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]

Nov 7th, 2019

[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

Central Place Foraging

- Central Place Foraging (CPF) is the discovery, collection, and transportation of targets to a location. A central-place foraging algorithm (CPFA) solves this task.

Central Place Foraging

- Central Place Foraging (CPF) is the discovery, collection, and transportation of targets to a location. A central-place foraging algorithm (CPFA) solves this task.
- This is a canonical task for teams of robots since it combines so many common robot subtasks (Winfield [1])

Central Place Foraging

- Central Place Foraging (CPF) is the discovery, collection, and transportation of targets to a location. A central-place foraging algorithm (CPFA) solves this task.
- This is a canonical task for teams of robots since it combines so many common robot subtasks (Winfield [1])
- Applications include automated agriculture, environmental sampling, planetary exploration including in-situ resource utilisation, robotic mining, etc, etc.

Central Place Foraging

- Central Place Foraging (CPF) is the discovery, collection, and transportation of targets to a location. A central-place foraging algorithm (CPFA) solves this task.
- This is a canonical task for teams of robots since it combines so many common robot subtasks (Winfield [1])
- Applications include automated agriculture, environmental sampling, planetary exploration including in-situ resource utilisation, robotic mining, etc, etc.
- This is also a fundamental task in biological systems such as the immune system and ant colonies.

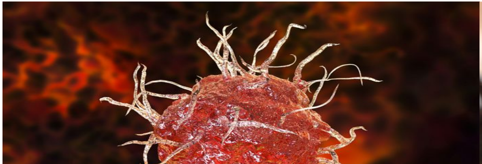


Central Place Foraging

- *Pogonomyrmex* sp. (Desert Harvester ants). Collect seeds.
- Swarmie robots generalise the central place foraging task. Allow us to design and test algorithms.
- We have built 100 Swarmies.

NASA Living off the Land - Extended Missions

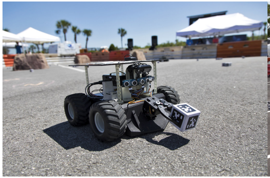

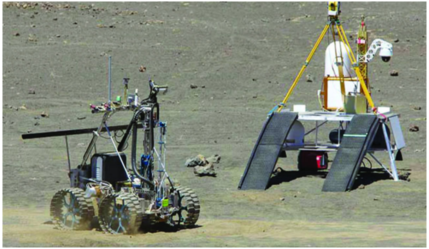

- NASA has developed fuel processing and mining robots for use on the Moon and Mars.
- Through the NASA Swarmathon Competition [2] we solicited more than 100 different CPFAs over 4 years from 40 universities.

NASA Swarmathon

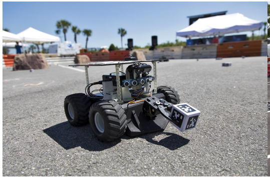


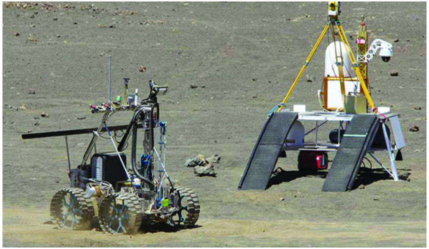

- NASA has developed fuel processing and mining robots for use on the Moon and Mars.
- Through the NASA Swarmathon Competition [2] we solicited more than 100 different CPFAs over 4 years from 40 universities.
- We can roughly divide the resulting CPFAs into four categories:
  - ⋆ Recruitment[1]
  - ⋆ Random Ballistic
  - ⋆ Rotating Spoke
  - ⋆ Spiral Pattern

NASA Swarmathon

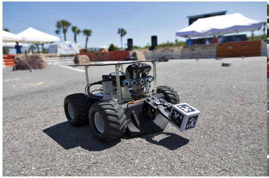

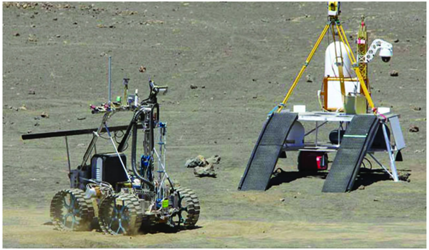

- NASA has developed fuel processing and mining robots for use on the Moon and Mars.
- Through the NASA Swarmathon Competition [2] we solicited more than 100 different CPFAs over 4 years from 40 universities.
- We can roughly divide the resulting CPFAs into four categories:
  - ⋆ Recruitment[1]
  - ⋆ Random Ballistic
  - ⋆ Rotating Spoke
  - ⋆ Spiral Pattern

[1] Surprisingly, information sharing algorithms have not done well so far.

