Spectral-Based Distributed Ergodic Coverage for Heterogeneous Multi-Agent Search

Guillaume Sartoretti[1,*], Ananya Rao[2,*], and Howie Choset[2]

[1] Mechanical Engineering Dept., National University of Singapore, SG 117575.
mpegas@nus.edu.sg,
WWW home page: http://www.marmotlab.org
[2] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15203, USA

Abstract. This paper develops a multi-agent heterogeneous search approach that leverages the sensing and motion capabilities of different agents to improve search performance (i.e., decrease search time and increase coverage efficiency). To do so, we build upon recent results in ergodic coverage methods for homogeneous teams, where the search paths of the agents are optimized so they spend time in regions proportionate to the expected likelihood of finding targets, while still covering the whole domain, thus balancing exploration and exploitation. This paper introduces a new method to extend ergodic coverage to teams of heterogeneous agents with varied sensing and motion capabilities. Specifically, we investigate methods of leveraging the spectral decomposition of a target information distribution to efficiently assign available agents to different regions of the domain and best match the agents' capabilities to the scale at which information needs to be searched for in these regions. Our numerical results show that distributing and assigning coverage responsibilities to agents based on their dynamic sensing capabilities leads to approximately 40% improvement with regard to a standard coverage metric (ergodicity) and a 15% improvement in time to search over a baseline approach that jointly plans search paths for all agents, averaged over 500 randomized experiments.

Keywords: multi-agent system, distributed search, heterogeneous teams

1 Introduction

With the rapid development of affordable robots with embedded sensing and computation capabilities, we are quickly approaching a point at which real-life applications will involve the deployment of hundreds, if not thousands, of robots [1, 2]. Among these applications, significant research effort has been devoted to multi-agent search [3, 4, 5, 6, 7], where deploying numerous agents can greatly improve the time-efficiency and robustness of search. In fact, deploying robots with various motion or sensing modalities can further improve the search performance, by leveraging the natural synergies between these capabilities (see

[*] These authors contributed equally to this work.

Fig. 1). Motivated by such problems, the main contribution of this work is to investigate the distribution of agents to search a domain at various spatial scales, based on their motion and sensing capabilities. Specifically, we build upon recent results in ergodic search processes [3, 8, 9, 10] to propose a mapping of available agents to a spectral-based decomposition of the search problem, to best match agents' capabilities to specific classes of areas of the search domain.

The approach in this paper is based on ergodic search processes, which, similar to other information-theoretic coverage methods [3, 4, 8, 11, 12], rely on an *a priori* information distribution, representing the likelihood of finding a target at any point over the search domain, to guide the search. In practice, this information distribution can be obtained from scouting missions or from expert knowledge, and is updated during search if inaccurate. If nothing is known about the targets' whereabouts, this distribution is usually considered uniform over the whole domain, i.e., targets could be anywhere with equal probability. Using this *a priori* information distribution, ergodic search processes optimize search paths over long time horizons for all agents.

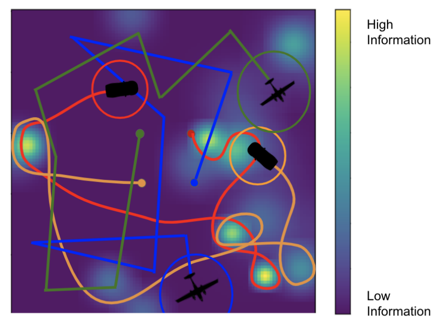

Fig. 1. Multi-agent search scenario involving two types of agents: differential-drive agents with short-range, high fidelity sensors (represented by the red and orange circles), and omnidirectional agents with long-range, low fidelity sensors (represented by the blue and green circles). The different colored lines represent the paths followed by the different agents. The underlying distribution shows the likelihood of finding targets throughout the domain.

In this work, we determine search paths via an optimization process using the ergodic metric [3]. The optimization of search paths, according to the ergodic metric, aims to drive agents to spend time in areas of the domain in proportion to the *a priori* likelihood of finding targets in these areas. This optimization is performed in the spectral domain, by minimizing the difference between the coefficients associated to the team's *time-average statistics* (i.e., fraction of the time spent in each area) and those of the information distribution. The contribution of this work is to exploit the spectral nature of the ergodic metric in search scenarios involving heterogeneous agents. To this end, we plan search paths for each agent type based on a smaller subset of the spectral coefficients associated with the information map, thus driving agents to search the domain at a spatial scale that best matches their motion and sensing capabilities.

