present the results of our analysis in Section IV. Finally we discuss our findings in Section V.

Related Work. Seminal contributions in search and distributed foraging have emerged in Operations Research [7], Physics [8], Computational Geometry [9], Ecology [10], and Robotics [11]–[13]. Central place foraging has been of fundamental interest to researchers of Swarm Intelligence because of its deep connections to social insect behaviour [14]. Generating an optimal search path that maximises the probability of detecting a target in non-trivial environments within a fixed time-frame is NP-complete, minimising the mean time to detection is NP-hard [15], [16]. Therefore search and CPF use heuristics.

CPF is a key robot swarm application [11], [17], [18], and research continues to focus on improving foraging algorithms and engineering swarm robot systems for foraging [19], [20].

Ghosh and Klein [21] provide a review of planar search algorithms, which is a critical component of CPFAs. Spiral search has long been known to be an optimal search strategy for *individual agents* searching for *single targets*, both from the standpoint of computational geometry [22] and more recently as a practical algorithm for real robots [23]. The optimality of spiral search for multiple agents searching for a *line* in the plane, and a point on a line, has also been proved [9]. We have examined the generalisation to multiple robots empirically [2]. To our knowledge this

[1]Computer Science Department, [2]Center for Advanced Research Computing, [3]Biology Department, University of New Mexico, Albuquerque, USA. [4]External Faculty Member, Santa Fe Institute, Santa Fe, USA.
*Correspondence: mfricke@unm.edu

is the first formal analysis of multi-agent spiral, and spoke CPFAs.

Rybski et al. [24], and Hecker and Moses [13] demonstrate the importance of site-fidelity in robot CPF in simulation and hardware experiments.

Feinerman et al. [25], also argue in favour of deterministic algorithms using formal analysis and establish a lower bound of $\Omega\left(Rf + R^2 f/N\right)$ for the time it takes to collectively forage adversarially placed targets. In a slightly relaxed but biologically motivated setting of uniformly random cluster placement around the collection zone, we show that SpiralCPFA circumvents this bound and takes only $O\left(R^2/N + R\max\{f/N, 1\}\right)$ time in expectation for complete collection.

II. Theoretical Analysis

A. Formal Model

We analyse holonomic robots and do not consider environmental effects, robot failures, sensor or actuation error, collisions or congestion. We considered these factors in our previous work in experiments with real robots and realistic simulations [3]. We assume a simple obstacle-free circular arena of radius R with a central, stationary, collection zone. Robots begin at the collection zone and transport the targets they discover to the collection zone. The location of the collection zone is known to all robots.

Targets in natural environments tend to be locally clustered [26]. Therefore we place m clusters, containing f targets each, uniformly at random within the arena. Thus, the expected distance of the clusters from the collection zone is $2R/3$. Targets are discoverable when robots are within a detection range, r which defines a detection area a around each target. Targets are stationary and depleted on collection. The number of targets initially available for collection is fixed but unknown to the robots.

Foraging is performed by N robots that move at a constant speed, s, and have limited internal memory. The robots know R and N. Each robot can transport one target at a time. After delivering a target to the collection zone, in SpiralCPFA and SpokeCPFA robots return to the point where the search pattern was interrupted. In the RandCPFA, robots choose with some probability p_s to either employ site-fidelity, i.e. return to the location of the last target discovered, or alternatively the robot chooses a random ballistic trajectory away from the collection zone. The robots do not communicate with each other in the algorithms we study (apart from initial determination of search trajectories in the deterministic algorithms, which depends on N). Robots can remember only the last location where they found a target and the location of the collection zone. Site-fidelity is inspired by the behaviour of foraging ants [27]. Note that because the SpiralCPFA and SpokeCPFA return to the point in their search pattern where they were interrupted by collecting a target, they implicitly implement site-fidelity.

Our formal model uses expected, rather than adversarial, target distributions and therefore allows tighter lower-bounds than those previously published. Table I summarises our notation.

TABLE I: Notation. All times are for complete collection of all targets.

χ	Price of Ignorance	T_s	Total search time
T_t	Total transport time	T_{tot}	Total foraging time
N	Number of searchers	s	Robot speed
R	Radius of the circular arena	r	Detection radius of robot
f	Total number of targets	m	Total number of clusters
a	Area of a single target	p_s	Probability of site-fidelity

1) The Omniscient Central Place Foraging Algorithm: We provide a formal analysis for each of the CPFAs. To do this we need an appropriate metric to measure relative success or failure. Our metric is the time for complete collection achieved by each algorithm compared to the time taken by an idealised perfect algorithm, as a function of the arena radius, R. A perfect CPFA, for complete collection, simply has to know the location of each target *a priori*. The difference between the CPFAs performance and the perfect algorithm is an effective measure of each algorithms' price of ignorance. The only additional requirement is that multiple robots are not simultaneously scheduled to collect the same target. This is achievable by providing the robots with the locations of targets and a centralised scheduling algorithm.

Each omniscient robot deploying this ideal algorithm will take exactly $2d/s$ units of time to collect a target, which is located at distance d from the collection zone. Since the expected distance is $d = 2R/3$ to a target, the total expected time for complete collection by this ideal algorithm is $4Rf/3Ns$.

2) Price of Ignorance: For a particular problem instance and a given foraging algorithm A, the price of ignorance metric, denoted $\chi(A)$, is defined as the time taken by A to collect all targets divided by that taken by the perfect algorithm:

$$\chi(A) = \frac{T_{\text{tot}}(A)}{(4Rf/3Ns)} = \frac{3NsT_{\text{tot}}(A)}{4Rf} \quad (1)$$

Thus, $\chi(A)$ must be at least 1 for any algorithm, and the most efficient algorithms are closest to 1. We summarize our main results for the price of ignorance of three CPFAs in Table II.

