# Generalized Multi-Agent Multi-Objective Ergodic Search

Bhaskar Vundurthy, Geordan Gutow, Akshaya Kesarimangalam Srinivasan,
Albert Xu, Burhanuddin Shirose, and Howie Choset

***Abstract*— Search and rescue missions often involve coordinating multiple agents with diverse objectives in a disaster area. To address the complexities of such scenarios, this paper presents a Multi-Agent Multi-Objective Ergodic Search (MA-MO-ES) method to optimize task allocation between agents and objectives for efficient exploration. The combinatorial nature of task allocation makes it computationally expensive to solve for optimal allocation using brute force. Apart from a large number of possible allocations, computing the cost of a task allocation is itself an expensive planning problem. Unlike previous approaches that limit the solution space by restricting the number of agents per task, this work allows for flexible agent assignments. We introduce a novel branch and bound algorithm that not only accommodates this generalization but also achieves an order of magnitude speedup when attempting to minimize the maximum ergodicity across all objectives. Additionally, in scenarios where one robot is paired with multiple objectives, finding an optimal trajectory for the robot involves balancing trade-offs between conflicting objectives, resulting in the generation of a Pareto optimal front. In this work, we leverage the minimax metric to label a single point on the Pareto optimal front as optimal and present an analytical approach via the minimum bounding sphere to compute this optimal minimax point without relying on the generation of the Pareto front.**

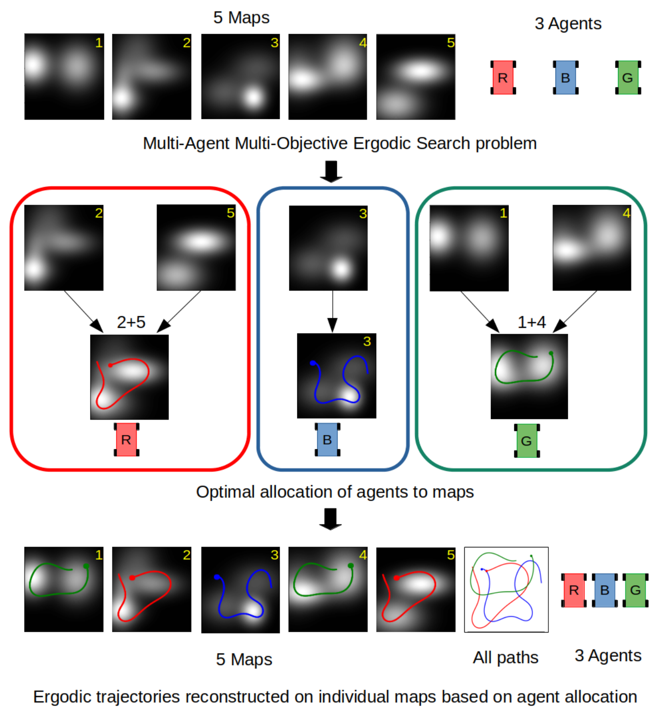


Fig. 1: This figure illustrates the Multi-agent Multi-objective Ergodic Search (MA-MO-ES) problem using five information maps $(1-5)$ that span the same physical region, and three agents (Red, Blue, and Green). The optimal allocation is shown with each agent in a different box. An agent's trajectory is optimized on a weighted average of its maps and evaluated by reconstructing the trajectory on each map.

## I. Introduction and Related Works

Applications such as search and rescue [1], surveillance [2], and planetary exploration [3] all require planning for multiple robots to collectively search a domain to gain information about it. In complex scenarios, there may be multiple competing forms of information to collect, which have different scales across the domain; which we refer to as information maps. In such a scenario, it is natural to cast the problem as a multi-objective optimization where each objective is to cover the domain subject to a single map. In a multi-agent setting, the problem further demands the appropriate allocation of agents to information maps to ensure effective coverage of all maps. We approach this problem as a Multi-Agent Multi-Objective Ergodic Search (MA-MO-ES) problem as shown in Fig. 1. The reader is referred to [4] for a detailed discussion about the prior work and a literature review on Multi-Robot Task Allocation (MRTA) problem, Ergodic Search, and Multi-Agent Multi-Objective Task Allocation and Path Planning.

