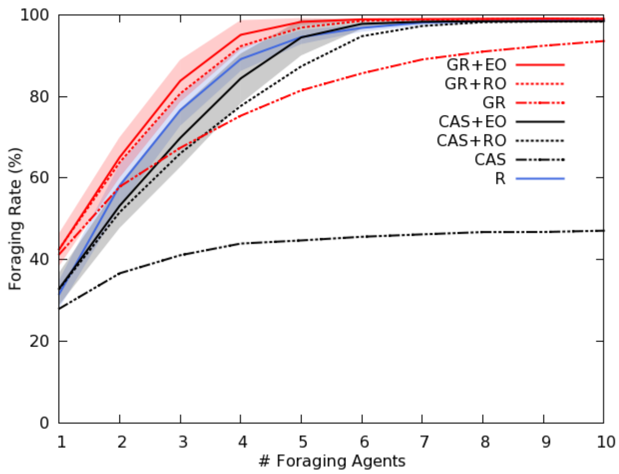

Figure 2: Comparison of our Greedy Rate (GR) algorithm against the benchmark of Continuous Area Sweeping (CAS) when resources follow the Bernoulli model of replenishment.

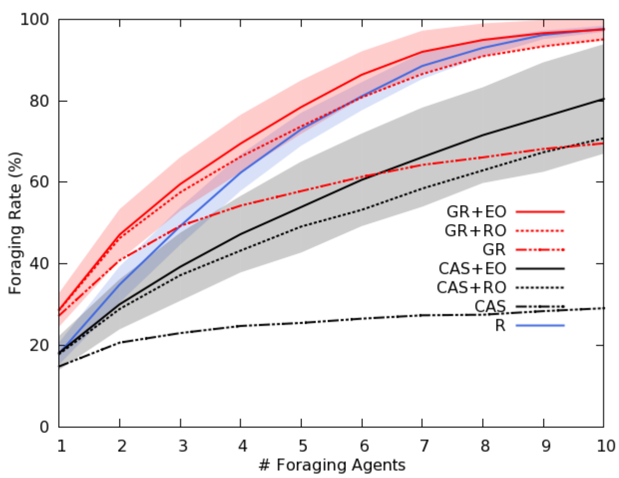

Figure 3: Comparison of our Greedy Rate (GR) algorithm against the benchmark of Continuous Area Sweeping (CAS) when resources follow the Poisson model of replenishment.

in the combination of both the foraging *and* information-gathering algorithms. Previous foraging algorithms generally do not consider incorporating new information from other sources (e.g., the reconnaissance agent), while our foraging algorithms exploit the fact the new information can arrive at any time, and thus improves the overall team foraging rate, as we describe in the next section.

7. EXPERIMENTS AND RESULTS

We describe the extensive experiments we conducted to analyze the performance of our algorithms we introduced in the previous sections. We compare against the baselines from sustainable foraging and continuous area sweeping, and evaluate our Expected Observation algorithm.

7.1 Experimental Setup

The foraging locations were randomly generated, following a uniform distribution over a square of size $N \times N$, and

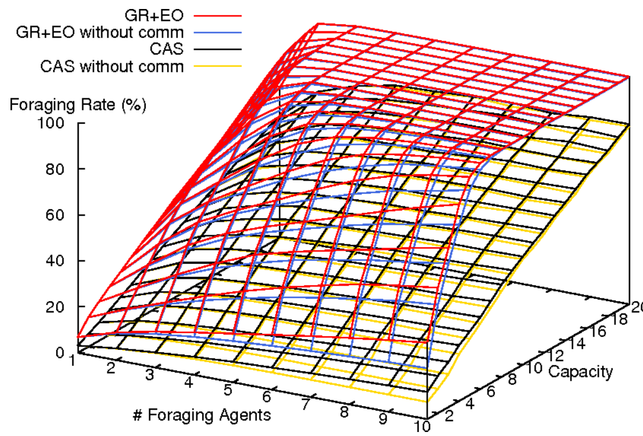

Figure 4: Effect of communication among foraging agents when resources follow the Bernoulli model of replenishment.