B. Distributed Archimedes Spiral Algorithm

In this section, we analyse the total foraging complexity of Distributed Archimedes Spiral Algorithm a variant of the algorithm proposed in [2]. In a nutshell, the robots search along a spiral path, starting at the collection zone. If the robot finds a target, it takes it back to the collection zone and then returns to the location where it found the target and resumes its search. If the robot hits an arena boundary, it completes the circuit at the arena boundary and then stops foraging. The spiral path is unique to each

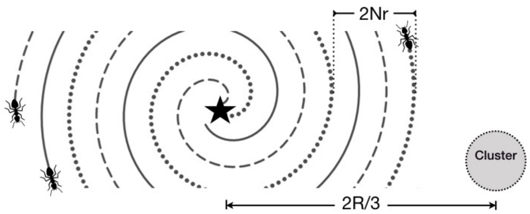

Fig. 1: Foraging path for $N = 3$ bots using SpiralCPFA, with a cluster placed at the expected distance of $2R/3$ from the collection zone.

robot foraging in the arena and collectively guarantees full arena coverage. Our main results are summarised in the theorem below:

Theorem 1. *(**SpiralCPFA Analysis**) The following holds for foraging using SpiralCPFA:*

(1) Using only a single robot, the expected price of ignorance of SpiralCPFA is at most $1 + \frac{2.25R}{rf} - \frac{1}{2f}$.

(2) Multiple robots reduce the expected price of ignorance of SpiralCPFA to at most $\sqrt{2} + \frac{1.67R}{rf} - \frac{N}{2f}$. This bound is tight up to constant factors.

Proof: [Proof Overview] The Archimedes' Spiral [28] is governed by the equation $\ell = \frac{r\theta}{\pi}$, where (ℓ, θ) is the polar co-ordinate pair of the spiral path traversed by the robot and r is the detection radius of the robot. Note that $2r$ is the separation between successive turns of the spiral, which is maintained in order to provide complete coverage of the arena.

1) Proof Sketch for Theorem 1 (1): Let $(\ell_{\max}, \theta_{\max})$ be the angular co-ordinates for the last point of the spiral that fits into the circular arena. Since the distance between $(0, 0)$ and $(\ell_{\max}, \theta_{\max})$ is R, we obtain $\theta_{\max} = \frac{R\pi}{r}$. Let L be the total length of the spiral path traversed by the robot. We can bound L as follows:

$$L \leq \frac{r}{\pi} \int_0^{\theta_{\max}} \sqrt{\theta^2 + 1}\, d\theta \leq \frac{r}{\sqrt{2}\pi} \theta_{\max}^2 < \frac{3R^2}{r} \qquad (2)$$

Note that in SpiralCPFA, every time the robot finds a target, it first collects that target back at the collection zone and then resumes its search from the point where it found the last target. Thus, to calculate the effective search time for a robot, we also consider the distance it needs to travel back to the cluster from the collection zone for each of the remaining $f - 1$ targets. This gives $E[T_s(\text{SpiralCPFA})] = \frac{L}{s} + (f-1)\frac{2R}{3s} \leq \frac{3R^2}{rs} + \frac{2R(f-1)}{3s}$. Note that the expected transport time is $\frac{2Rf}{3s}$ since it takes $\frac{2R}{3s}$ time in expectation to transport a single target back to the collection zone.

Hence, the expected total completion time for SpiralCPFA is at most $\frac{3R^2}{rs} + \frac{2R(f-1)}{3s} + \frac{2Rf}{3s}$. This gives the expected price of ignorance as stated in the theorem.

2) Proof Sketch for Theorem 1 (2): SpiralCPFA can be easily adapted for multiple robots by ensuring that each robot travels along a separate Archimedes' Spiral path in a way that no two spirals ever intersect. For example, in the steady state, Figure 1 shows the foraging trajectory for $N =$ 3 robots. Based on this, each of the robots will now traverse a path where the distance between successive spirals is $2rN$, governed by the equation $\ell_i = \frac{rN}{\pi}\theta_i + (i-1)\frac{r}{\pi}$. Here, (ℓ_i, θ_i) are the polar co-ordinates of the spiral path along which the i^{th} robot travels in the arena for $i \in \{1, 2, \ldots, N\}$.

Let $(\ell_{i,\max}, \theta_{i,\max})$ be the angular co-ordinates for the last point traversed by the i^{th} robot in the arena and L_i be the length of the Archimedes' Spiral so traversed. Similar to the analysis for a single robot, we obtain the following $\theta_{i,\max} = \left(\frac{R\pi}{r} - (i-1)\right)\frac{1}{N}$ and $L_i \leq \pi R^2/\sqrt{2}rN$.

To compute the expected total transport time, note that each robot traverses an area of $2r$ per unit length of the path it travels during the search. Thus, the total area traversed by the i^{th} robot is $2rL_i \leq \sqrt{2}\pi R^2/N$. This corresponds to $\frac{\sqrt{2}}{N}$ fraction of the arena and hence, the expected number of targets collected by any robot is at most $\frac{\sqrt{2}f}{N}$. This way, we obtain $E[T_t(\text{SpiralCPFA})] = \frac{2\sqrt{2}Rf}{3Ns}$.