Our Goal: We want proof!

Our observations in real robots and simulations give a general efficiency ordering[1]:

1. Spiral Algorithms
2. Rotating Spoke
3. Stochastic Walks

There are many confounding factors in experiments with real robots. Maybe people who coded spiral search happened to be better at localisation and pick and place on average, etc.

We want a formal explanation of the ordering!
No one has ever performed a complexity analysis of ANY CPFA

[1] Fricke IROS 2016 [3], Ackerman ICRA 2018, [2], and Qi ICRA 2019 [4]

Randomised Ballistic Algorithms

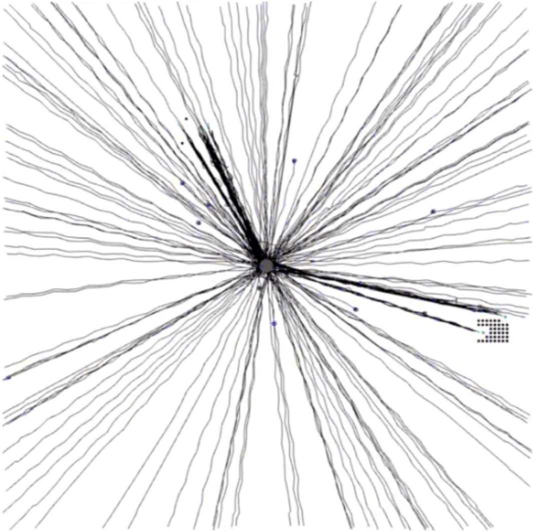

- Random algorithms are diverse but in general they do not guarantee complete coverage and allow for areas to be revisited.
- We have repeatedly found that the best stochastic CPFAs use ballistic motion.
- Robots move away from the central collection zone in a random direction.
- Motion is ballistic and is based on Levin [5].
- Alternatives include Brownian motion and Lévy walks. We find that CPFA Lévy walks converge to Ballistic motion. [6]

Spoke Algorithms

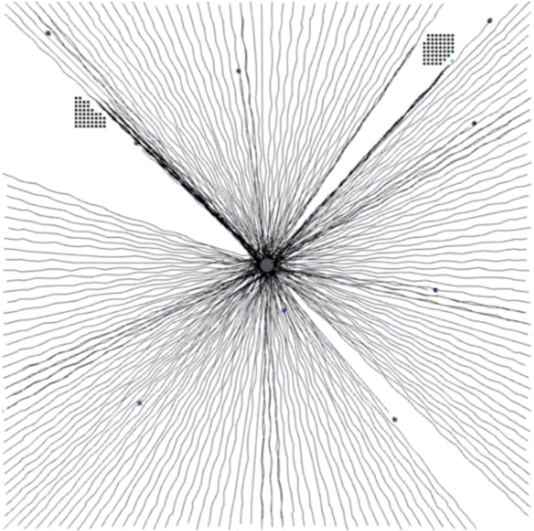

- Southwest Indian Polytechnic Institute: Third place year 1, First place year 2, Second place year 3.
- Robots divide the search space into equal regions.
- Minimises collisions between robots.
- Ensures complete coverage.

Spiral Algorithms

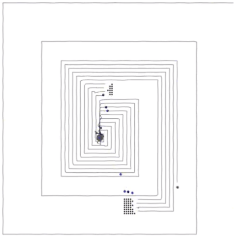

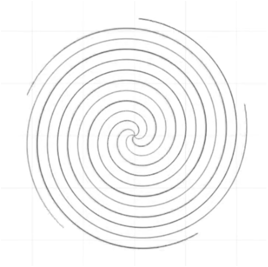

- Durham Technical College: Year 3 winner, Cabrillo College: Year 4 winner.
- Spirals can have different geometries, but as long as they are space filling the same asymptotic performance.
- Complete coverage no revisiting sites.
- Single searcher case is currently used by the US Coast Guard for search and rescue.
- See Fricke, IROS, 2016 [3] for Square Spiral.