The key contributions of this proposed work include:

1) Generalized Multi-Agent Multi-Objective Ergodic Search (MA-MO-ES) Problem Formulation
2) Novel Branch and Bound-based solution approach
3) An order of magnitude improvement in solution computation time when compared to a naive extension of prior work [4]
4) New solution approach leverages robots that start at the same initial location for further speedup
5) Minimum Bounding Sphere-based scalarization approach that computes the minimax point on the Pareto optimal front without computing the entire Pareto optimal front when a single robot is paired with multiple information maps

Bhaskar Vundurthy, Geordan Gutow, Akshaya Kesarimangalam Srinivasan, Albert Xu, Burhanuddin Shirose, and Howie Choset are with Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA. (email: {pvundurt, ggutow, akesarim, albertx, bshirose, choset}@andrew.cmu.edu).

## II. Problem Description

### A. Mathematical Preliminaries

Let $\mathscr{W} = [0,L_1] \times [0,L_2] \times \cdots \times [0,L_\nu] \subset \mathbb{R}^\nu$ denote a $\nu$-dimensional workspace that is to be explored by the robots. Each robot has an (identical) $n$-dimensional state space $Q = \mathscr{W} \times V$ $(n \geq \nu)$. $V$ is comprised of the robot state components, such as velocities or orientations, that do not affect what the sensor sees. Let $q_i : [0,T] \to Q$ denote a trajectory of the $i^{th}$ robot in its state space with $T \in \mathbb{R}^+$ representing the time horizon. Let the set of all state space trajectories be $H$. Let $P : Q \to \mathscr{W}$ project the state space into the workspace. The robots have deterministic dynamics given by $\dot{q}_i(t) = f(q_i(t), u_i(t))$, where $u_i(t) \in \mathbb{R}^m$ is the control input of the $i^{th}$ robot.

Let $c(x,q_i) : \mathscr{W} \times H \to [0,1]$ denote the time-averaged statistics of a trajectory $q_i$, which is defined as:

$$c(x,q_i) = \frac{1}{T}\int_0^T \delta(x - P(q_i(\tau)))d\tau, \quad (1)$$

where $\delta$ is a Dirac function. Let $\phi : \mathscr{W} \to \mathbb{R}$ denote a static information map that describes the amount of information at each location in the workspace. In this work, each information map is a probability distribution with $\int_{\mathscr{W}} d\phi = 1$ and $\phi(x) \geq 0, \forall x \in \mathscr{W}$. The quality with which a robot trajectory $q_i$ covers an information map $\phi$ can be quantified using an ergodic metric [5]:

$$\begin{aligned}\mathscr{E}(\phi,q_i) &= \sum_{k=0}^{K} \lambda_k (c_k - FC_k)^2 \quad (2)\\ &= \sum_{k=0}^{K} \lambda_k \left(\frac{1}{T}\int_0^T F_k(q(\tau))d\tau - FC_k\right)^2\end{aligned}$$

where (i) $FC_k = \int_{\mathscr{W}} \phi(x)F_k(x)dx$ represents the $k^{th}$ Fourier coefficient of the information map, $F_k(x) = \frac{1}{h_k}\Pi_{j=1}^{\nu}\cos(\frac{k_j \pi x_j}{L_j})$ is the cosine basis function for index $k \in \mathbb{N}^\nu$ and $K$ is the number of Fourier bases considered, (ii) $c_k$ denotes the $k^{th}$ Fourier coefficient of $c(x,q_i)$, (iii) $h_k$ denotes the normalization factor as defined in [5], and (iv) $\lambda_k = (1+||k||^2)^{-\frac{\nu+1}{2}}$ denotes the weight for each corresponding Fourier coefficient. When an agent must explore multiple maps, planning can instead consider a weighted sum (or "scalarization") of all the maps. Different choices of weights will result in different ergodic metric, evaluated on the individual maps.