Next, we bound the expected total search time. Similar to the discussion for the single robot, we can generalise the total expected search time for multiple robots as follows:

$$\begin{aligned} E[T_s(\text{SpiralCPFA})] &= \max_i \left(\frac{L_i}{s} + E[T_t(\text{SpiralCPFA})]_i - \frac{2R}{3s}\right) \\ &\leq \max_i \frac{L_i}{s} + E[T_t(\text{SpiralCPFA})] - \frac{2R}{3s} \end{aligned} \qquad (3)$$

Note that it holds that $E[T_{\text{tot}}(\text{SpiralCPFA})] = E[T_t(\text{SpiralCPFA})] + E[T_s(\text{SpiralCPFA})]$. This gives $E[T_{\text{tot}}(\text{SpiralCPFA})] \leq \max_i \frac{L_i}{s} + 2E[T_t(\text{SpiralCPFA})] - \frac{2R}{3s} < \frac{\pi R^2}{\sqrt{2}rNs} + \frac{4\sqrt{2}Rf}{3Ns} - \frac{2R}{3s}$. The bound on the price of ignorance, as stated in the lemma, then immediately follows.

Lower Bound Result: We also show that our analysis for multiple robots is tight. We omit the full proof here and only state our main result. We show that the expected price of ignorance for SpiralCPFA is at least $1 + \left(\frac{3\pi R}{8rf} - \frac{9N}{4}\right)$. To prove this, we compute a lower bound on the total length of the path traversed by any robot in the arena. Concretely, we show that $L_i \geq \frac{\pi R^2}{2Nr} - \frac{R(N-1)}{N}$. This allows us to bound the expected search time from below.

C. Spoke Central Place Foraging Algorithm

The SpokeCPFA is reminiscent of an algorithm used by the Southwestern Indian Polytechnic Institute (SIPI) in the NASA Swarmathon competition [4]. This algorithm placed the SIPI team in the top three for three years running. In the SIPI CPFA the searchers move radially away from the collection zone until they find a target or reach a predetermined limit. When a searcher returns to the collection zone it increments its angle of departure slightly for the next ballistic run. The radial search progresses around the collection zone like the hands of a clock. Successive turns are sufficiently small that a pile is likely to be encountered on the next spoke, effectively implementing site-fidelity.

Such a sweeping mechanism will cover the entire space in the arena. To ensure this we set the sweep angle in a way that the maximum distance between any two spokes is at most twice the detection radius of the robots. Moreover, the robot can move along one spoke and return along the next to further reduce the time by a factor of 2. Thus, we set the sweep angle $\Theta = 4\sin^{-1}(r/R) \geq 4r/R$ for a total of $2\pi/\Theta \leq \pi R/(2r)$ sweeps. Hence, the distance travelled by the robot to sweep the entire arena is $4R\pi/\Theta \leq \pi R^2/r$. Note that the search time also takes into consideration the time taken by the robot to travel back to the last location where it found the target, which in expectation is $2Rf/(3s)$, resulting in a total search time of $\pi R^2/r + 2Rf/(3s)$. Also, the total transport time for f targets is $2Rf/(3s)$ in expectation. This enables us to obtain the expected price of ignorance of SpokeCPFA as $1 + \frac{3\pi R}{4rf}$.

Note that the time taken for the general case where multiple targets can exist along a single spoke is never more than the bound obtained here. This is because the distance covered by the robot that takes multiple trips along a single spoke (for, say, f' targets along the spoke) until it reaches the end of the arena is upper bounded by that covered by f' single trips along different spokes.

Also note that the price of ignorance for SpokeCPFA with multiple robots is exactly the same as that obtained for a single robot. This is because each robot can be assigned a disjoint sector with angle $2\pi/N$ radians and made to forage only in that sector. Thus, N robots reduces T_{tot} by a factor of N, the same reduction as the ideal foraging algorithm (from Equation 1) and hence, the price of ignorance remains the same, independent of the total number of robots.

D. Randomized Ballistic Central-Place Foraging Algorithm

We now move on to a randomized foraging algorithm based on the SpokeCPFA discussed in Section II-C. Similar to the analytical model proposed by Levin et al. in [29], this algorithm models the random walk around the collection zone as ballistic trajectories performed in randomly chosen directions. More concretely, the robot chooses a random direction from the collection zone and moves in a straight line path along that direction until either a target is found (either in direct line of sight or within the detection radius) or a pre-specified distance is reached (which in this case is the arena radius R). We refer to this as the RandCPFA (or, RandCPFA in short) in this paper. We state the main results for this algorithm in the theorem below and provide a proof sketch thereafter.

Theorem 2. ***(RandCPFA Analysis)*** *The following holds for stochastic foraging using RandCPFA:*

(1) Using only a single robot, the expected price of ignorance of RandCPFA is at most $\frac{\pi R}{r} - \frac{1}{2}$ without any site fidelity exhibited by the robot or any consideration of the effect of food depletion.

(2) Deploying site fidelity is expected to strictly decrease the price of ignorance to at most $\frac{m}{f}\left(\frac{\pi R}{r}\right) + \frac{1}{2}\left(1 - \frac{m}{f}\right)$, without any consideration of the effect of food depletion.

(3) For arenas having multiple target clusters, the depletion of these piles over time (due to foraging) causes the price of ignorance to increase to at most $\frac{(2\pi \ln 2)NR}{r} + \frac{1}{2}$, when no site fidelity is deployed by the robot.

(4) With both site fidelity and depletion effects into consideration, the expected price of ignorance of RandCPFA is at most $\frac{\pi R}{r} - \frac{N}{2f}$. This holds true even when the robot employ site fidelity and recruitment strategies with multiple robot, assuming all the robot move independently and identically of each other and collision effects are not considered.