This paper is organized as follows: Section 2 discusses recent advances in multi-agent search and in coordination of heterogeneous multi-agent systems. We then provide a brief background of ergodic search processes in Section 3. Section 4 details our spectral-based decomposition of a search problem and of the available agents. We then present and discuss the results of our systematic set of experiments in Section 5. There, we observe that our best agent distribution approach leads to a 40% increase in coverage performance (thus generally leading to more time-efficient search), averaged over 500 randomized experiments, over a baseline that jointly plans search paths for the whole team regardless of their individual abilities. Section 6 offers concluding remarks.

2 Prior Work

2.1 Multi-Agent Search

Current active search methods generally fall into one of three main categories: geometric, gradient-based, and trajectory optimization-based approaches. Geometric methods, e.g., lawnmower patterns, can be good search strategies in order to uniformly cover a domain in which there is near-uniform probability of finding a target [13, 14]. Since these approaches exhaustively cover the search domain, they are also the logical choice in cases where there is no *a priori* information about the targets' locations.

An information map, or information distribution, is defined to be a probability distribution representing the likelihood of a target being found at each location in the domain. When such *a priori* information is available (and, usually, non-uniform), more advanced search processes can be created that leverage this information map in order to improve search according to some metric, such as time to find all targets.

For example, in gradient-based, or "information surfing", methods [4, 11, 12], agents guide their movement in the direction of the derivative of the information map around their positions to greedily maximize the short-term information gain. That is, agents are always driven in the direction of the greatest information gain, which naturally leads them to areas where the likelihood of finding a target is maximized. Information surfing can be implemented in a fully decentralized manner, since it does not require tight coordination between agents, and potential fields can be introduced to help distribute agents to different areas of the domain. However, gradient-based approaches generally do not rely on the uncertainty associated with the information distribution, which can lead to areas left unexplored, as this uncertainty can help differentiate areas of low-information that have not been explored from areas with no information to be gained. Gradient-based approaches are also very sensitive to noise in the information map, as the gradient cannot be estimated accurately in these situations, and suffer from greedily over-exploiting local information maxima.

Optimization-based approaches look at search as an information gathering maximization problem, which is then solved by planning (usually joint) paths for the agents. Several recent works in coverage methods [9, 3, 10, 8] rely on sampling-based path planning, where a large number of paths are sampled and the best path is chosen based on a cost metric. Optimization-based approaches can combine both the predicted information distribution as well as its associated uncertainty into the cost function that drives the optimization. However, these approaches generally do not scale well for large multi-agent systems since they remain centralized. Even for sampling-based approaches, the number of paths that need to be sampled to find near-optimal search paths grows exponentially with the number of agents, although growing the number of samples linearly with the team size seems to experimentally provide good-quality search paths [9, 3].

2.2 Heterogeneity in MAS

Most search methods developed for homogeneous groups of agents (i.e., agents with similar capabilities) do not support groups of agents with heterogeneous capabilities (for example, a set of agents with different sensing or motion capabilities), or struggle with the increased computational complexity [15, 16, 17, 18, 19]. Many previous works involving heterogeneous agents concentrate on offering initial, usually centralized and non-scalable solutions to the problems that they mainly focus on defining [20, 21, 22, 23]. Some works have considered using auction-based mechanisms for task assignment in heterogeneous groups [24, 25, 26], while others have proposed agent redistribution based on given sets of their capabilities [27, 28]. Other works have considered using robots with the best communication or coordination capabilities as "leader agents" to plan for and coordinate the other agents with lesser capabilities [6, 29].

3 Background on Ergodic Search Processes

Ergodic search processes [8] produce trajectories for multi-agent systems, such that agents spend time in each area of the domain proportional to the expected amount of information present in this area. To this end, the spatial time-average statistics of an agent's trajectory $\gamma_i : (0,t] \rightarrow \mathcal{X}$, quantifies the fraction of time spent at a position $\boldsymbol{x} \in \mathcal{X}$, where $\mathcal{X} \subset \mathrm{I\!R}^{\mathrm{d}}$ is the d-dimensional search domain. For N agents, the joint spatial time-average statistics of the set of agents trajectories $\{\gamma_i\}_{i=1}^{N}$ is defined as [8]

$$C^t(\boldsymbol{x}, \gamma(t)) = \frac{1}{Nt} \sum_{i=1}^{N} \int_0^t \delta(\boldsymbol{x} - \gamma_i(\tau)) \, d\tau, \tag{1}$$

where δ is the Dirac delta function.