### B. Problem Description

Given a set of $m$ **information maps** $\mathscr{M} = \{M_1, M_2, \ldots, M_m\}$, and a set of $r$ **robots** $\mathscr{R} = \{R_1, R_2, \ldots, R_r\}$, consider the power sets $\mathscr{P}(\mathscr{M})$ and $\mathscr{P}(\mathscr{R})$ of sets $\mathscr{M}$ and $\mathscr{R}$ as follows:
$\mathscr{P}(\mathscr{M}) = \{\emptyset, \{M_1\}, \{M_2\}, \ldots, \{M_1, M_2\}, \ldots, \{M_1, \ldots, M_m\}\}$
$\mathscr{P}(\mathscr{R}) = \{\emptyset, \{R_1\}, \{R_2\}, \ldots, \{R_1, R_2\}, \ldots, \{R_1, \ldots, R_r\}\}$.
A **pairing** $P_{ij}$ can then be defined as follows:
$P_{ij} = \{(m_i^{\mathscr{P}}, r_j^{\mathscr{P}}) : m_i^{\mathscr{P}} \in \mathscr{P}(\mathscr{M}), r_j^{\mathscr{P}} \in \mathscr{P}(\mathscr{R})\}$
An **allocation** $\mathscr{A}$ is a set of pairings:
$\mathscr{A} = \{P_{ij}, P_{i'j'}, \ldots\}$
Note that pairings can also be represented with a single index for brevity: $\mathscr{A} = \{P_1, P_2, \ldots, P_k\}$ where $k \leq 2^m \times 2^r$.

To make the problem tractable, we impose the following five constraints on the allocations:

1) Every pairing in an allocation should contain at least one robot and at least one map.
$$\implies |m_i^{\mathscr{P}}| \geq 1 \ \wedge \ |r_j^{\mathscr{P}}| \geq 1 \ \forall \ P_{ij} \in \mathscr{A}$$
Examples that violate constraint 1:
$\mathscr{A}_1 = \{(\{M_1\}, \emptyset), (\{M_2\}, \{R_1, R_2\})\}$,
$\mathscr{A}_2 = \{(\emptyset, \{R_1\}), (\{M_1, M_2\}, \{R_2\})\}$

2) All maps and robots belong to at least one pairing within an allocation.
$$\implies \left(\bigcup_{P_{ij}\in\mathscr{A}} m_i^{\mathscr{P}} = \mathscr{M}\right) \wedge \left(\bigcup_{P_{ij}\in\mathscr{A}} r_j^{\mathscr{P}} = \mathscr{R}\right)$$
Examples that violate constraint 2:
$\mathscr{A}_3 = \{(\{M_1\}, \{R_1, R_2\})\}$ - $M_2$ is not part of any pairing,
$\mathscr{A}_4 = \{(\{M_1, M_2\}, \{R_1\})\}$ - $R_2$ is not part of any pairing

3) One map can not be part of two different pairings:
$$\implies (m_i^{\mathscr{P}} \cap m_{i'}^{\mathscr{P}}) = \emptyset \ \forall \ i \neq i', P_{ij}, P_{i'j} \in \mathscr{A}$$
Examples that violate constraint 3:
$\mathscr{A}_5 = \{(\{M_1, M_2\}, \{R_1\}), (\{M_1\}), \{R_2\}\}$ - $M_1$ is part of both pairings

4) One robot can not be part of two different pairings:
$$\implies (r_j^{\mathscr{P}} \cap r_{j'}^{\mathscr{P}}) = \emptyset \ \forall \ j \neq j', P_{ij}, P_{ij'} \in \mathscr{A}$$
Examples that violate constraint 4:
$\mathscr{A}_6 = \{(\{M_1\}, \{R_1, R_2\}), (\{M_2\}, \{R_1\})\}$ - $R_1$ is part of both pairings

5) Multiple information maps can not be paired with multiple robots:
$$\implies |m_i^{\mathscr{P}}| = 1 \vee |r_j^{\mathscr{P}}| = 1 \ \forall \ P_{ij} \in \mathscr{A}$$
Examples that violate constraint 5:
$\mathscr{A}_7 = \{(\{M_1, M_2\}, \{R_1, R_2\})\}$

**Problem:** Determine allocation $\mathscr{A}_{opt}$ that minimizes the maximum ergodicity for all maps, while respecting constraints $1-5$:

$$\mathscr{A}_{opt} = \underset{\mathscr{A}}{\text{argmin}}\left(\max_{m_i^{\mathscr{P}} \in \mathscr{M}, P_{ij}\in\mathscr{A}} \mathscr{E}_{m_i}^{q_{r_j}}\right)$$