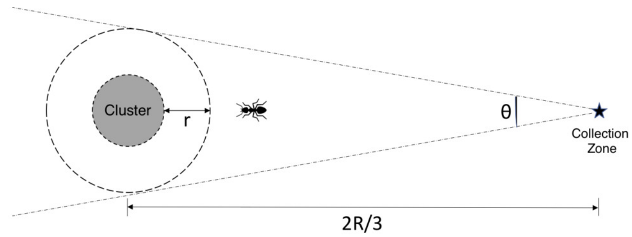

Fig. 2: Relationship between the angular region in which the cluster is detected in a single ballistic walk and the area of the cluster/ detection radius of the robot.

Proof: [Proof Overview] We begin by computing the likelihood p of finding at least one target cluster in a random chosen direction from the collection zone (see Figure 2). It can be shown that $p = 1 - \left(1 - \frac{\theta}{2\pi}\right)^m$, where $\theta = 2\sin^{-1}\left(\frac{3}{2R}\left(r + \sqrt{\frac{fa}{m\pi}}\right)\right)$. This is because the target cluster is detected as soon as its outer boundary falls within the detection radius of the robot [5]. The inverse sine expression can be simplified to obtain the following two bounds:

$$\theta \geq 3r/R \tag{4}$$

$$3r/(2\pi R) \leq p \leq 3mr/(\pi R) \tag{5}$$

1) Proof Sketch for Theorem 2 (1): If the target is not found, the robot travels for $2R/s$ units of time before reaching the collection zone again, else, it travels for expected $2R/3s$ units of time before finding a target cluster. Assuming only a single robot in the system, recall that with probability p, the robot finds at least one target cluster. This causes the expected search time for the robot to be the following:

$$T_s = \frac{1}{s}\sum_{i=0}^{\infty}(1-p)^i p\left(2Ri + \frac{2R}{3}\right) = \frac{2R(3-2p)}{3ps} \tag{6}$$

Without site fidelity, then to find all f targets, the robot needs to make f such identical and independent trips, resulting in a total search time of fT_s.

Now, observe that once a target cluster is found, the robot moves in a straight line path to the collection zone, each costing $2R/3s$ units of time. Thus, the total transport time taken by the robot to collect all targets in the arena is $T_t = \frac{2fR}{3s}$. Adding the total search time and the transportation time, we obtain the total foraging time for RandCPFA to be $T_{\text{total}} = fT_s + T_t = \frac{2Rf(3-2p)}{3ps} + \frac{2fR}{3s} = \frac{2Rf(3-p)}{3ps}$. This allows us to bound the price of ignorance using Equation (1) as follows:

$$\chi\,(\textsc{RandCPFA}) = \left(\frac{2Rf(3-p)}{3ps}\right)\left(\frac{3s}{4Rf}\right) = \frac{3-p}{2p} \qquad (7)$$

Using the lower bound on p from Equation (5), we obtain the the result as stated.

We obtain a tighter bound using a coupon-collector argument in Section II-E. However, it becomes non-trivial to extend this argument for the case of site-fidelity and target depletion and hence, we continue the remaining proofs in this theorem using the same technique as discussed above.

2) Proof Sketch for Theorem 2 (2): We say that the robot exhibits site fidelity $p_s \in [0,1]$ if each time it finds a target at a particular site, it returns to that site for its next trip with probability p_s. Otherwise, it chooses a random direction to move to. We then denote $\chi(p_s)$ to be the price of ignorance given site fidelity p_s.

Ignoring the effects of target depletion, the expected search time to find a target cluster for the first time is the same as before. However, for each of the remaining $\frac{f}{m}-1$ targets in that pile, the robot takes in expectation $p_s(2R/3s)+(1-p_s)T_s$ units of time. This is because the robot returns to the same location with probability p_s and performs a Ballistic search otherwise. Hence, the expected search time to completely consume a target cluster with site fidelity is $T_s+(f/m-1)\,(p_s(2R/3s)+(1-p_s)T_s)$, which can be rearranged to $(1+(f/m-1)(1-p_s))T_s+p_s(f/m-1)(2R/3s)$. The transport time for these targets remains the same. Hence, the total expected time taken by the robot to collect all the targets in the arena is $(mp_s+f(1-p_s))\,T_s + p_s\,(f-m)\,\frac{2R}{3s}$. Using Equations (5) and (6), we show that the expected price of ignorance is at most $\frac{m}{f}\left(\frac{\pi R}{r}\right)+\frac{1}{2}\left(1-\frac{m}{f}\right)$, as stated in the theorem.

3) Proof Sketch for Theorem 2 (3): For mathematical simplicity, we perform our analysis assuming only one target per pile ($f=m$) and no site fidelity (so that each robot collects m/N piles in expectation). We begin with computing how the probability of finding target clusters degrades as they start depleting. Let p_n be the probability that in a single trip of the robot, at least one target cluster is encountered, assuming the arena currently has n unexplored target clusters. Then, $p_n \approx \frac{n\theta}{2\pi}$ (when θ is sufficiently small compared to $2\pi/m$). Given this value of p_n, the (expected) search time is at most $\frac{2R(1+2p_n)}{3s(1-p_n)}$ to find a single target cluster. Thus, to find the time to search for all the m/N target clusters, we use an integral approximation and assume $N<m$ to obtain the following (here T_s^{depl} denotes the total (expected) search time with depletion effects and $q_n = 1-p_n$):

$$T_s^{depl} \le \sum_{k=1}^{m/N}\left(\frac{2R(1+2p_{m-kN})}{3s(1-p_{m-kN})}\right) \le \frac{8\pi m R^2 \ln 2}{3rs} \qquad (8)$$