Formally, the agents' time-averaged trajectory statistics is optimized against the expected information distribution over the whole domain, by matching their spectral decompositions. This is obtained by minimizing the ergodic metric $\Phi(\cdot)$, expressed as the weighted sum of the difference between the spectral coefficients of these two distributions [8]:

$$\Phi(\gamma(t)) = \sum_{k=0}^{m} \lambda_k \left| c_k(\gamma(t)) - \xi_k \right|^2, \tag{2}$$

where c_k and ξ_k are the Fourier coefficients of the time-average statistics of the set of agents' trajectories $\gamma(t)$ and the desired spatial distribution of agents respectively, and λ_k are the weights of each coefficient difference. In practice, $\lambda_k = \sqrt{(1 + \|k\|^2)^{-(d+1)}}$ is usually defined to place higher weights on the lower frequency components, which correspond to larger spatial-scale variations in the information distribution.

The goal of ergodic coverage is to generate optimal controls $\boldsymbol{u}^*(t)$ for each agent, whose dynamics is described by a function $f\colon \mathcal{Q} \times \mathcal{U} \to \mathcal{T}\mathcal{Q}$, such that

$$\begin{aligned} \boldsymbol{u}^*(t) \quad & = \arg\min_{\boldsymbol{u}} \Phi(\gamma(t)), \\ \text{subject to } \dot{\boldsymbol{q}} & = f(\boldsymbol{q}(t), \boldsymbol{u}(t)), \\ \|\boldsymbol{u}(t)\| \quad & \leq u_{max} \end{aligned} \tag{3}$$

where $\boldsymbol{q} \in \mathcal{Q}$ is the state and $\boldsymbol{u} \in \mathcal{U}$ denotes the set of controls. Eq.(3) can either be solved by discretizing the exploration time and solving for the optimal control input at each time-step [8], by trajectory optimization to plan feed-forward trajectories over a specified time horizon [30], or by using sampling-based motion planners [31], where it is straightforward to pose additional constraints such as obstacle avoidance.

4 Distributed Heteregeneous Ergodic Search

This work investigates the coordination of a team of heterogeneous agents during search from two key fronts. First, we look at how the spectral decomposition of the information distribution can be interpreted in order to search regions at different spatial scales. In other words, lower frequency components typically describe the distributions in broad strokes, while higher frequency ones are responsible for filling in the details. Second, we study how agents should be assigned to different spectral bands of the information decomposition, and formulate an assignment that reasons about the agents' varying capabilities to cover a domain more efficiently.

4.1 Spectral Bands of the Information Distribution

In this work, we rely on the spectral decomposition of the information map to guide the search task assignment for agents with heterogeneous motion and sensing capabilities. We recall that, in the spectral decomposition of the information distribution Eq.(2), lower-frequency coefficients correspond to larger-scale variations in the spatial distribution of information, while higher-frequency coefficients correspond to smaller-scale variations.

Building upon this observation, we propose to define M spectral bands (i.e., sets of frequency coefficients within particular ranges), with M the number of agent types in the heterogeneous team. Each band can be seen as a separate (although not completely independent) **search subtask** that can be distributed to a specific type of agent based on its motion/sensing capabilities to search at a given spatial scale. In this work, we break down the overall set of spectral coefficients into M successive bands of equal length, but other decompositions of the set of coefficients into bands could be considered, and will be investigated in future works. To help visualize the different search subtasks, resulting from such a choice of bands, we can reconstruct partial representations of the information distribution, each based on a single band of coefficients, as shown in Fig. 2.

To formally define and use these search subtasks, we modify the ergodic metric Eq.(2) to rely on a specific band of coefficients only:

$$\Phi(\gamma(t)) = \sum_{k=c_1}^{c_2} \lambda_k \left| c_k(\gamma(t)) - \xi_k \right|^2 , \tag{4}$$

where c_1 and c_2 define the first and last coefficients of the spectral band. Note that the same result could be achieved by setting $\lambda_k = 0 \; \forall k < c_1, k > c_2$.

4.2 Assignment of Agents to Spectral Bands

In order to find an optimized set of paths according to the new ergodic coverage metric Eq.(4), we must first find an optimized assignment of agent types to spectral bands which maximizes the search performance. To this end, we note that these bands can first be distributed to heterogeneous agents based on their sensing capabilities. For example, agents with low-fidelity, high-range sensing capabilities should generally be assigned low-frequency spectral coefficients, as expressed in Eq.(5) in order to perform large-scale, broad-stroke exploration. Conversely, agents with high-fidelity, low-range sensing capabilities should likely be assigned high-frequency spectral coefficients, as expressed in Eq.(6) in order to perform detailed, small-scale exploration.

$$\Phi(\gamma(t)) = \sum_{k=c_1}^{\frac{c_1+c_2}{2}} \lambda_k \left| c_k(\gamma(t)) - \xi_k \right|^2 , \tag{5}$$

$$\Phi(\gamma(t)) = \sum_{k=\frac{c_1+c_2}{2}}^{c_2} \lambda_k \left| c_k(\gamma(t)) - \xi_k \right|^2 , \tag{6}$$

