

Balancing Robotic Search and Survival: A Game-Theoretic Framework for Ergodic Search in Contested Domains

Darwin Mick¹, Graduate Student Member, IEEE, Howie Choset¹, Fellow, IEEE, and Bhaskar Vundurthy¹, Member, IEEE

Abstract—Autonomous search in contested environments presents a fundamental conflict: agents must find targets as quickly as possible while minimizing exposure to intelligent adversaries. Existing coverage algorithms, while efficient, generate trajectories that are easily predicted and exploited by hostile actors. Conversely, standard pursuit-evasion games typically optimize for capture or escape, neglecting the agents’ need to persist in the environment to search. This paper presents a unified game-theoretic framework for multi-agent ergodic search in adversarial environments. We formulate the problem as a general-sum differential game where agents minimize a composite cost of spectral ergodic error and a penalty on adversary proximity, while adversaries minimize distance to the nearest agent. We compute local open-loop Nash equilibria strategies within a receding horizon control framework, synthesizing reactive behaviors that balance long-term coverage with immediate survival. Extensive numerical experiments reveal two key strategic insights: 1) communication topologies determine efficiency of player behaviors, and 2) in resource-constrained scenarios, adversaries achieve higher disruption by delaying capture rather than capturing immediately, fundamentally altering the optimal defensive policy.

Index Terms—Robotics, optimal control, game theory.

I. INTRODUCTION

AUTONOMOUS systems have demonstrated increasing capability in systematic search, supported by a substantial body of work on optimized path planning for target detection, frequently leveraging probabilistic representations of target location [1], [2]. However, previous work relies on the key assumption that the environment is free of intelligent mobile adversaries that actively seek to disrupt the search task [1], [2], [3]. This paper addresses the critical complication that arises when the area is patrolled by hostile *adversaries* whose

Received 1 March 2026; revised 9 May 2026; accepted 27 May 2026. Date of publication 15 June 2026; date of current version 23 June 2026. This article has not been presented at a conference. Recommended by Senior Editor P. Tesi. (Corresponding author: Darwin Mick.)

Darwin Mick is with the Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: dmick@andrew.cmu.edu).

Howie Choset and Bhaskar Vundurthy are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: choset@andrew.cmu.edu; pvundurt@andrew.cmu.edu).

Digital Object Identifier 10.1109/LCSYS.2026.3703796

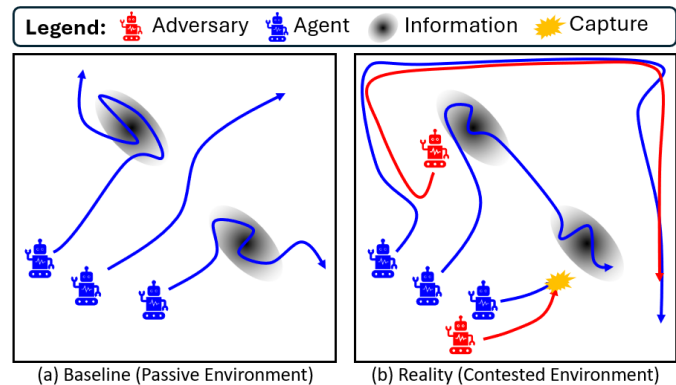


Fig. 1. Strategies for search in contested environments. (a) In a benign environment, three agents evenly distribute the search task. (b) The differential game formulation yields an emergent division of labor in the presence of adversaries: one agent is pursued (limiting its search), one is captured (stopping its search), the remaining agent naturally attempts to maximize search, demonstrating implicit cooperation without explicit role assignment.

objective is to impede the search of *agents* (Fig. 1). This introduces a fundamental tension for agents between fulfilling the complex spatial search objective while simultaneously evading adversaries.

To resolve this fundamental tension, we propose a unified control framework that couples ergodic search [4] with differential game theory [5]. Ergodic search is a powerful information-based search technique that dynamically prioritizes where to send agents based on information encoded by probability distributions [4]. Differential game theory provides a “survival” mechanism that models the dynamic interaction among agents and adversaries to enable agents to deviate from optimal search trajectories for survival [5].