Spiral Algorithms

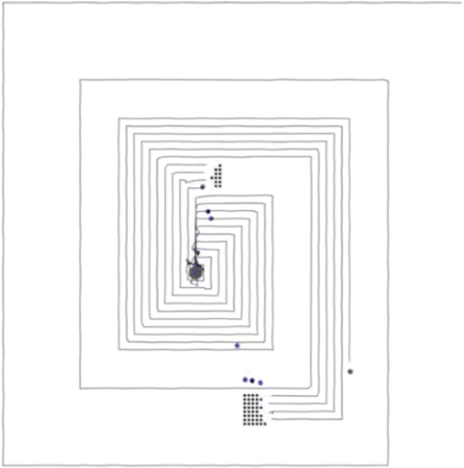

- Durham Technical College: Year 3 winner, Cabrillo College: Year 4 winner.
- Spirals can have different geometries, but as long as they are space filling the same asymptotic performance.
- Complete coverage no revisiting sites.
- Single searcher case is currently used by the US Coast Guard for search and rescue.
- See Fricke, IROS, 2016 [3] for Square Spiral.

The Omniscient CPFA

For comparison we define an ideal CPFA that knows where all the targets are.

- We examine the time to complete collection. Therefore the omniscient algorithm's total cost is the sum of the time taken to retrieve all the targets, i.e. the total transport time, T_{t}.

The Omniscient CPFA

For comparison we define an ideal CPFA that knows where all the targets are.

- We examine the time to complete collection. Therefore the omniscient algorithm's total cost is the sum of the time taken to retrieve all the targets, i.e. the total transport time, T_{t}.
- R is the radius of the foraging area.
- f is the number of targets, N is the number of robots.
- s is the robots' average speed.
- So the Omniscient CPFA has zero search cost, denoted T_{s}.
- Two trips per target, f, and because the search area is circular we have $\frac{2}{3}R$; $\times 2$ for trips is $\frac{4R}{3}$, divided by the number of robots, N, and speed, s.
- Therefore, $T_{\mathrm{t}} = \frac{4Rf}{3Ns}$

The Price of Ignorance

Given the Omniscient CPFA we can define the Price of Ignorance metric, $\chi(A)$ to be the ratio of the time algorithm A takes to collect all targets to the Omniscient CPFA's time.

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

Thus, the closer $\chi(A)$ is to 1 the more efficient A is.

Unlike the Omniscient CPFA real algorithms must have their robots search for targets. We therefore defined r to be the range at which a robot can detect a target.

Theorems: Random Spoke CPFAs

Theorem 1 (Random CPFA's Price of Ignorance)

1. $T_{tot} \in O\left(\frac{4\pi R^2 f}{3Nrs} - \frac{2R}{3s}\right)$
2. $\chi(A) \in O\left(\frac{2\pi R}{rf} \log\left(\frac{2\pi R}{3r}\right) + \frac{1}{2}\right)$ *(Coupon Collector).*

Theorem 1 (Random CPFA's Price of Ignorance)

1. $T_{tot} \in O\left(\frac{4\pi R^2 f}{3Nrs} - \frac{2R}{3s}\right)$
2. $\chi(A) \in O\left(\frac{2\pi R}{rf} \log\left(\frac{2\pi R}{3r}\right) + \frac{1}{2}\right)$ *(Coupon Collector).*
3. *Target depletion increases the price of ignorance:* $\chi(A) \in O\left(\frac{(2\pi \ln 2)NR}{r} + \frac{1}{2}\right)$.
4. *But when site fidelity (self-recruitment) is introduced the cost is reduced again:* $\chi(A) \in O\left(\frac{\pi R}{r} - \frac{N}{2f}\right)$.

Site-fidelity gives you many of the advantages of pheromone recruitment without some of the downsides. Desert harvester ants seem to use it instead of pheromones.

Theorems: Rotating Spoke CPFAs

The rotating spoke algorithm sweeps segments of a disk resulting in a price of ignorance that does not depend on N, but only on the radius of the foraging area, the number of clusters and the robot sensor range.