Before we dive into the proposed solution approach to this generalized MA-MO-ES problem, it is worth mentioning how all kinds of pairing are handled. When one robot is paired with one information map, we refer to it as Single Agent Single Objective Ergodic Search (SA-SO-ES) and utilize techniques from [6]. When one robot is paired with multiple information maps, we need to leverage techniques from multi-objective optimization. We refer to this variant as Single Agent Multi Objective Ergodic Search (SA-MO-ES) and lean on results from [7]. This work also presents an approach to extend SA-MO-ES to Multi-Agent Multi-Objective Ergodic Search (MA-MO-ES) where multiple maps are paired with multiple agents. This extension, once again, leverages multi-objective optimization. As a result, such extension would always underperform when compared to an allocation where the pairing has at most one map or at most one robot [4]. For this reason, we will not consider pairings with multiple robots and multiple maps. Lastly, when multiple robots are paired with one information map, we refer to it as Multi Agent Single Objective Ergodic Search (MA-SO-ES) and once again, use the stochastic optimization technique presented in [6] to obtain the ergodic trajectory for all the robots involved.

## III. Key Insights into the Proposed Method

It is worth noting that this formulation relaxes an additional constraint 6 placed in prior work [4]:

6) Every pairing can have at most one robot:

$$|r_j^{\mathscr{P}}| = 1, \ \forall \ P_{ij} \in \mathscr{A}$$

Counterexample: $\mathscr{A}_8 = \{(\{M_1\}, \{R_1, R_2\})\}$

The relaxation of this constraint allows for an improvement in optimality. Consider the following example (see Fig. 2) where three maps are to be allocated to three robots such that the maximum ergodicity on the three maps is minimized.

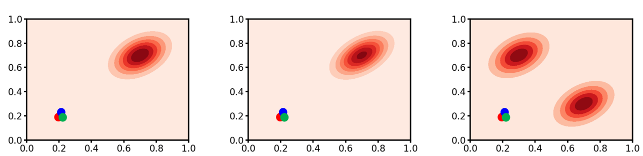

(a) Map 1 ($M_1$) (b) Map 2 ($M_2$) (c) Map 3 ($M_3$)

Fig. 2: Three information maps $M_1, M_2, M_3$ and three robots $R(red), B(blue), G(green)$ with same initial location

Prior work [4] presents a Branch and Bound approach to this problem where the robots are treated as levels and the branching is carried out on the maps (see Fig. 3). For instance, in the first level, maps $M_1$, $M_2$, and $M_3$ are allocated to robot $R$. In the next level, one of the two remaining maps are allocated to $G$. While pairing two maps with a single agent does not violate any constraints, it would exhaust two maps leaving none for the third agent, thus violating Constraint 1. Consequently, for this specific example, we do not show any pairings between multiple maps and a single robot in Fig. 3. In this example, since all agents start at the same location, any allocation that pairs one robot to one map irrespective of the robot yields the same result, as is illustrated in the leaf nodes of the graph created in Fig. 3.

The tree shown in Fig. 3 enforces Constraint 6. When attempting to relax constraint 6 with a similar structure, this framework poses a significant challenge. Adding multiple robots to a single map does not necessarily reduce ergodicity. As a result, the feature leveraged by [4] to prune nodes as they are created based on an incumbent is no longer helpful.

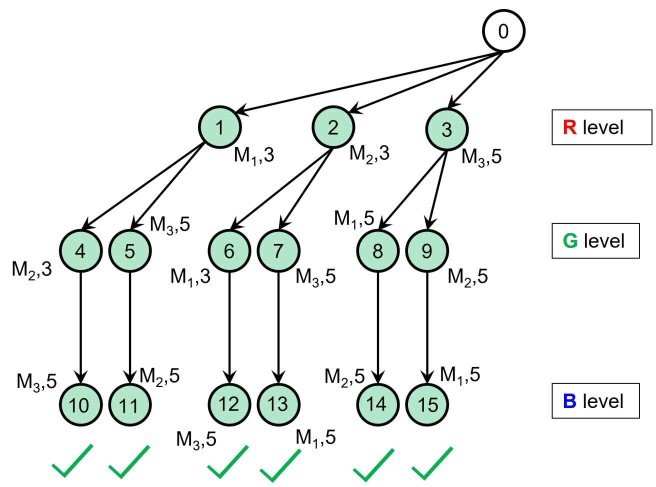


Fig. 3: Branch and bound presented in [4] for the example shown in Fig. 2. At each node, $M_1, 5$ indicates the pairing of $M_1$ with the robot in its respective layer resulting in an ergodicity value of $5 \times 10^{-6}$.