Coupling the objectives of ergodic search with the survival constraints of a differential game creates a unique challenge not addressed in standard pursuit-evasion literature [6]. In classical evasion scenarios, agents minimize risk by maximizing distance from threats while sometimes fleeing to safe regions [7]. In our problem, the search mission requires agents to revisit high-information regions without the possibility of a permanent safe haven. This constraint necessitates a unique cost function that continuously balances the global pressure to search against the local imperative to survive.

Additionally, the control framework must be robust to the varying degrees of communication inherent in contested environments — ranging from complete signal loss to active coordination. Because information availability fundamentally alters optimal strategies [8], we systematically analyze how team performance evolves across a spectrum of communication topologies rather than assuming a static configuration.

Finally, we investigate how these operational constraints influence the adversaries' optimal strategy. As such, we analyze the critical tradeoff between a game of attrition (capture) and a game of area denial (pursuit). Our results reveal that under resource constraints, a simple "pursuit" strategy is often superior to a "capture" strategy, as capturing an agent removes the pursuer from the field, whereas pursuit limits the remaining agents access to high-value information.

Three primary contributions formalize these insights:

- 1) A differential game formulation that captures the interaction between searching agents and adversaries, enabling agents to balance search and survival objectives.
- 2) Optimal control strategies and resultant game outcomes under varying communication topologies.
- 3) A characterization of adversary's best response as a function of relative team size.

II. RELATED WORK

We first review search strategies and strategic adversarial modeling via differential games, then identify the research gap addressed. Note, we use the term "search" to describe agents gathering spatial information from the environment, rather than the traditional pursuit-evasion context where pursuers search for evading opponents.

A. Search and Coverage Algorithms

Classical Coverage Path Planning (CPP) [1], [2] guarantees complete inspection but uses deterministic patterns (e.g., raster scan) that render agents predictable and vulnerable to intelligent adversaries. Informative Path Planning (IPP) [9], [10] provides probabilistic prioritization but often behaves myopically. Ergodic search [4] offers a robust middle ground: combining CPP's global guarantees with IPP's probabilistic weighting while generating unpredictable trajectories. Although prior work extends ergodic search to static obstacles [11] and complex geometric constraints [12], scenarios involving intelligent mobile adversaries remain unaddressed.

B. Differential Games

Differential games provide the standard framework for modeling adversarial interactions, particularly for pursuit-evasion scenarios [6], target defense and rendezvous [7], and reach-avoid problems [13], but general solution methods are often intractable for real-time applications. On the other hand, Linear-Quadratic Differential Game (LQDG) formulations assume linear system dynamics and quadratic costs [8] yielding closed-form feedback solutions making them ideal for real-time multi-agent planning [7], [14].

Differential Games are commonly modeled as either zero-sum games where one player's gain is offset by its opponent's loss or general sum games where all players can benefit or lose. These formulations typically optimize for point-based objectives (reaching a state or defending a boundary) [15] and do not address objectives where the agent's goal is defined by

a statistic over its entire trajectory (e.g. searching a probability map). To the best of our knowledge, no prior work integrates the spatiotemporal objectives of ergodic search directly into the cost function of a differential game.

III. PROBLEM FORMULATION

To capture the complex interaction between agents and adversaries, we define the system dynamics and ergodic metric. We then analyze the metric's sensitivity to search interruption, motivating the asymmetric objectives of the competing teams.

A. System Modeling

A team of agents and a team of adversaries operate in a bounded two-dimensional workspace, $\mathcal{W} = [0, L_1] \times [0, L_2] \subset \mathbb{R}^2$. The set of agents is denoted by $\mathcal{A} = \{a_1, \dots, a_A\}$ and the set of adversaries as $\mathcal{G} = \{g_1, \dots, g_G\}$. Assuming homogeneous dynamics across both teams, we define the joint set of players as $\mathcal{P} = \mathcal{A} \cup \mathcal{G}$. This allows us to unify the system notation: for any player $p_i \in \mathcal{P}$, the complete trajectory is $x_i : [0, T] \rightarrow \mathcal{X}$ where \mathcal{X} is an m -dimensional state space. T denotes the total duration of the game and the state of player with index i at time t is $x_i(t) \in \mathcal{X}$. The control input of player with index i is defined as $u_i : [0, T] \rightarrow \mathcal{U}$, where $u_i(t)$ represents the input at time t within the control space \mathcal{U} . All players follow the same dynamics $\dot{x}_i(t) = f(x_i(t), u_i(t))$.