Similarly, different motion models can also be used as a basis for task distribution between the agents. For instance, we believe that faster agents, relying on lower-frequency coefficients of the decomposition, could perform a coarse exploration of the domain. On the other hand, by relying on higher-frequency

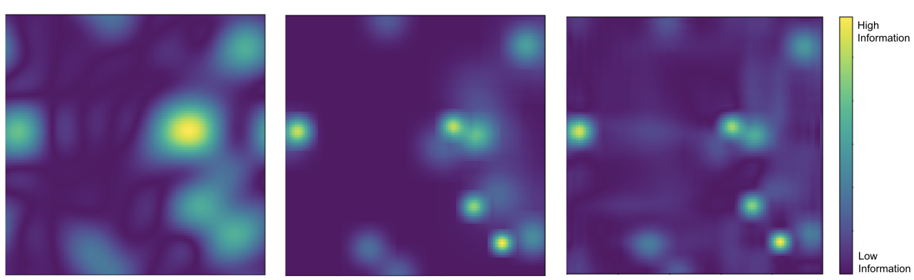

Fig. 2. Example spectral reconstruction of a given map (center), based on only its lower-order coefficients only (left), or higher-order ones only (right). Yellow regions correspond to regions of high information, while darker blue regions correspond to regions of low information (here, high/low likelihood of finding targets).

coefficients, slower agents would be naturally driven to perform a smaller-scale, detailed search. A similar intuitive assignment can be made when dealing with agents with varying motion constraints. For example, given a team composed of omnidirectional and maximal-curvature-constrained agents, our hypothesis is that the former should rely on lower-frequency coefficients, while the latter can more easily chain “spots” of information obtained from high-frequency bands.

5 Results and Discussion

We present our systematic investigation of four different ways agent types can be assigned to search subtasks, by relying on a large set of simulation experiments composed of fixed, randomly generated search problems. We compare these assignment methods according to various standard search metrics, such as the time to find all targets and the effectiveness of coverage (using the ergodic metric), showing that our optimal assignment can yield up to 40% increase in these metrics. Our results rely on sampling-based trajectory optimization, but we emphasize that our investigation should extend to other optimization methods.

5.1 Agent’s Sensing and Motion Models

The sensor footprint of each agent is modeled as a Gaussian distribution centered at the agent’s position, whose variance prescribes a circular observation range $\rho > 0$. At each point within this observation range, we use the Gaussian probability density function to represent the likelihood of detecting a target at each time step. We consider a mix of agents with low-range, high-fidelity sensors (i.e., a Gaussian of low variance and thus higher maximal detection probability at its center), and agents with high-range, low-fidelity sensors (i.e., a Gaussian of larger variance and thus lower maximal peak detection likelihood).

In addition to the different sensing models, we also consider two types of agents’ motion models. The first model is a simple first order integrator that represents omnidirectional agents, such as quad-rotor UAVs or legged ground robots. We further consider agents with differential drive constraints (i.e., resulting in curved paths with a maximum curvature), such as fixed-wing airplanes or wheeled ground vehicles. We sample paths for the agents by sequencing path primitives - straight lines of various directions and lengths for the omnidirectional agents, and curves from a finite collection with various curvatures and lengths for the differential agents. Agents plan long trajectories, execute these paths for 10 timesteps, update the map using their observations, and then replan. We further rely on a cross-entropy planner [31] to optimize the paths of all agents via 3 levels of sample refinement with a total of $15 \cdot N$ samples (where N is the total number of agents, $N = 10$ in practice).

5.2 Experiment Details

Scenarios Randomization We compare the performance of various assignment methods, with that of a baseline that plans paths for all agents by relying on

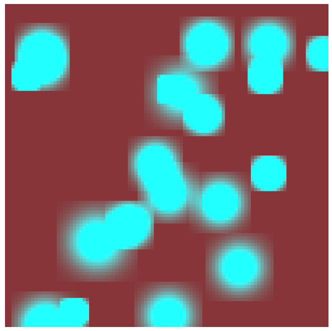

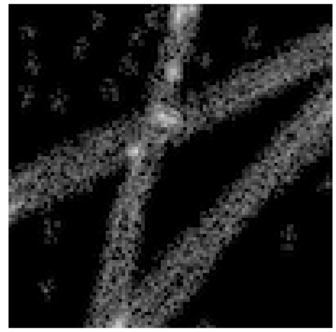

Fig. 3. Examples of the two classes (Gaussian Mixture Models (left), and road networks (right)) of randomly generated information maps for evaluating the proposed approach.

the overall distribution maps (i.e., no decomposition into bands or assignments), through 500 randomized search scenarios. These scenarios vary the locations of targets and the initial information maps (as randomly generated Gaussian mixture models, or road-network inspired information maps). Additionally, for each experiment, a randomly generated team of 10 agents is formed by selecting both the sensing and motion model with equal probability for each agent. Team compositions, starting positions, initial information maps, and target locations are kept identical among experiments with different controllers, to ensure our results are comparable.