Using Equation (4), we obtain the following:

$$\chi \le \frac{3Ns}{4Rm}\left(\frac{8\pi m R^2\ln 2}{3rs}+\frac{2Rm}{3Ns}\right) \le \frac{(2\pi\ln 2)NR}{r}+\frac{1}{2} \qquad (9)$$

4) Proof Sketch for Theorem 2 (4): Let p_f be the probability of discovering the target cluster in one Ballistic walk by the robot when the number of targets in the target cluster is f. We know that $p_f = \theta_f/2\pi$, where θ_f has the same form as in Equation (4). Let $T_{s,f}^{\mathrm{NoSF}}$ be the expected time (without site fidelity) for as single robot to find the target cluster when it has f targets. Note that a single robot collects $1/N$ fraction of the total targets in expectation. Then, we obtain the following:

$$T_{s,f}^{\mathrm{NoSF}} = \sum_{i=1}^{f/N} T_{s,i}^{\mathrm{NoSF}} = \sum_{i=1}^{f/N}\frac{2R(3-2p_i)}{3sp_i} \le \frac{4Rf}{3Ns}\left(\frac{\pi R}{r}-1\right) \qquad (10)$$

where the last inequality follows from the bound in Equation (5). Adding the effect of site fidelity (where p_s is the probability that the robot will return to the same location that it last located a target), we denote by $T_{s,f}(p_s)$ the expected search time in this case. Then, we obtain the following:

$$\begin{aligned} T_{s,f}(p_s) &= T_{s,f}^{\mathrm{NoSF}} + \sum_{i=1}^{(f/N)-1}\left(\frac{2Rp_s}{3s}+T_{s,i}^{\mathrm{NoSF}}(1-p_s)\right) \\ &\le \frac{4Rf}{3Ns}\left(\frac{\pi R}{r}-1\right)+\frac{2Rf}{3Ns}-\frac{2R}{3s} \end{aligned} \qquad (11)$$

Note that the bound above holds for $p_s = 1$ (perfect site fidelity). Adding the expected transport time, this gives the upper bound on the total expected foraging time of RandCPFA as $\frac{4Rf}{3Ns}\left(\frac{\pi R}{r}\right)-\frac{2R}{3s}$. This causes the price of ignorance to be at most $\frac{\pi R}{r}-\frac{N}{2f}$.

E. Tighter Analysis for RandCPFA

We now discuss a method to obtain a tighter bound for the expected foraging time for RandCPFA, ignoring the effects of target depletion and site fidelity. Observe that in the proof for Theorem 2, we assumed that the expected time it takes to collect f targets is at most f times the expected time to collect a single target. This has an oversimplification effect that the robot now makes multiple passes over the entire arena until complete collection is achieved.

However, observe that as soon as the entire arena is foraged once, all targets would have been collected and no further passes are required. This allows us to use the expected time to cover the entire arena as a tighter upper bound on the search time for RandCPFA, compared to that obtained in Theorem 2 (1). To obtain this bound, we bound the number of distinct Ballistic trajectories by $2\pi/\theta$, where θ is the same as that in Equation (4). Since each trajectory can be traversed multiple times independently and identically, the expected number of times each trajectory is traversed at least once is at most $\frac{4\pi}{\theta}\log\left(\frac{2\pi}{\theta}\right)$, which is obtained similar to the analysis of coupon collector problem [30]. The maximum time it takes to complete each trajectory is $2R/s$, and hence, a tighter upper bound on the total foraging time is now at most $\left(\frac{2R}{s}\right)\left(\frac{4\pi}{\theta}\right)\log\left(\frac{2\pi}{\theta}\right)+\frac{2Rf}{3s} \le \left(\frac{8\pi R^2}{3sr}\right)\log\left(\frac{2\pi R}{3r}\right)+\frac{2Rf}{3s}$. This allows us to bound the price of ignorance by $\left(\frac{2\pi R}{rf}\right)\log\left(\frac{2\pi R}{3r}\right)+\frac{1}{2}$, which is significantly lower that in Theorem 2 (1) (see Figure 4a for an empirical comparison).

III. Empirical Methods

To validate the formal analysis, we ran experiments using the ARGoS simulator for a variant of each algorithm implemented for a square arena with $N = 10$ robots. The "foot-bot" robot, programmed to move like the iAnt robot from [13], was used for all experiments. The experiments used one and ten robots and measured the time to complete collection in an environment with $f = 256$ targets arranged in $m = 4$ clusters, of 64 items each, placed uniformly at random in the arena. Robots were able to detect and pick up a target if they came within $r = 0.2$ meters of it. For the RandCPFA experiments the robot always use site fidelity ($p_s = 1$). Results from these experiments are shown in Section IV Figure 3b.

IV. Results

Figure 3a shows the relative price of ignorance for each algorithm given in Section II. For $R < 29.28$ SpokeCPFA has a lower price of ignorance in comparison to the SpiralCPFA. Our empirical results, shown in Figure 3b, show a similar crossover point, but at a much larger arena radius ($R \approx 100$) due to collision effects caused by congestion. As predicted, our experiments show that SpiralCPFA has the lowest cost of ignorance for large arenas. In addition to congestion in small arenas, SpokeCPFA and SpiralCPFA also performs better due to the higher density of targets. Furthermore, in simulation, the variance in price of ignorance is very low for SpiralCPFA. For SpokeCPFA and RandCPFA the price of ignorance is highly variable, demonstrating that performance is sensitive to the placement of clusters.