We define a capture event at time $t \leq T$ if the Euclidean distance between an agent a and an adversary g satisfies $\min_{g \in \mathcal{G}} \|x_a(t) - x_g(t)\|_2 \leq \delta$, where $\delta > 0$ is the capture radius.

In this work, we explore game dynamics under two distinct interaction models: one where a successful capture removes both the agent and adversary from the game, and a "no capture" model where adversaries continuously track agents to restrict their movement without initiating capture. We evaluate the strategic differences between these two mechanics in Sec. V-C.

B. Ergodic Measure & Sensitivity

Let $\phi : \mathcal{W} \rightarrow [0, 1]$ represent a static information map that describes the value of information at any position in the workspace. That is, $\int_{\mathcal{W}} d\phi = 1$ and $\phi(q) \geq 0, \forall q \in \mathcal{W}$. To evaluate multi-agent coverage, let a team trajectory \mathbf{x}_A be the concatenated trajectories of all agents $a \in \mathcal{A}$. For an information map ϕ and team trajectory \mathbf{x}_A , we define the ergodic metric as in [16]:

$$\mathcal{E}(\phi, \mathbf{x}_A) = \sum_{k=0}^K \lambda_k \left(\frac{1}{T} \int_0^T F_k(\mathbf{x}_A(t)) dt - FC_k \right)^2 \quad (1)$$

where $FC_k = \int_{\mathcal{W}} \phi(q) F_k(q) dq$ is the k th Fourier coefficient of the information map, $F_k(q) = \frac{1}{h_k} \prod_{n=1}^v \cos\left(\frac{k_n \pi q_n}{L_n}\right)$ is the cosine basis function for index $k \in \mathbb{N}$. h_k is the normalization factor as defined in [4]. $\lambda_k = (1 + \|k\|^2)^{-\frac{v+2}{2}}$ denotes the weight for each corresponding Fourier coefficient.

Consequently, achieving low ergodicity requires both visiting high-information regions and persisting in them for a duration proportional to their probability density suggesting premature termination of the trajectory—such as capture by an adversary—will penalize search performance. To quantify this cost, in Fig. 2 we analyzed the difference in ergodicity

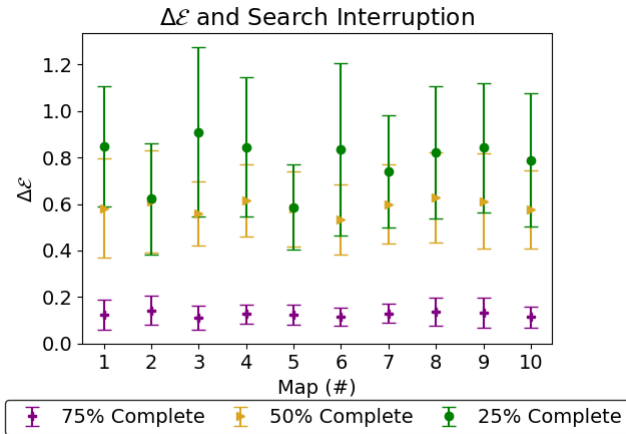


Fig. 2. The average change in ergodicity ($\Delta\mathcal{E} = \mathcal{E}_{\text{interrupted}} - \mathcal{E}_{\text{complete}}$) over 10 maps and 10 initial positions where errorbars represent one standard deviation. Results demonstrate that interruption degrades the quality of the information-based search being conducted by the agents.

of interrupted and complete trajectories ($\Delta\mathcal{E} = \mathcal{E}_{\text{interrupted}} - \mathcal{E}_{\text{complete}}$). Non-negative $\Delta\mathcal{E}$ values in all cases imply that early termination degrades search quality establishing a critical design constraint: agents cannot treat evasion as a secondary objective; they must actively maximize separation distance from adversaries to ensure time for search.