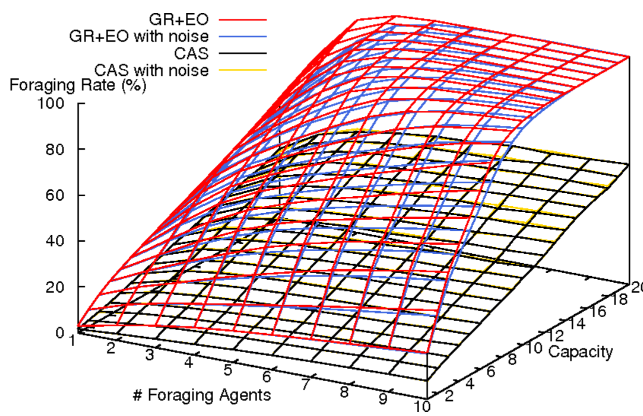

Figure 5: Effect of observational noise when resources follow the Poisson model of replenishment.

either followed the Bernoulli, Poisson, or the stochastic Logistic models. The agents' initial positions were randomly generated to be also be within the $N \times N$ square. In each experiment, we simulated $T = 1000$ timesteps, and recorded the number of resources $v_{0,T}$ foraged to the home l_0.

We varied the number of agents n from 1 to 10, and the capacities from 1 to 20. We chose 10 and 20 because it was sufficient for a Random foraging algorithm (i.e., agents that randomly select their destinations) to forage almost all resources in the Bernoulli and Poisson scenarios. We set the number of locations $|\mathcal{L}| = 20$ (since it is twice the number of agents) and set the number of location that the reconnaissance agent could visit at each timestep to be $M = \frac{|\mathcal{L}|}{4}$.

As a baseline, we assumed that foraging agents were capable of limited communication when they were within $\frac{N}{10}$ distance, and could communicate their destinations and payloads. We assumed that observations are not noisy as a baseline. We investigate the effects of no communication and noisy observations in the Bernoulli and Poisson models.

7.2 Experiments with Bernoulli and Poisson Models of Replenishment

We compared our Greedy Rate (GR) algorithm against the Continuous Area Sweeping (CAS) algorithm [1]. As a baseline, we used a Random foraging (R) algorithm where agents randomly select their destination.

Both our GR algorithm and the CAS algorithm can use information gathered by the reconnaissance agent, and we compared our Expected Observation (EO) algorithm to a Random Observation (RO) algorithm, where M locations would be randomly chosen by the reconnaissance agent.

Figures 2 and 3 show the performance of our algorithms when the capacities of the agents are 10. Since the number of resources generated in simulation varied, we measured the percentage of resources that were successfully foraged at the end of the experiment. The solid red, black and blue lines show Greedy Rate with Expected Observation (GR+EO), Continuous Area Sweeping with Expected Observation (CAS+EO), and Random Foraging (R), and the shaded areas show the standard deviations of these algorithms. The dotted and dashed lines show other combinations of foraging and reconnaissance algorithms.

As the number of agents increase, R outperforms both GR and CAS ($p = 1 \times 10^{-24}$ and $p = 3 \times 10^{-27}$ with a 2-tailed T-test on R vs GR in Bernoulli and Poisson respectively), primarily because the agents do not share their models, so agents tend to head to similar locations. Even though agents coordinate when possible, the limited range of communication causes inefficiencies in foraging.

However, the introduction of a reconnaissance agent substantially improves both GR and CAS. Our EO algorithm outperforms the RO algorithm ($p = 1 \times 10^{-15}$ and $p = 2 \times 10^{-39}$ for GR+EO vs GR+RO on Bernoulli and Poisson respectively), and GR+EO outperforms CAS+EO and R ($p = 3 \times 10^{-26}$ and $p = 1 \times 10^{-29}$ for Bernoulli, and $p = 5 \times 10^{-36}$ and $p = 1 \times 10^{-31}$ for Poisson). It is interesting to note that CAS+EO performs substantially better than the baseline CAS. In general, adding a single reconnaissance agent with EO provides a much higher benefit than increasing the number of foraging agents.