Figure 4a shows that performance in simulation is considerably better than the upper bound given in section II-D. Figure 4b shows that our empirical results match our formal analysis well for the SpokeCPFA and even better for the SpiralCPFA, especially for large R. Because robots constantly return to the collection zone without retrieving a target, the SpokeCPFA diverges from the theoretical bound in simulation. In contrast, the SpiralCPFA experiments suffer less from congestion in the large arena because robots *only* return to the collection zone when they have retrieved a resource. We believe the deviation from the theoretical model for SpiralCPFA is almost entirely due to congestion at small R.

V. Discussion

Our theoretical analysis and experiments in simulation help to explain the effectiveness of spiral algorithms observed in previous simulations and experiments with physical robots [2], [3]. In our formal analysis SpiralCPFA is expected to outperform both SpokeCPFA and RandCPFA. We attribute this to the non-uniformity of SpiralCPFA, where the robot continuously eliminates area to search for targets with every step and no two robots have overlapping search areas. This is not the case for SpokeCPFA which repeatedly visits points near the centre of the arena (Figure 5b). Overlap in search is even worse in RandCPFA since each successive ballistic walk is made independently of previously foraged areas. We plot an approximation of the fraction of area eliminated for search over time for the three algorithms in Figure 5a. We observe that RandCPFA indeed chooses to ignore all previously obtained information and performs each ballistic walk as if it were foraging in a completely unknown arena. On the other hand, SpiralCPFA shows a monotonic increase in the information gained about what parts of the arena do not contain any more targets and chooses never to traverse those areas again. SpokeCPFA shows a more shallow increase due to overlap near the centre.

TABLE II: Summary of various analytical results in this paper. See Table I for notation. Ranked according to the least upper-bound on expected performance (lower is better).

Algorithm	**Rank**	**Collection Time**	**Price of Ignorance**
SpiralCPFA	1	$\frac{\pi R^2}{\sqrt{2}Nrs}-\frac{2R}{3s}+\frac{4\sqrt{2}Rf}{3Ns}$	$\frac{3\pi R}{4\sqrt{2}rf}-\frac{N}{2f}+\sqrt{2}$
SpokeCPFA	2	$\frac{\pi R^2}{Nrs}+\frac{4Rf}{3Ns}$	$\frac{3\pi R}{4rf}+1$
RandCPFA	3	$\frac{4\pi R^2 f}{3Nrs}-\frac{2R}{3s}$	$\frac{\pi R}{r}-\frac{N}{2f}$

Note that our empirical results for the deterministic algorithms are consistent with those predicted by our formal analysis, which is not true for the case of RandCPFA. We attribute this to the deterministic nature of the SpiralCPFA and SpokeCPFA that facilitates accurate formal analysis. On the other hand, the analysis for RandCPFA is challenging even with our simplifying assumptions. We make a first attempt to provide a tighter analysis in Section II-E, but that analysis gives a tight bound on searching every location in the arena, and all clusters can be found in less time. We leave a tighter analysis to future work.

As clusters become smaller and disappear (due to collection of targets), it becomes more challenging for the robots to find the remaining targets. To the best of our knowledge, we make the first attempt to model the effect of depletion in collaborative foraging and prove that when food depletion is taken into account, one can expect a multiplicative blowup up to $\Omega(N)$ in the price of ignorance. Thus, a larger number of bots on one hand collect food items faster (by working in parallel), but they also deplete the arena at a rate which significantly slows down foraging as time progresses. The effect of target depletion is further exacerbated in our experiments because robots may bore a "tunnel" through the clusters, resulting in two smaller clusters that must be rediscovered before all resources can be collected, an effect not captured by the model.

Our analysis shows that site fidelity strongly counteracts the effect of target depletion allows RandCPFA to remain efficient. This has been observed empirically [31], [32].

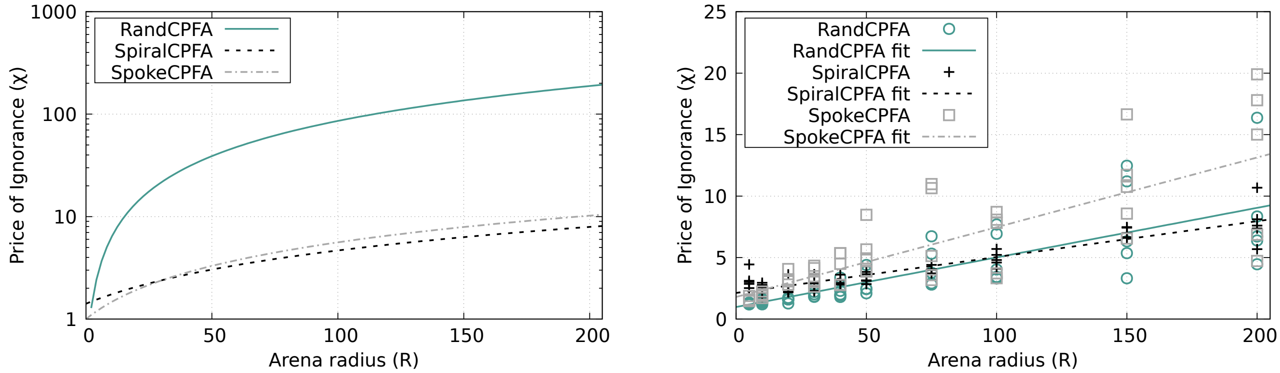

(a) (b)

Fig. 3: (a) Comparison of CPFAs from formal analysis. Note that the y-axis is on a log scale. (b) Price of Ignorance for simulated robots ($N = 10$). Collisions cause the SpiralCPFA to underperform for smaller arena sizes.