Agent Assignments In our experiments, we group together all agents with the same motion and sensing constraints, thus yielding four independent agent groups. We let A_0 be the set of agents with low-fidelity, high range sensing capabilities and fast, omni-directional motion models and A_1 be the set of agents with low-fidelity, high range sensing capabilities and slower, curve-constrained motion models. Similarly, we let A_2 and A_3 be the sets of agents with high-fidelity, low-range sensing capabilities and omni-directional motion models and low-fidelity, high-range sensing capabilities and curve-constrained motion models respectively. Finally, we decompose the information map into a set C of M spectral coefficient bands - $C_0, ..., C_{M-1}$.In our experiments, we considered $M = 4$. Assignments, i.e., mappings from agent types to spectral bands can be expressed as:

$$h : A_i \longrightarrow X, \qquad X \in \{C_0, C_1, C_2, C_3\} \tag{7}$$

Our optimal assignment, based on the intuition built in Section 4.2, can be expressed as:

$$h_{\text{optimal}}(A_i) = C_i \tag{8}$$

In order to investigate the effectiveness of the optimal assignment Eq.(8), we compared its performance to that of more naive assignments Eq.(9), Eq.(10), as well as to that of an adversarial assignment Eq.(11). Finally, we compare these

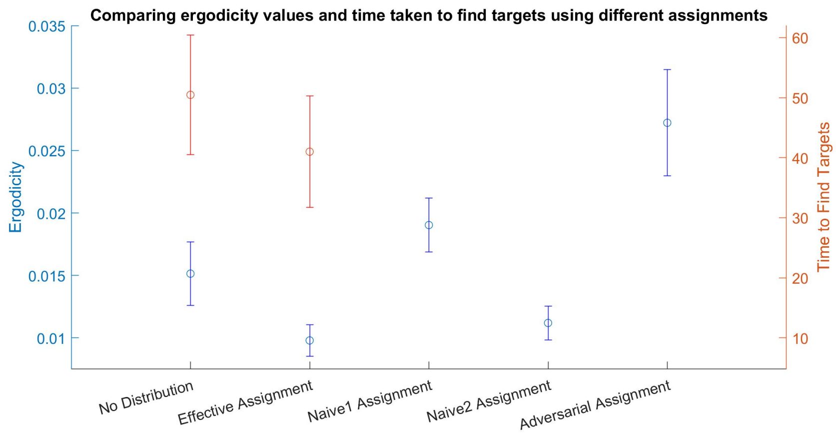

Fig. 4. Search performance comparison between the different agent assignments and the baseline, in terms of coverage performance (using the ergodic metric, lower is better) and time to find all targets (lower is better).

results with a baseline that assumes all agents are identical (homogeneous team), i.e., all the agents rely on the same, full spectral distribution of the information map.

$$h_{\text{naive1}}(A_i) = \begin{cases} C_0 & i = 1 \\ C_1 & i = 0 \\ C_2 & i = 3 \\ C_3 & i = 2 \end{cases} \tag{9}$$

$$h_{\text{naive2}}(A_i) = \begin{cases} C_0 & i = 2 \\ C_1 & i = 3 \\ C_2 & i = 0 \\ C_3 & i = 1 \end{cases} \tag{10}$$

$$h_{\text{adversarial}}(A_i) = C_{3-i} \tag{11}$$

Performance Metrics and Sensitivity to Hyper-parameters We ran another set of experiments in order to investigate the sensitivity of this approach to hyper parameters. There, we focused on the ratio of agents of different capabilities, the total number of agents exploring the domain and the number of samples used in each time step of the sample-based path planner.

All of these experiments were run on Gaussian information distributions and road-network inspired information distributions (Fig.3) in an effort to simulate potential use cases. The results of these two types of information distributions are reported separately as their metrics distributions are significantly different, although the overall performance improvement is similar.