Theorem 2 (Spoke CPFA's Price of Ignorance)

1. $T_{tot} \in O\left(\frac{\pi R^2}{Nrs} + \frac{4Rf}{3Ns}\right)$
2. *In expectation* $\chi(A) \in O\left(\frac{3\pi R}{4rf} + 1\right)$

Theorems: Spiral CPFAs

Expected transport time: $T_t \leq \frac{2\sqrt{2}Rf}{3Ns}$

Theorem 3 (Spiral CPFAs Price of Ignorance)

1. $T_{tot} \in \Theta\left(\frac{\pi R^2}{\sqrt{2}Nrs} - \frac{2R}{3s} + \frac{4\sqrt{2}Rf}{3Ns}\right)$
2. $\chi(A) \in \Theta\left(\sqrt{2} - \frac{N}{2f} + \frac{3\pi R}{4\sqrt{2}rf}\right).$

Testing the Analysis: Random CPFA

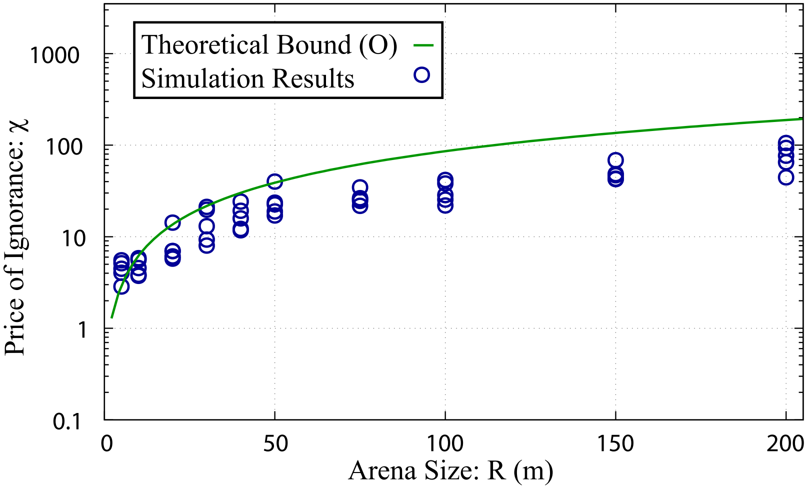

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64. Collisions.

Testing the Analysis: Spoke CPFA

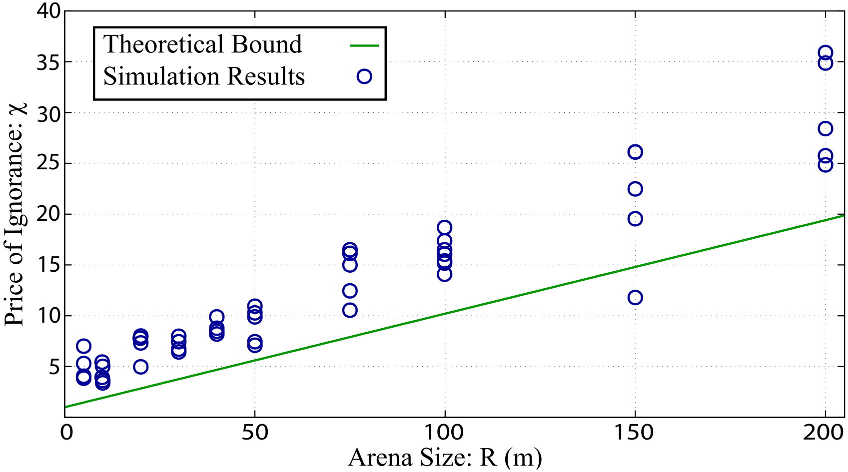

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64. Collisions.

Testing the Analysis: Spiral CPFA

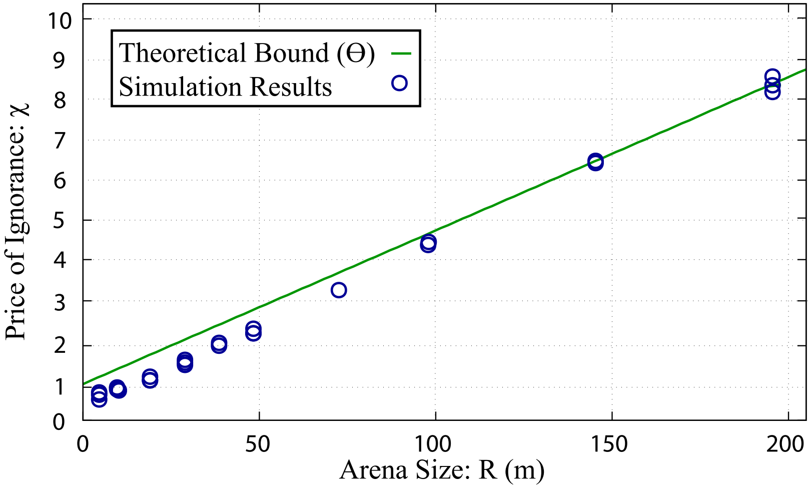

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64. Collisions.