We investigated having no communication (among the foraging agents) and noisy observations. Figures 4 and 5 shows the effects as the number of agents and their capacities vary. While a lack of communication and noisy observations affect our algorithms, the effect is minimal (a median of 0.3% and 2.2% respectively for communication and noise, thus illustrating that our algorithms are robust to a lack of communication and noisy observations.

In addition, Figures 4 and 5 clearly illustrate the efficacy of our GR+EO algorithms over the baseline CAS, across all numbers of foraging agents and agent capacities. We chose 10 to be the maximum number of foraging agents, and 20 to be the maximum capacity, because our algorithm have already hit the 100% foraging rate before that point. In contrast, CAS does not reach 100% even with 10 foraging agents with a capacity of 20 each.

7.3 Experiments with Stochastic Logistic Model

We compared our Adaptive Sleep (AS), Adaptive Sleep with Target Change (ASTC) algorithms against Sustainable Foraging (SF) [17], and a Random (R) foraging algorithm as baseline. Only our algorithms could use information gathered by the reconnaissance agent.

Fig. 6 shows the algorithms' performance when the agent capacities are 20, and the stochastic Logistic noise is $\sigma_e = 0.08$. The shaded regions show the standard deviations of

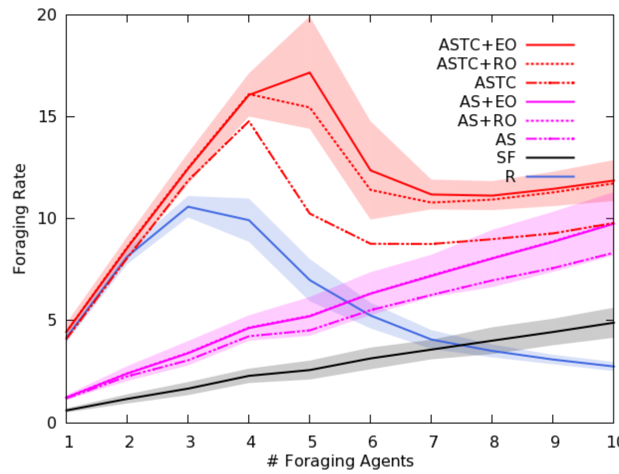

Figure 6: Comparison of our Adaptive Sleep (AS) and Adaptive Sleep with Target Change (ASTC) algorithms against the benchmark of Sustainable Foraging (SF) when resources follow the stochastic Logistic model of replenishment.

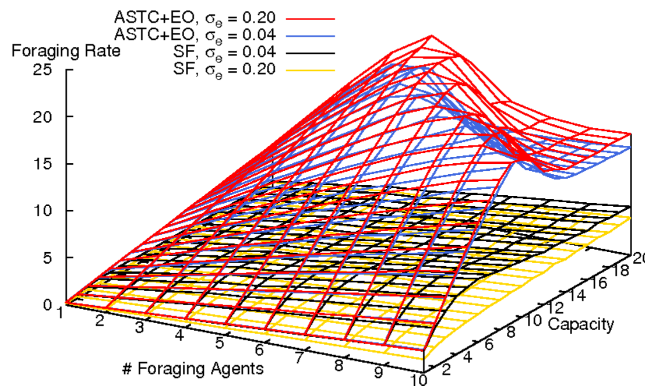

Figure 7: Effect of noise in the stochastic Logistic model.

ASTC+EO, AS+EO, SF and R. SF and AS both increase linearly, since the agents select a single destination; our AS algorithm outperforms SF ($p = 7 \times 10^{-61}$).