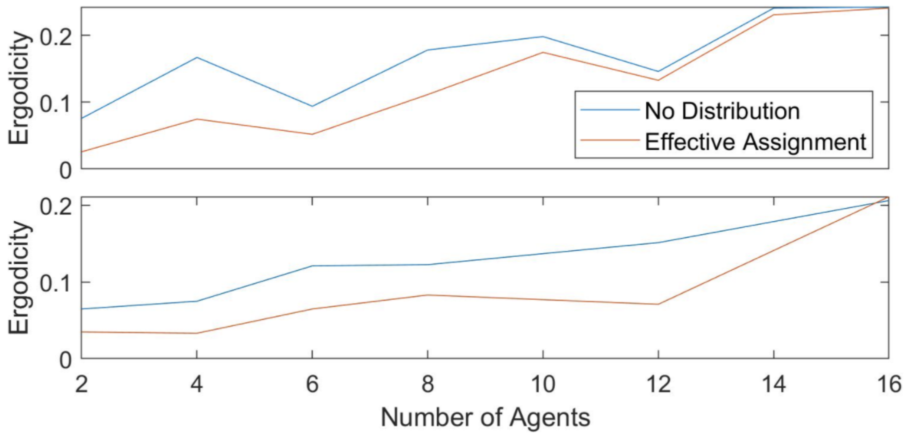

Fig. 5. Sensitivity to the team size, comparing the coverage performance between our optimal assignment and the baseline. Gaussian maps (top), and road network (bottom). As expected, note the improved performance for smaller teams.

5.3 Experimental Results

When looking at the results of the different assignments in term of overall coverage performance, measured via the ergodic metric, the first observation we can make is that our optimal assignment, expressed in Eq.(4.2) results in approximately 40% improvement over the baseline, while naive heuristics show $\pm 5\%$ improvement and the adversarial heuristic yields approximately 25% deterioration in performance over the baseline approach (Fig.4). These results verify the intuition built in Section 4.2. Our results also confirm an important point: more effective coverage of the domain leads to finding targets faster.

As expressed in Section 4.2, we believe that lower order spectral bands preserve broad domains of information. Therefore, agents with high-range sensing capabilities and less constrained, motion models would be better suited to coarse exploration as they are capable of covering larger areas quickly, with higher uncertainty. On the other hand, higher order spectral bands preserve edges and details. Lacking more general information about the map means that there will be more “false positive” areas, that is, more domains that show higher information in this scale of spatial variation but not in the original. So, agents with high-fidelity, low-range sensing capabilities seem better-suited to rely on the higher order spectral bands, because, high-fidelity might only lead to false positives that have less impact on the search. Smaller, more concentrated areas of information in this spatial scale are thus better explored by agents with curve-constrained motion models.

In the naive heuristics, agents rely on spectral bands that are well-suited to either their sensing or motion capabilities but not to both. For example, in the first naive heuristic Eq (9), agents with low-fidelity, high range sensing capabilities and slow, curve-constrained motion models rely on the lowest frequency

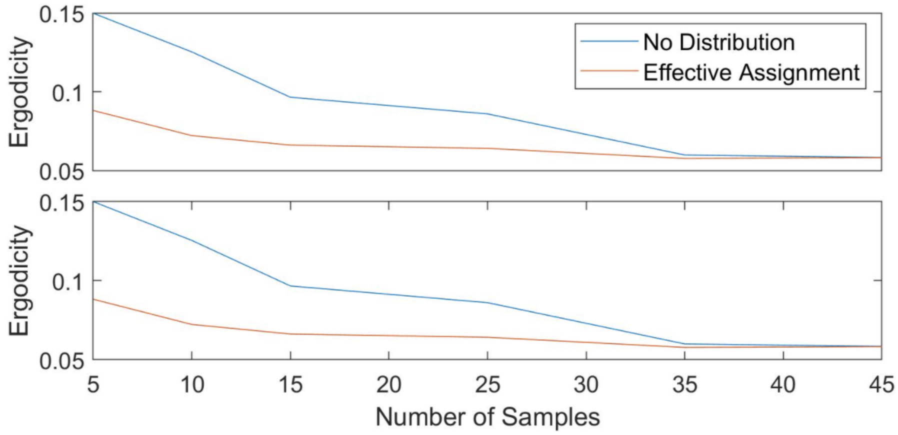

Fig. 6. Sensitivity to the number of sampled paths, comparing the coverage performance between our optimal assignment and the baseline. Gaussian maps (top), and road network (bottom). There again, and as expected, note that our distributed search approach specifically improves performance with small numbers of samples.

spectral bands, which are well suited to their sensing capabilities but not to their motion model. This kind of partially suitable mapping results in a performance almost equivalent to the baseline, as this assignment drives agents to, on average, explore the information distribution in a similar manner as in the baseline (i.e., there are no expected advantages over the baseline, but no clear downsides either).