Testing the Analysis: Spiral CPFA

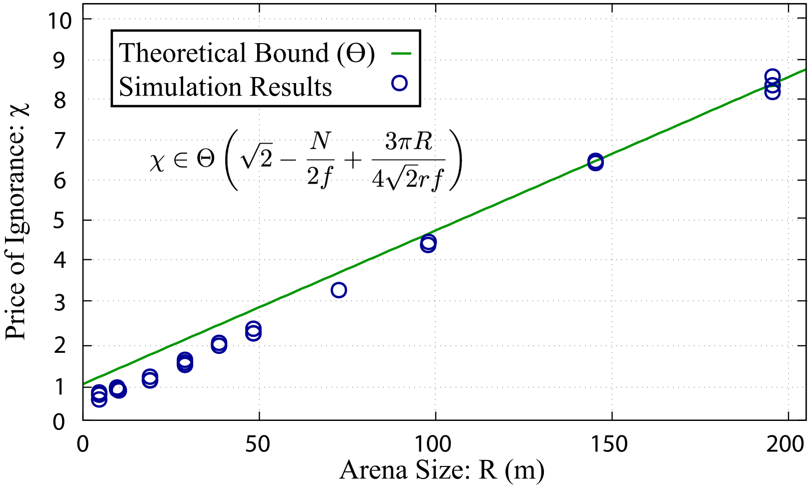

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64.

Testing the Analysis: Spiral CPFA

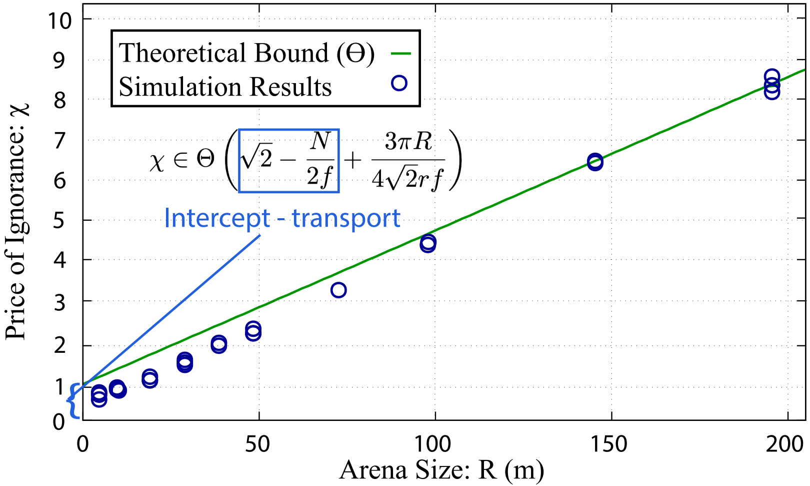

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64.

Testing the Analysis: Spiral CPFA

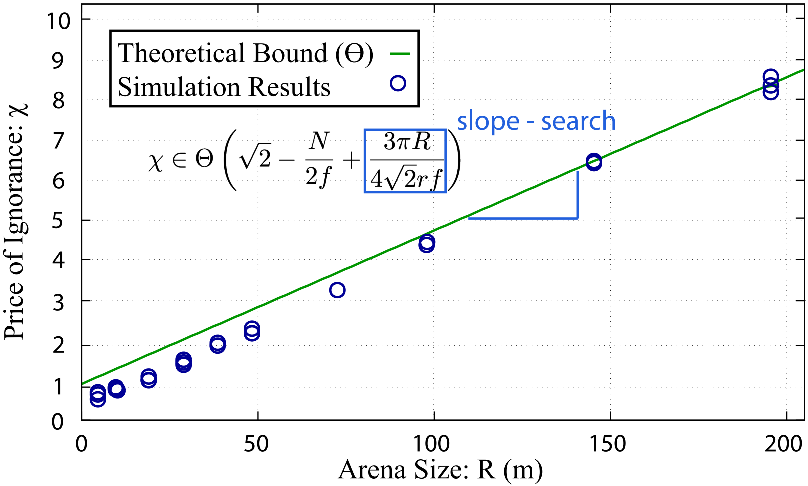

Argos Robot Simulator [7], 5 trials per arena size, 10 robots per trial, 4 clusters of 64.