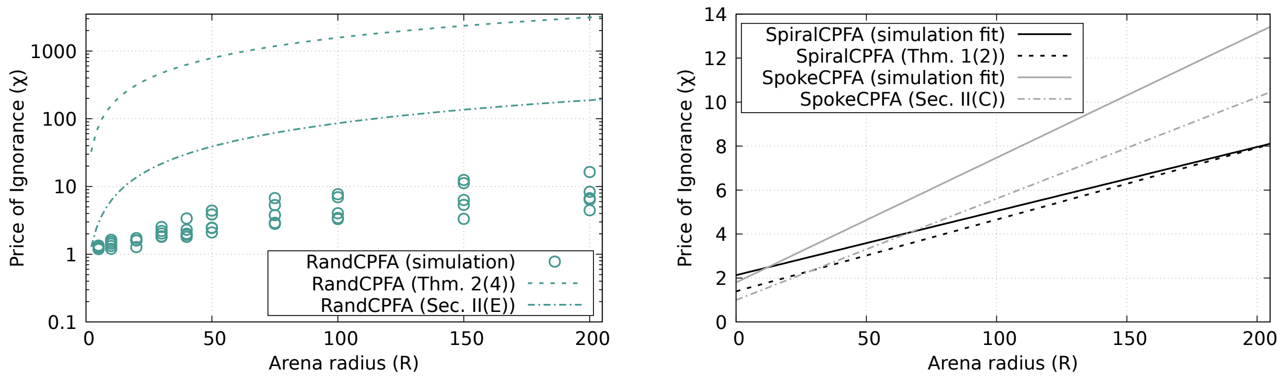

(a) (b)

Fig. 4: Comparison of observed price of ignorance to the upper bounds derived in Section II. (a) Theory vs. simulation for RandCPFA. (b) Theory vs. simulation for SpiralCPFA and SpokeCPFA, only the linear fit is shown for clarity.

Our results provides one of the first formal arguments for phenomenon.

Steels [17], and Balch and Arkin [33] showed that communication of target locations by swarm members improves performance in distributed foraging tasks. However their results were based on simulation only. As demonstrated by Rybski et al. [24] implementation of swarm communication in real systems may not show the expected improvement. Communication by real error-prone robots interpreting a noisy environment can harm foraging performance by misdirecting foraging resources, a phenomena termed *misinformation cost* by Pitonakova et al. [34].

To illustrate how communication among the robots can help improve foraging, consider the following scenario in which robots carrying out the RandCPFA have a special marker which they can place in the arena at any time. This marker is only visible to the nearby robots and once detected, indicates a desired direction to them. Assuming only a single pile in the arena, when some robot finds a target in this pile, it places the marker near the collection zone in the correct direction and all other bots then forage the remaining targets in optimal time. To further optimise the algorithm, all bots can search the first target in parallel and the one that finds it first places the marker. Note that this effect is different from site fidelity since only one robot returns to the last visited location in the latter. The expected distance covered in this marker-based algorithm is $O(R^2/N + Rf/N)$, which closes the gap between the asymptotic performance of RandCPFA and SpiralCPFA. Steels [17] implements a similar approach using "crumbs", as do Hecker and Moses [13] using virtual waypoints.

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. 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. From a biological perspective, we provide analytical insight into naturally observed phenomena of site fidelity [31]. 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.

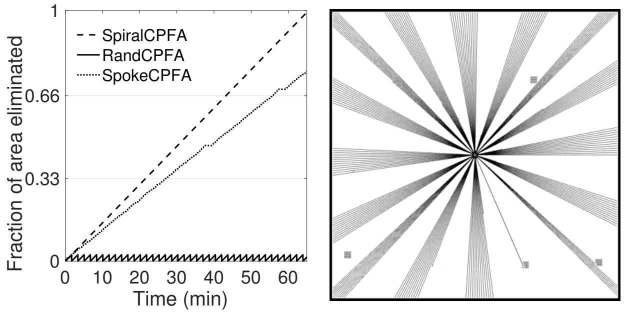

(a) (b)

Fig. 5: (a) Fraction of area explored by the robot over time. SpiralCPFA never re-searches the same area, which reduces the price of ignorance. (b) Visualisation of SpokeCPFA in ARGoS. Black lines are drawn where robots have searched. A gradient is visible showing that areas near the nest have been re-searched multiple times.

VI. Acknowledgements

We thank the Southwest Indian Polytechnic Institute and Durham Technical Community College teams along with the rest of the NASA Swarmathon participants for their central-place foraging algorithm submissions. We also thank Jared Saia for valuable input and discussion on theoretical aspects of this work.

VII. Funding

This work was supported by a James S. McDonnell Foundation Complex Systems Scholar Award, funding from the UNM Vice President for Research, and NASA MUREP #NNX15AM14A funding for the Swarmathon.

References

[1] A. F. T. Winfield, "Foraging Robots," in *Encyclopedia of complexity and systems science*, R. A. Meyers, Ed. New York: Springer, 2009, pp. 3682–3700.

[2] G. Fricke, J. Hecker, A. Griego, L. Tran, and M. Moses, "A Distributed Deterministic Spiral Search Algorithm for Swarms," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, 2016, pp. 4430–4436.

[3] L. Qi, A. D. Griego, G. M. Fricke, and M. E. Moses, "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, 2019.