Finally, in the adversarial heuristic, agents rely on spectral bands worst-suited to their sensing and motion capabilities. As expected, this mismatch leads to a decrease in performance as agents struggle to search at the spatial scale assigned to them, since it doesn't match their capabilities. Some side-by-side comparison videos of these methods in example scenarios can be found at http://bit.ly/DARS21-HetMASearch.

As expected, we further note that improvement in performance over the centralized approach decreases as the number of agents covering the domain increases (Fig.5). When there are fewer agents available to cover a domain, the coverage efficacy of each agent's path strongly influences the overall coverage of the domain by the team, since each agent is effectively responsible for a larger portion of the domain. However, when a large number of agents are covering a domain, each agent is effectively responsible for a smaller portion of the domain, so the effectiveness of the path of each agent has a negligible impact on the team's coverage of the domain.

Our results also indicate that improvement in performance over the centralized approach decreases as the number of samples taken in each step of the sample-based path planner increases (Fig.6). We know that the path primitives sampled by each agent at each step of the sample-based path planner depend on the spatial scale at which the agent is searching. When there are fewer samples being considered at each step, the spatial distribution that the agent is relying on has greater influence on the effectiveness of the sampled paths than when a large number of samples are drawn. That is, in the limit where sampling is performed on a near-infinite number of paths, decomposing the information map should not lead to an improved solution (since the best sampled paths will be globally optimal, and can also be found by the baseline approach). Thus, our results show that distributing the correct spectral bands (i.e., search subtasks) to the agents has more impact on the achieved coverage of a domain when planning paths using small numbers of samples. We envision this to be a significant advantage, especially for robot deployments that necessitate real-time planning and re-planning capabilities, where planning time is mainly controlled by the number of samples to be drawn.

6 Conclusion

In this paper, we investigated the idea of leveraging the spectral nature of a state-of-the-art coverage metric, the ergodic metric, to improve the heterogeneous multi-agent search of a domain by matching the agents' motion and sensing constraints to specific search subtasks. These subtasks were defined as performing search at different spatial scales, by relying on a limited subset of the spectral coefficients that represent the overall information map. After building intuition on the link between sensing and motion models and the different search scales based on subsets of coefficients, we proposed an agent assignment method to map agent types to specific search subtasks (i.e., subsets of spectral coefficients). In our systematic numerical tests, we compared our optimal assignment to naive and adversarial assignments, as well as to a baseline that plans paths for all agents regardless of their individual capabilities, and showed our distributed ergodic search approach lead to significantly improved performance (up to 40%), both in terms of coverage efficiency and time to find all targets. Additionally, our distributed approach allows sampled-based optimization methods to require a smaller number of samples to find high-quality search paths, and improves the performance of smaller agent teams, which might maximize its impact to real-world multi-robot deployments.

This work paves the way for new heterogeneous multi-agent search methods, where synergies among agents could be automatically identified and leveraged to improve the efficacy of the process. In particular, future works will approach the general problem of assigning any type of agent the right set of spectral coefficients. To this end, and for general cases where human intuition/experience cannot suffice, we believe that machine learning based methods could offer us the tool to learn such a data-driven mapping. Furthermore, the work presented in this paper considered centralized subtask assignment and path planning, but our future work will seek decentralized task assignment (and potentially ergodic path planning) solutions, to really allow such distributed heterogeneous multi-agent search methods to scale to large teams, and ultimately allow large-scale real-life deployments.

Bibliography

[1] Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. Science **345**(6198), 795–799 (2014). DOI 10.1126/science.1254295

[2] Howard, A., Parker, L.E., Sukhatme, G.S.: Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. The International Journal of Robotics Research **25**(5-6), 431–447 (2006)

[3] Ayvali, E., Salman, H., Choset, H.: Ergodic coverage in constrained environments using stochastic trajectory optimization. In: International Conference on Intelligent Robots and Systems, pp. 5204–5210. IEEE (2017)

[4] Lanillos, P., Gan, S.K., Besada-Portas, E., Pajares, G., Sukkarieh, S.: Multi-uav target search using decentralized gradient-based negotiation with expected observation. Information Sciences **282**, 92–110 (2014)

[5] Murphy, R.R.: Disaster robotics. MIT press (2014)

[6] Chand, P., Carnegie, D.A.: Mapping and exploration in a hierarchical heterogeneous multi-robot system using limited capability robots. Robotics and autonomous Systems **61**(6), 565–579 (2013)

[7] Chung, T.H., Hollinger, G.A., Isler, V.: Search and pursuit-evasion in mobile robotics. Autonomous robots **31**(4), 299–316 (2011)