Consequently, there can be no pruning achieved until every node is expanded all the way until its leaf node.

To this end, we propose a modification to the structure of this branch and bound tree where the levels are now made of maps instead of robots. This brings back the ability to prune nodes as soon as they are created if they prove to be more expensive than the incumbent, severely improving the solution computation time. An illustration of the branch and bound tree with this modification is presented in Fig. 4.

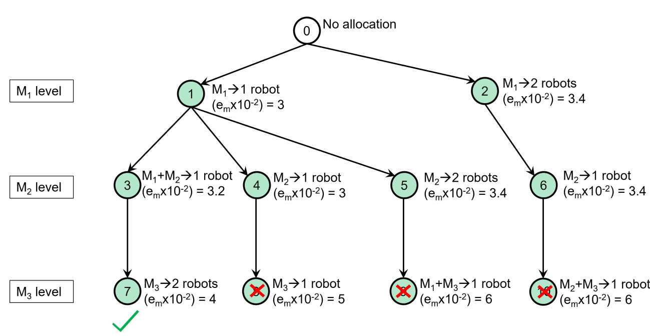


Fig. 4: Modified Branch and bound tree to solve the example shown in Fig. 2 for the generalized MA-MO-ES problem.

As is illustrated in Fig. 4, the ergodicity for the optimal allocation dropped from $5 \times 10^{-2}$ from Fig. 3 to $4 \times 10^{-2}$ with the new approach, resulting in an allocation that is optimal is impossible to obtain using the former branch and bound tree. The resulting trajectories on the three information maps are presented in Fig. 5.

Our approach also uniquely allows for scenarios where all robots start simultaneously, further increasing computational speed. Such an approach has not been discussed in prior works. Table I illustrates the specific benefits offered by the proposed approach in the last column. The first and second columns indicate prior work [4] and a naive extension of [4] to relax constraint 6, respectively.

The specific speedups obtained over [4] are presented in

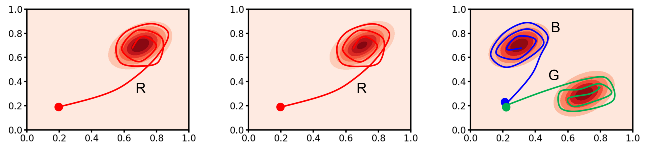


(a) $R$ for $M_1$ (b) $R$ for $M_2$ (c) $B,G$ for $M_3$

Fig. 5: Trajectories of robots $R,B,G$ on three information maps $M_1,M_2,M_3$

| | [4] | Extension to [4] | Proposed |
|---|---|---|---|
| **Optimal Solution** | No | Yes | Yes |
| **Nodes Explored** | 15 | 30 | 10 |
| **Leverages identical robots** | No | No | Yes |

TABLE I: Key benefits of the proposed framework

Table II. The results show an order of magnitude speedup when 4 robots are paired with a varied number of information maps ranging from 4 maps to 9 maps.

| **#maps** | **4** | **5** | **6** | **7** | **8** | **9** |
|---|---|---|---|---|---|---|
| **Extension to [4]** | 98 | 274 | 678 | 1923 | 6324 | 14763 |
| **Proposed** | 27 | 48 | 91 | 192 | 425 | 884 |

TABLE II: Planner run times (in sec.) for 4 robots

## IV. Minimum Bounding Sphere (MBS) Scalarization

In this work, we propose an approach for SA-MO-ES that focuses on computing the optimal scalarized information map directly. This is done without the need to calculate the entire Pareto-Optimal front, and the goal is to minimize the worst-case coverage or highest ergodic metric on any information map. The highest ergodic metric achieved on the individual maps depends on the scalarized information map used to optimize the trajectory of the agent. When considering the minmax optimality criterion, the best-scalarized information map is determined by minimizing the maximum distance between the scalarized map and the individual maps. This concept is similar to solving the minimum enclosing circle problem. Thus, we can compute the required scalarized map as the center of the minimum bounding sphere of the information maps in Fourier space.

The minimum bounding sphere for a set of points $\{p_1,p_2,\cdot,p_n\}$ is unique and is the sphere with minimum volume such that all the points lie within or on the surface of the sphere.