Conclusions

- For the first time we have a formal understanding of some common CPFAs.

Conclusions

- For the first time we have a formal understanding of some common CPFAs.
- Simulations validate the theoretical bounds.

Conclusions

- For the first time we have a formal understanding of some common CPFAs.
- Simulations validate the theoretical bounds.
- The performance rankings for 100 m foraging arena from theory are:
  1. Spiral CPFAs (4×Perfect)
  2. Spoke CPFAs (10×Perfect)
  3. Random CPFAs (100×Perfect)

  Which is in alignment with the NASA Swarmathon competition results.

Conclusions and Future Work

- See our paper for the full proofs and simulation details. We would love to see improvements on our random and spoke bounds.
- We need algorithmic complexity analysis of information sharing CPFAs. That they are sub-optimal is entirely based on empirical work.
- A way to escape the bounds given here is to introduce heterogeneous robot teams. We have an ICRA submission analysing that situation in review now.

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 retrieval tasks. In biology, immune systems searching for pathogens and ant-colonies searching for targets can also be understood by analysing CPFAs.

Empirical work in real robots leads us to investigate three simple algorithms: the Distributed Archimedean 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, a swarm robot foraging competition that allowed us to make direct comparisons among many different foraging algorithms in simulations and in physical robots [4].

Our 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 (returning to the location of the last found target) and the adverse effects of oversampling (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 an idealised omniscient algorithm[†]. This quantifies the penalty paid by each algorithm for not knowing target locations, which for the complete collection task is the major determinant of CPFA performance. Our preliminary results using this metric appeared in [8].

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 idealised scenarios. To test whether the ranking of upper-bounds holds, we run experiments using the Autonomous Robots Go Swarming (ARGoS) [9] simulator for each of the CPFAs. In combination, the asymptotic analysis and ARGoS simulations 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 [10], Physics [11], Computational Geometry [12] and Robotics [13]–[15] (among other areas). Central place foraging has been of fundamental interest to researchers of Swarm Intelligence because of its deep connections to social insect behaviour [16]. 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, and minimising the mean time to detection is NP-hard [17], [18]. Therefore, search and CPF use heuristics.

Rybski et al. [19], and Hecker and Moses [15] demonstrate that site-fidelity can dramatically reduce search times in robot CPF in simulation and hardware experiments. So far this effect has mostly been argued empirically [20], and our result provides some of the first theoretical evidence to explain this phenomenon.

Ghosh and Klein [21] provide a review of planar search algorithms, which is a critical component of CPFAs. Spiral

[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

[†]We note that the price of ignorance is similar to the notion of *competitiveness* for online algorithms [6], which has been used in the context of foraging in [7]. However, in our application, rather than measuring competitiveness with respect to the amount of advice available (as in [7]) we measure competitiveness with respect to the knowledge of the locations of resources and in the ability of the foraging algorithm to avoid repeatedly searching the same locations in the foraging arena. Additionally, we assume a uniformly random pile placement, which is different from the adversarial setting often analysed for online algorithms.

Conclusions and Future Work

- Why do desert harvester ants use random ballistic search when it is inefficient?
- Our spoke results suggest that by rotating by a **very irrational** number such as the golden ratio you can improve asymptotic efficiency.

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 retrieval tasks. In biology, immune systems searching for pathogens and ant-colonies searching for targets can also be understood by analysing CPFAs.

Empirical work in real robots leads us to investigate three simple algorithms: the Distributed Archimedean 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, a swarm robot foraging competition that allowed us to make direct comparisons among many different foraging algorithms in simulations and in physical robots [4].

Our 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 (returning to the location of the last found target) and the adverse effects of oversampling (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 an idealised omniscient algorithm†. This quantifies the penalty paid by each algorithm for not knowing target locations, which for the complete collection task is the major determinant of CPFA performance. Our preliminary results using this metric appeared in [8].