[8] Mathew, G., Mezić, I.: Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. Physica D: Nonlinear Phenomena **240**(4), 432–442 (2011)

[9] Ayvali, E., Ansari, A., Wang, L., Simaan, N., Choset, H.: Utility-guided palpation for locating tissue abnormalities. Robotics and Automation Letters (2017)

[10] Miller, L.M., Silverman, Y., MacIver, M.A., Murphey, T.D.: Ergodic exploration of distributed information. IEEE Transactions on Robotics **32**(1), 36–52 (2016)

[11] Baxter, J.L., Burke, E., Garibaldi, J.M., Norman, M.: Multi-robot search and rescue: A potential field based approach. In: Autonomous robots and agents, pp. 9–16. Springer (2007)

[12] Wong, E.M., Bourgault, F., Furukawa, T.: Multi-vehicle bayesian search for multiple lost targets. In: International Conference on Robotics and Automation, pp. 3169–3174. IEEE (2005)

[13] Choset, H.: Coverage for robotics–a survey of recent results. Annals of mathematics and artificial intelligence **31**(1), 113–126 (2001)

[14] Ablavsky, V., Snorrason, M.: Optimal search for a moving target - A geometric approach. In: AIAA Guidance, Navigation, and Control Conference and Exhibit. AIAA (2000)

[15] Yan, Z., Jouandeau, N., Cherif, A.A.: A survey and analysis of multi-robot coordination. International Journal of Advanced Robotic Systems **10**(12), 399 (2013)

[16] Dias, M.B., Zlot, R., Kalra, N., Stentz, A.: Market-based multirobot coordination: A survey and analysis. Proceedings of the IEEE **94**(7), 1257–1270 (2006)
[17] Kim, J.H., Vadakkepat, P.: Multi-agent systems: a survey from the robot-soccer perspective. Intelligent Automation & Soft Computing **6**(1), 3–17 (2000)
[18] Tang, F., Parker, L.E.: ASyMTRe: Automated synthesis of multi-robot task solutions through software reconfiguration. In: Proceedings - IEEE International Conference on Robotics and Automation, pp. 1501–1508 (2005)
[19] Parker, L.E.: ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots (1994)
[20] Prorok, A., Hsieh, M.A., Kumar, V.: Formalizing the impact of diversity on performance in a heterogeneous swarm of robots. In: International Conference on Robotics and Automation, pp. 5364–5371. IEEE (2016)
[21] Dahl, T.S., Matarić, M., Sukhatme, G.S.: Multi-robot task allocation through vacancy chain scheduling. Robotics and Autonomous Systems **57**(6-7), 674–687 (2009)
[22] Jones, E.G., Browning, B., Dias, M.B., Argall, B., Veloso, M., Stentz, A.: Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In: International Conference on Robotics and Automation, pp. 570–575. IEEE (2006)
[23] Koes, M., Nourbakhsh, I., Sycara, K.: Heterogeneous multirobot coordination with spatial and temporal constraints. In: AAAI, vol. 5, pp. 1292–1297 (2005)
[24] García, P., Caamaño, P., Duro, R.J., Bellas, F.: Scalable task assignment for heterogeneous multi-robot teams. International Journal of Advanced Robotic Systems **10** (2013). DOI 10.5772/55489
[25] Gerkey, B.P., Matarić, M.J.: Sold!: Auction methods for multirobot coordination. IEEE Transactions on Robotics and Automation **18**(5), 758–768 (2002)
[26] Zlot, R., Stentz, A.T., Dias, M.B., Thayer, S.: Multi-robot exploration controlled by a market economy. In: International Conference on Robotics and Automation, vol. 3, pp. 3016–3023 (2002)
[27] Prorok, A., Hsieh, M.A., Kumar, V.: Fast redistribution of a swarm of heterogeneous robots. In: EAI International Conference on Bio-inspired Information and Communications Technologies, pp. 249–255. ICST (2016)
[28] Halász, A., Hsieh, M.A., Berman, S., Kumar, V.: Dynamic redistribution of a swarm of robots among multiple sites. In: International Conference on Intelligent Robots and Systems, pp. 2320–2325. IEEE (2007)
[29] Grabowski, R., Navarro-Serment, L.E., Paredis, C.J., Khosla, P.K.: Heterogeneous teams of modular robots for mapping and exploration. Autonomous Robots **8**(3), 293–308 (2000)
[30] Miller, L.M., Murphey, T.D.: Trajectory optimization for continuous ergodic exploration. In: American Control Conf. (ACC), 2013, pp. 4196–4201. IEEE (2013)
[31] Kobilarov, M.: Cross-entropy motion planning. The Int. J. Robot. Res. **31**(7), 855–871 (2012)