In an ideal scenario, a trajectory optimized towards the center of the minimum bounding sphere will precisely converge to that distribution and the ergodic metric in the trajectory optimization converges to zero. By definition of minimum bounding spheres, the maximum distance of the center of the sphere to any of the information maps is equal to the radius of the sphere (derived later in this section). As a result, the maximum ergodicity of this trajectory on any of the individual information maps will be proportional to the square of the radius of the sphere. Importantly, this result holds true regardless of the starting position of the agent being considered and is the minmax ergodic metric that can be achieved when the agent is assigned multiple information maps. Consequently, for SA-MO-ES, we compute the optimal scalarized map by determining the center of the minimum bounding sphere (MBS) of the information maps in Fourier space.

Now, let us suppose the trajectory optimized does not achieve an ergodic metric of zero against the center of the minimum bounding sphere. Even in this scenario, we can bound the maximum ergodic metric of the trajectory on the maps as stated in Theorem 1. Consider the following variables defined in Table III.

*Theorem 1:* Consider an agent tasked to cover a domain subject to information maps $\Phi = \{\phi^{(1)},\phi^{(2)},\cdots\phi^{(n_o)}\}$. Let the weighted Fourier coefficients (features) of each element of $\Phi$ be represented as $M_s = \{M_1,M_2,\cdots M_{n_o}\}$ where $M_i = feature(\phi^{(i)})$ and hence $M_i \in \mathbb{R}^{K^2}$. Let $C$ and $r$ represent the center and radius of the minimum bounding sphere of $M_s$. Consider $q_C^*$ to be the trajectory obtained by ergodic trajectory optimization against $getMap(C)$. Let the ergodic metric of $q_C^*$ on $getMap(C)$ be $E_C$. Then the ergodic metric of $q_C^*$ on any $\phi^{(i)} \in maps$ is at most $E_C + r^2 + 2r\sqrt{E_C}$.

*Proof:* [Proof of Theorem 1] The minimum bounding sphere of the feature vectors in $M_s$ is represented by the center and radius $C$ and $r$ respectively. The center $C$ is a feature vector in $\mathbb{R}^{K^2}$ and can be reconstructed to an information map, using $getMap$, that can be used to optimize for a trajectory. Let the feature vector of the average of $\Phi$ be $W$. The locally optimal ergodic trajectory of an agent when optimized against $getMap(C)$ and $getMap(W)$ is $q_C^*$ and $q_W^*$

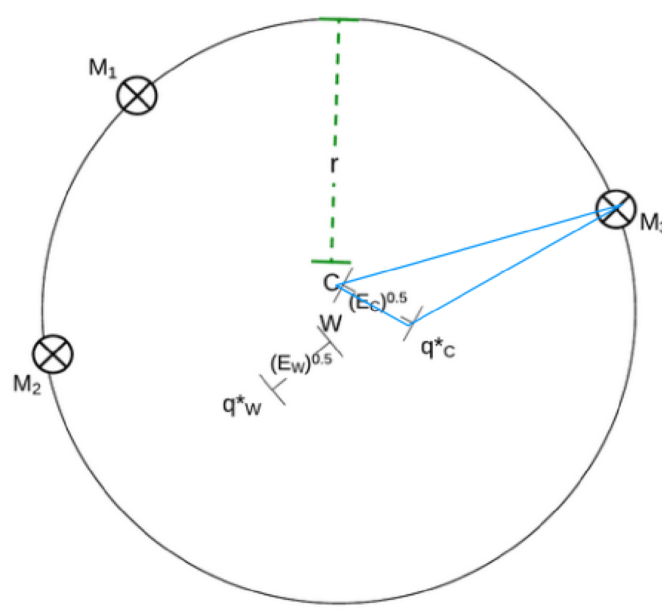


(a)

Fig. 6: **Minimum bounding sphere scalarization:** The figure represents three information maps, as feature vectors in Fourier space $\{M_1,M_2,M_3\}$, that are assigned to one agent. The trajectories $q_C^*$ and $q_W^*$ are the locally optimal ergodic trajectory of the agent when optimized against the information maps represented by $getMap(C)$ and $getMap(W)$ respectively.