[4] S. M. Ackerman, G. M. Fricke, J. P. Hecker, K. M. Hamed, S. R. Fowler, A. D. Griego*, J. C. Jones*, J. J. Nichol*, K. W. Leucht, and M. E. Moses, "The Swarmathon: An Autonomous Swarm Robotics Competition," in *International Conference on Robotics and Automation (ICRA), 2018 IEEE/RSJ International Conference on*. IEEE.

[5] D. Levin, S. Forrest, S. Banerjee, C. Clay, J. Cannon, M. Moses, and F. Koster, "A spatial model of the efficiency of {T} cell search in the influenza-infected lung," *Journal of theoretical biology*, vol. 398, pp. 52–63, 2016.

[6] C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, and F. Ducatelle, "ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems," *Swarm intelligence*, vol. 6, no. 4, pp. 271–295, 2012.

[7] L. D. Stone, *Theory of optimal search*. Academic Press New York, 1975.

[8] F. Bartumeus and J. Catalan, "Optimal search behavior and classic foraging theory," *Journal of Physics A: Mathematical and Theoretical*, vol. 42, no. 43, p. 434002, 2009.

[9] R. Baeza-Yates and R. Schott, "Parallel searching in the plane," *Computational Geometry*, vol. 5, no. 3, pp. 143–154, 1995.

[10] G. H. Pyke, "Optimal foraging theory: a critical review," *Annual review of ecology and systematics*, pp. 523–575, 1984.

[11] A. Campo and M. Dorigo, "Efficient multi-foraging in swarm robotics," in *Advances in Artificial Life*. Springer, 2007, pp. 696–705.

[12] W. Liu and A. Winfield, "Modelling and optimisation of adaptive foraging in swarm robotic systems," *The International Journal of Robotics Research*, 2010.

[13] J. P. Hecker and M. E. Moses, "Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms," *Swarm Intelligence*, vol. 9, no. 1, pp. 43–70, 2015.

[14] S. Garnier, J. Gautrais, and G. Theraulaz, "The biological principles of swarm intelligence," *Swarm Intelligence*, vol. 1, no. 1, pp. 3–31, 2007.

[15] K. E. Trummel and J. R. Weisinger, "Technical Note—The Complexity of the Optimal Searcher Path Problem," *Operations Research*, vol. 34, no. 2, pp. 324–327, 1986.

[16] A. Sarmiento, R. Murrieta-Cid, and S. Hutchinson, "A multi-robot strategy for rapidly searching a polygonal environment," in *Ibero-American Conference on Artificial Intelligence*. Springer, 2004, pp. 484–493.

[17] L. Steels, "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, 1990, pp. 8–14.

[18] A. Drogoul and J. Ferber, "From Tom Thumb to the Dockers: Some experiments with foraging robots," *From Animals to Animats*, vol. 2, p. 451, 1993.

[19] A. F. T. Winfield, "Towards an engineering science of robot foraging," *Distributed autonomous robotic systems*, vol. 8, pp. 185–192, 2009.

[20] M. Dorigo, M. Birattari, and M. Brambilla, "Swarm robotics," *Scholarpedia*, vol. 9, no. 1, p. 1463, 2014.

[21] S. K. Ghosh and R. Klein, "Online algorithms for searching and exploration in the plane," *Computer Science Review*, vol. 4, no. 4, pp. 189–201, 2010.

[22] R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins, "Searching in the plane," *Information and Computation*, vol. 106, no. 2, pp. 234–252, 1993.

[23] M. Meghjani, S. Manjanna, and G. Dudek, "Multi-target rendezvous search," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 2596–2603.

[24] P. E. Rybski, A. Larson, H. Veeraraghavan, M. Anderson, and M. Gini, "Performance evaluation of a multi-robot search & retrieval system: Experiences with MinDART," *Journal of Intelligent and Robotic Systems*, vol. 52, no. 3-4, pp. 363–387, 2008.

[25] O. Feinerman, A. Korman, Z. Lotker, and J.-S. Sereni, "Collaborative search on the plane without communication," in *Proceedings of the 2012 ACM symposium on Principles of distributed computing*. ACM, 2012, pp. 77–86.

[26] M. E. Ritchie, *Scale, heterogeneity, and the structure and diversity of ecological communities*. Princeton University Press, 2009.

[27] T. P. Flanagan, K. Letendre, W. R. Burnside, G. M. Fricke, and M. E. Moses, "Quantifying the Effect of Colony Size and Food Distribution on Harvester Ant Foraging," *PloS one*, vol. 7, no. 7, p. e39427, 2012.

[28] H. Holland, "The archimedes spiral," *Nature*, vol. 179, no. 4556, p. 432, 1957.

[29] D. Levin, "The environment constrains successful search strategies in natural distributed systems," 2016.

[30] R. Motwani and P. Raghavan, *Randomized algorithms*. Cambridge university press, 1995.

[31] D. P. Mersch, A. Crespi, and L. Keller, "Tracking individuals shows spatial fidelity is a key regulator of ant social organization," *Science*, vol. 340, no. 6136, pp. 1090–1093, 2013.

[32] J. P. Hecker and M. E. Moses, "Beyond Pheromones: Evolving Adaptable, Flexible, and Scalable Ant-Inspired Robot Swarms," 2014.

[33] T. Balch and R. C. Arkin, "Communication in reactive multiagent robotic systems," *Autonomous Robots*, vol. 1, no. 1, pp. 27–52, 1994.

[34] L. Pitonakova, R. Crowder, and S. Bullock, "The Information-Cost-Reward framework for understanding robot swarm foraging," *Swarm Intelligence*, vol. 12, no. 1, pp. 71–96, 2018.