R outperforms AS and SF when the number of agents are small, primarily because changing destinations allows resources to replenish at a higher rate. However, as the number of agents increase, R's performance begins to plummet as locations become over-foraged and the replenishment rate lowers. ASTC combines the benefits of AS and R, allowing agents to choose a destination, and also visit other unassigned locations. Thus, the shape of the ASTC curve is similar to R, albeit at a much higher foraging rate.

The introduction of a reconnaissance agent improves the foraging rate. EO and RO perform similarly with the AS algorithm. For ASTC, EO and RO perform similarly when the number of agents $n < 5$ ($p = 0.14$), but EO outperforms RO when there are $n \geq 5$ ($p = 8 \times 10^{-6}$). The constant 5 corresponds to the number of locations the reconnaissance agent visits: EO determines the best locations to visit, compared to RO's random choice. When $n < 5$, there is a high probability that RO visits all agents' locations.

Fig. 7 shows the effect of stochastic Logistic noise ($\sigma_e = 0.04$ to 0.20). SF performs poorly as the noise increases, but ASTC+EO performs better with higher noise, showing that our algorithm takes advantage of the noise (noise creates a probability of generating resources ahead of schedule).

7.4 Summary of Experimental Analysis

Across the Bernoulli and Poisson models of replenishment, our Greedy Rate (GR) algorithm outperforms the baseline Continuous Area Sweeping (CAS). Further, the addition of the reconnaissance agent improves the performance of both GR and CAS, since additional information is provided to both algorithms. Our Expected Observation (EO) algorithm outperforms the Random Observation (RO) algorithm, showing that although the reconnaissance agent can visit $\frac{|\mathcal{L}|}{4}$ of the locations each timestep, selecting which locations to visit still plays a very important role. Random selection (which will visit every location every 4 timesteps on average) improves the team foraging rate, but not as much.

In addition, it is important to note that our EO algorithm significantly improves the CAS algorithm's foraging rate, so our information-gathering algorithm is not specific to our foraging algorithms, but can be applied to any foraging algorithm that makes use of new information. Also, our algorithms are robust to noise in observations, and performs with minimal degradation when communication among the foraging agents are unavailable.

Similarly, for the stochastic Logistic model of replenishment, our Adaptive Sleep (AS) and Adaptive Sleep with Target Change (ASTC) algorithms outperform the baseline of Sustainable Foraging (SF), across all numbers of foraging agents and agent capacities. Our algorithm is robust to the noise in the stochastic Logistic model, and the EO algorithm improves our foraging algorithms significantly. The ASTC algorithm incorporates both the features of the AS algorithm (to maximize the foraging rate at the assigned location) and the Random algorithm (to exploit the resource replenishment at unassigned locations).

Thus, our experiments show that our algorithms are distributed and require little communication among the foraging agents, and are robust to noise in observations, a lack of communication, and noise in the models. We outperform the baselines significantly, and the addition of the reconnaissance agent improves the multi-agent team's foraging rate, even for foraging algorithms that we did not create.

8. CONCLUSION

We formally defined the continuous foraging problem, where agents visit known foraging locations to collect and deliver resources to a home location. The resources replenish over time, and we defined three models of resource replenishment: the Bernoulli and Poisson models where resources replenish probabilistically (e.g., mail entering a mailbox), and a stochastic Logistic model where the rate of resource replenishment depends on the number of existing resources (e.g., a population of fish).

We considered two types of agents: foraging agents that actively forage resources, and a reconnaissance agent that cannot forage items, but can visit a subset of the locations to determine the number of resources available.

We contributed algorithms for the foraging and reconnaissance agents, and to evaluate our algorithms, we performed experiments in simulation, benchmarking against existing algorithms in sustainable foraging and continuous area sweeping. We showed that our algorithms outperform the existing ones even without the use of the reconnaissance agent. Further, we demonstrated that the reconnaissance agent further improves the foraging rate of the multi-agent team, even in the presence of noisy observations and no communication among the foraging agents, thus illustrating the benefits of our algorithms.

Acknowledgments

This work was supported by the A*STAR Computational Resource Centre through the use of its high performance computing facilities. The authors also thank Wei Li Cheah for her help with visualization of the results.