TABLE III: **MBS scalarization variables definition:** Variables and functions used in the proof

| Variable | Explanation |
|---|---|
| $n_o$ | Number of assigned information maps |
| $\mathscr{M} = \{\phi^{(1)}, \phi^{(2)}, \cdots\}$ with $\phi^{(i)} : \mathscr{W} \to \mathbb{R}$ | Set of information maps or information distributions |
| $K \in \mathbb{R}$ | Number of Fourier basis functions used |
| F: 2D distribution $\to \mathbb{R}^{K^2}$ | Fourier transform: Outputs vector of Fourier coefficients |
| $\lambda \in \mathbb{R}^{K^2}$ | Weight for each corresponding Fourier coefficient |
| feature: 2D distribution $\to \mathbb{R}^{K^2}$ | Outputs a vector of weighted Fourier coefficients: feature(A) = $\lambda^{0.5} \odot F(A)$ |
| C $\in \mathbb{R}^{K^2}$ | Center of the MBS of features of assigned maps, Can be reconstructed into an information map |
| r $\in \mathbb{R}$ | Radius of the MBS of features of assigned maps |
| getMap: $R^{K^2} \to \mathscr{M}$ | Reconstructs the information map from a feature vector |
| W $\in \mathbb{R}^{K^2}$ | Feature of the average of the assigned information maps, i.e., $W = feature(\frac{1}{n_o} \sum_{i \in [1,n_o]} \phi^{(i)})$ |
| $Q = \{q, q : [0,T] \to \mathscr{W}\}$ | Set of feasible trajectories for agent considered ($T \in \mathscr{R}^+$: time horizon) |
| $\mathscr{E}$: $\mathscr{M} \times Q \to \mathbb{R}$ | Computes the ergodic metric of $q \in Q$ on $\phi \in \mathscr{M}$ |
| $q^*_\phi$: $[0,T] \to \mathscr{W}$ | The locally optimal ergodic trajectory on map $q^*_\phi$ = argmin$_{\forall q \in Q} \mathscr{E}(\phi, q)$ |
| $E_\phi \in \mathbb{R}$ | The minimum ergodic metric achieved by the agent on the map $E_\phi$ = min$_{\forall q \in Q} \mathscr{E}(\phi, q)$ |
| $\mathscr{E}_w$: $\mathscr{M}^{n_o} \times Q \to \mathbb{R}$ | Computes the highest ergodic metric of $q \in Q$ on the maps: $\mathscr{E}_w = \max_{i \in [1,n_o]} \mathscr{E}(\phi^{(i)}, q)$ |

with ergodic metric $E_C$ and $E_W$ respectively:

$$\mathscr{E}(getMap(C), q^*_C) = E_C$$
$$\mathscr{E}(getMap(W), q^*_W) = E_W$$

Next, we can derive the relation between Euclidean distance in the feature space and the ergodic metric. In the feature space, the Euclidean distance between two distributions is computed as:

$$dist(A,B) = \sqrt{\sum (feature(A) - feature(B))^2} \quad (3)$$
$$= \sqrt{\sum (\lambda^{0.5} \odot f(A) - \lambda^{0.5} \odot f(B))^2} \quad (4)$$
$$= \sqrt{\sum_{k \in K} \lambda_k (f(A)_k - f(B)_k)^2} \quad (5)$$

If we consider $A \in \mathscr{M}$ and $B \in q$, then the corresponding ergodic metric can be computed as:

$$\mathscr{E}(A,B) = \sum_{k \in K} \lambda_k (f(A)_k - f(B)_k)^2$$

Thus equation (5) becomes:

$$dist(A,B) = \sqrt{\mathscr{E}(A,b)} \quad (6)$$
$$dist^2(A,B) = \mathscr{E}(A,B) \quad (7)$$

If $\mathscr{E}(getMap(C), q^*_C) \neq 0$, then the trajectory did not converge to $C$. In that case, the center of the minimum bounding sphere $C$, the feature vector of one of the information maps $M_i$, and the agent's trajectory $q^*_C$ form a triangle as illustrated in Figure 6 for the 2D case. Then using triangle inequality we can write:

$$dist(M_i, q^*_C) \leq dist(M_i, C) + dist(C, q^*_C)$$

Since, the distances are positive, by squaring both sides we get:

$$dist^2(M_i, q^*_C) \leq dist^2(M_i, C) + dist^2(C, q^*_C) + 2dist(M_i, C)dist(C, q^*_C) \quad (8)$$