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 idealised scenarios. To test whether the ranking of upper-bounds holds, we run experiments using the Autonomous Robots Go Swarming (ARGoS) [9] simulator for each of the CPFAs. In combination, the asymptotic analysis and ARGoS simulations 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 [10], Physics [11], Computational Geometry [12] and Robotics [13]–[15] (among other areas). Central place foraging has been of fundamental interest to researchers of Swarm Intelligence because of its deep connections to social insect behaviour [16]. 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, and minimising the mean time to detection is NP-hard [17], [18]. Therefore, search and CPF use heuristics.

Rybski et al. [19], and Hecker and Moses [15] demonstrate that site-fidelity can dramatically reduce search times in robot CPF in simulation and hardware experiments. So far this effect has mostly been argued empirically [20], and our result provides some of the first theoretical evidence to explain this phenomenon.

Ghosh and Klein [21] provide a review of planar search algorithms, which is a critical component of CPFAs. Spiral

[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

† We note that the price of ignorance is similar to the notion of *competitiveness* for online algorithms [6], which has been used in the context of foraging in [7]. However, in our application, rather than measuring competitiveness with respect to the amount of advice available (as in [7]) we measure competitiveness with respect to the knowledge of the locations of resources and in the ability of the foraging algorithm to avoid repeatedly searching the same locations in the foraging arena. Additionally, we assume a uniformly random pile placement, which is different from the adversarial setting often analysed for online algorithms.

Acknowledgements

Abhinav Aggarwal

Will Vining

Diksha Gupta

Melanie Moses

JAMES S. MCDONNELL FOUNDATION

NASA

UNM

Thank you

Questions?

References i

- A. F. T. Winfield, "Foraging Robots," in *Encyclopedia of complexity and systems science*, R. A. Meyers, Ed. New York: Springer, 2009, pp. 3682–3700.
- 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.
- G. M. Fricke, P. H. Joshua, D. G. Antonio, T. T. Linh, and E. M. Melanie, "A Distributed Deterministic Spiral Search Algorithm for Swarms," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on.* Copyright 2016, IEEE reprinted with permission from the authors., 2016, pp. 4430–4436.

- 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.
- 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.
- G. M. Fricke, "Search in T cell and Robot Swarms: Balancing Extent and Intensity," Ph.D. dissertation, University of New Mexico, 2017.

- C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. D. Caro, and F. Ducatelle, "ARGoS: a modular, multi-engine simulator for heterogeneous swarm robotics," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, pp. 5027–5034.
- J. P. Hecker and M. E. Moses, "Beyond Pheromones: Evolving Adaptable, Flexible, and Scalable Ant-Inspired Robot Swarms," 2014.
- 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.

Information Sharing

Strategies that relied on communication tended to do poorly.

- Balancing when to respond to recruitment signals vs exploring is a classic explore/exploit trade-off. One that is further complicated by stochastic environments and noisy sensors.

Information Sharing

Strategies that relied on communication tended to do poorly.

- Balancing when to respond to recruitment signals vs exploring is a classic explore/exploit trade-off. One that is further complicated by stochastic environments and noisy sensors.
- None of the CPFAs that used recruitment performed well in the competition over its 4 years.

Information Sharing

Strategies that relied on communication tended to do poorly.

- Balancing when to respond to recruitment signals vs exploring is a classic explore/exploit trade-off. One that is further complicated by stochastic environments and noisy sensors.
- None of the CPFAs that used recruitment performed well in the competition over its 4 years.
- A recruitment algorithm developed in our own lab that does perform relatively well relied on offline learning to estimate the the optimal trade-off for each target distribution (Hecker and Moses, *Swarm Int.*, 2015[8]).

Information Sharing

Strategies that relied on communication tended to do poorly.

- Balancing when to respond to recruitment signals vs exploring is a classic explore/exploit trade-off. One that is further complicated by stochastic environments and noisy sensors.
- None of the CPFAs that used recruitment performed well in the competition over its 4 years.
- A recruitment algorithm developed in our own lab that does perform relatively well relied on offline learning to estimate the the optimal trade-off for each target distribution (Hecker and Moses, *Swarm Int.*, 2015[8]).
- There is one kind of recruitment always helps performance: self recruitment or *site fidelity*. Site fidelity is implicit in the spoke and spiral algorithms since they return to the place their pattern was interrupted by finding a target.

For a formal treatment of misinformation costs see Pitonakova et al. *Swarm Int.*