REFERENCES

[1] M. Ahmadi and P. Stone. Continuous Area Sweeping: A Task Definition and Initial Approach. In *Proc. Int. Conf. Advanced Robotics*, pages 316–323, 2005.

[2] S. Alers, D. Bloembergen, D. Hennes, S. de Jong, M. Kaisers, N. Lemmens, K. Tuyls, and G. Weiss. Bee-inspired Foraging in an Embodied Swarm. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 1311–1312, 2011.

[3] M. Dorigo and M. Birattari. Ant colony optimization. In *Encyclopedia Machine Learning*, pages 36–39. 2010.

[4] B. P. Gerkey and M. J. Mataric. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *J. Robotics Research*, 23(9):939–954, 2004.

[5] H. Hakoyama and Y. Iwasa. Extinction Risk of a Density-dependent Population Estimated from a Time Series of Population Size. *J. Theoretical Biology*, 204(3):337–359, 2000.

[6] N. Hoff, A. Sagoff, R. Wood, and R. Nagpal. Two Foraging Algorithms for Robot Swarms using only Local Communication. In *Proc. Int. Conf. Robotics and Biomimetics*, pages 123–130, 2010.

[7] N. Lemmens, S. Jong, K. Tuyls, and A. Nowe. Bee Behaviour in Multi-agent Systems. In *Adaptive Agents and Multi-Agent Systems III*, volume 4865, pages 145–156. 2008.

[8] K. Lerman, C. Jones, A. Galstyan, and M. J. Mataric. Analysis of Dynamic Task Allocation in Multi-Robot Systems. *J. Robotics Research*, 25(3):225–241, 2006.

[9] S. Liemhetcharat and M. Veloso. Modeling and Learning Synergy for Team Formation with Heterogeneous Agents. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 365–375, 2012.

[10] S. Liemhetcharat and M. Veloso. Weighted Synergy Graphs for Role Assignment in Ad Hoc Heterogeneous Robot Teams. In *Proc. Int. Conf. Intelligent Robots and Systems*, pages 5247–5254, 2012.

[11] S. Liemhetcharat and M. Veloso. Forming an Effective Multi-Robot Team Robust to Failures. In *Proc. Int. Conf. Intelligent Robots and Systems*, 2013.

[12] S. Liemhetcharat and M. Veloso. Synergy Graphs for Configuring Robot Team Members. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 111–118, 2013.

[13] S. Liemhetcharat and M. Veloso. Weighted Synergy Graphs for Effective Team Formation with Heterogeneous Ad Hoc Agents. *Journal of Artificial Intelligence*, 208(2014):41–65, 2014.

[14] W. Liu, A. Winfield, J. Sa, J. Chen, and L. Dou. Towards Energy Optimization: Emergent Task Allocation in a Swarm of Foraging Robots. *Adaptive Behavior*, 15(3):289–305, 2007.

[15] L. Panait and S. Luke. A Pheromone-Based Utility Model for Collaborative Foraging. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, pages 36–43, 2004.

[16] D. A. Shell and M. J. Mataric. On Foraging Strategies for Large-Scale Multi-Robot Systems. In *Proc. Int. Conf. Intelligent Robots and Systems*, pages 2717–2723, 2006.

[17] Z. Song and R. Vaughan. Sustainable Robot Foraging: Adaptive Fine-Grained Multi-Robot Task Allocation for Maximum Sustainable Yield of Biological Resources. In *Proc. Int. Conf. Intelligent Robots and Systems*, pages 3309–3316, 2013.

[18] L. Vig and J. Adams. Market-based Multi-Robot Coalition Formation. In *Proc. Int. Symp. Distributed Autonomous Robotics Systems*, pages 227–236, 2006.

[19] E. Yoshida, T. Arai, J. Ota, and T. Miki. Effect of Grouping in Local Communication System of Multiple Mobile Robots. In *Proc. Int. Conf. Intelligent Robots and Systems*, pages 808–815, 1994.