Substituting equation (7) in equation (8), we get:

$$\mathscr{E}(M_i, q^*_C) \leq dist^2(M_i, C) + E_C + 2 * dist(M_i, C) * \sqrt{E_C} \quad (9)$$

By definition of the minimum bounding sphere, the maximum distance of any feature vector from the center of the sphere is equal to the radius of the sphere. Further, since equation (9) is true for all $M_i \in M_s$, we can bound the highest ergodic metric achieved on individual maps to be:

$$\mathscr{E}_w(M_s, q^*_C) \leq r^2 + E_C + 2r\sqrt{E_C} \quad (10)$$

■

And a consequence of theorem 1 is the statement in the next corollary.

*Corollary 1.1:* The highest ergodic metric of $q^*_C$ on any map does not exceed $E_C + 2r\sqrt{E_C}$ more than that achieved by $q^*_W$ on the information maps.

Since $r^2$ is the minimum maximum ergodic metric that can be achieved when $M_s$ are assigned to a single agent (by definition of minimum bounding spheres), we also have:

$$\mathscr{E}_w(M_s, q_W^*) \geq r^2 \tag{11}$$

Thus, substituting equation (11) in equation (10), we can get an expression showing that the highest ergodic metric on the maps for the trajectory optimized against the center of the minimum bounding sphere does not exceed $E_C + 2r\sqrt{E_C}$ more than that achieved by the trajectory optimized against the average of the information maps.

$$\mathscr{E}_w(M_s, q_C^*) \leq \mathscr{E}_w(M_s, q_W^*) + E_C + 2r\sqrt{E_C}$$

This shows that even when the time average statistics of the trajectory does not converge exactly to the center of the minimum bounding sphere, the minmax metric achieved will be bounded by $E_C + r^2 + 2r\sqrt{E_C}$, where $E_C$ is the ergodic metric of the trajectory optimized against the center of the minimum bounding sphere, $r^2$ is the best minmax ergodic metric achievable and $r$ is the radius of the minimum bounding sphere.

## V. Conclusion and Future Work

The multi-agent multi-objective ergodic search problem is an allocation problem that is NP-hard to solve optimally. An exhaustive search algorithm, while optimal in allocation, quickly becomes intractable in terms of runtime as the number of agents and maps increases. We thus present a branch and bound algorithm that helps reduce the average runtime by a factor of 16 while still providing the optimal allocation scheme. Further, we present an analytical approach to compute the minimax point on the Pareto optimal front for a Single Agent Multi-Objective Ergodic (SA-MO-ES) Search pairing. This eliminates the need to compute the entire Pareto front for every pairing of SA-MO-ES kind. Future steps for this work include leveraging the minimum bounding sphere approach to create clusters of maps that can then be paired with robots as part of an allocation. We conjecture that such a clustering strategy will generate an incumbent solution that is close to the optimal, allowing for significant pruning of the search tree in our branch and bound algorithm and leading to even higher computation speeds.

## References


[1] J. Berger, N. Lo, and M. Noel, "A new multi-target, multi-agent search-and-rescue path planning approach," *International Journal of Computer and Information Engineering*, vol. 8, no. 6, pp. 978 – 987, 2014.

[2] F. M. D. Fave, S. Canu, L. Iocchi, D. Nardi, and V. A. Ziparo, "Multi-objective multi-robot surveillance," *2009 4th International Conference on Autonomous Robots and Agents*, pp. 68–73, 2000.

[3] S. G. Satpute, P. Bodin, and G. Nikolakopoulos, "Cooperative planning for multi-site asteroid visual coverage," *Advanced Robotics*, vol. 35, no. 21-22, pp. 1332–1346, 2021.

[4] A. K. Srinivasan, G. Gutow, Z. Ren, I. Abraham, B. Vundurthy, and H. Choset, "Multi-agent multi-objective ergodic search using branch and bound," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 844–849.

[5] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, pp. 432–442, 2011.

[6] E. Ayvali, H. Salman, and H. Choset, "Ergodic coverage in constrained environments using stochastic trajectory optimization," 2017.

[7] Z. Ren, A. K. Srinivasan, B. Vundurthy, I. Abraham, and H. Choset, "A pareto-optimal local optimization framework for multiobjective ergodic search," *IEEE Transactions on Robotics*, pp. 1–12, 2023.