Crucially, the adversary does not share this trade-off between search and evasion. This asymmetry in objectives — where agents balance search against survival while adversaries optimize solely for proximity — necessitates a general-sum formulation rather than a standard zero-sum framework.

Reflecting the asymmetry, we define the game objectives:

Problem 1: Given the multi-agent system constraints in Section III-A, determine the optimal control sequences u_i^* for all $p_i \in \mathcal{P}$ that satisfy the dynamics constraints and establish a local open-loop Nash equilibrium where:

- 1) The agent team \mathcal{A} minimizes a composite objective of spectral ergodic error and adversary proximity.
- 2) Each adversary $g \in \mathcal{G}$ acts as a pursuer, minimizing the separation distance to its assigned target agent.

IV. METHODOLOGY

A. Player Objectives & Cost Functions

The agents have three primary, competing objectives: ergodic search, capture evasion, and control effort. We formulate this trade-off as a constrained trajectory optimization problem minimizing a linear combination of the objectives:

$$\begin{aligned} & \underset{\mathbf{u}_{\mathcal{A}}}{\operatorname{argmin}} \alpha \mathcal{E}(\phi, \mathbf{x}_{\mathcal{A}}) - \beta \int_0^T \sum_{a=1}^A \|x_a(t) - x_{g,\text{evade}}(t)\|_2^2 dt \\ & + \int_0^{T-1} \mathbf{u}_{\mathcal{A}}(t)^\top \mathbf{R} \mathbf{u}_{\mathcal{A}}(t) dt \\ \text{s.t. } & \dot{x}_a = f(x_a(t), u_a(t)), \mathbf{x}_{\mathcal{A}}(0) = \bar{\mathbf{x}}_{\mathcal{A}}(0) \\ & x_a(t) \in \mathcal{X}, u_a(t) \in \mathcal{U}, c(\mathbf{x}_{\mathcal{A}}) \in \mathcal{W}, \forall a \in \mathcal{A} \end{aligned} \quad (2)$$

where α and β are scaling weights for the ergodicity and adversary proximity objectives, respectively, and \mathbf{R} is a positive definite diagonal matrix that penalizes the control input. The term $x_{g,\text{evade}}(t)$ represents the position of the nearest

adversary g , determined by $\operatorname{argmin}_{g \in \mathcal{G}} \|x_a(t) - x_g(t)\|_2$. Furthermore, $\mathbf{u}_{\mathcal{A}}$ is the concatenated controls of all agents $a \in \mathcal{A}$. $c(\cdot)$ is an inequality constraint bounding the agents within the workspace and $\bar{\mathbf{x}}_{\mathcal{A}}(0)$ defines initial positions.

The primary goal of an adversary is to pursue and capture an agent. Because adversaries do not share a global spatial objective (unlike the agents' joint ergodic metric), each adversary $g \in \mathcal{G}$ independently minimizes its own decentralized cost function:

$$\begin{aligned} & \underset{u_g}{\operatorname{argmin}} \eta \int_0^T \|x_g(t) - x_{a,\text{pursue}}(t)\|_2^2 dt \\ & + \int_0^{T-1} u_g(t)^\top \mathbf{R} u_g(t) dt \\ \text{s.t. } & \dot{x}_g = f(x_g(t), u_g(t)), x_g(0) = \bar{x}_g(0) \\ & x_g(t) \in \mathcal{X}, u_g(t) \in \mathcal{U}, c(x_g) \in \mathcal{W} \end{aligned} \quad (3)$$

where η is a positive scalar that penalizes distance from the pursued agent, and $x_{a,\text{pursue}}(t)$ denotes the position of that targeted agent a . By default, adversaries greedily target the nearest agent, determined via $\operatorname{argmin}_{a \in \mathcal{A}} \|x_a(t) - x_g(t)\|_2$. While we introduce mechanisms for adversaries to explicitly coordinate these discrete target assignments in Sec. V-B, their continuous trajectory optimization remains completely decentralized.

B. Local Nash Equilibrium & Convergence

Having defined the individual player objectives, we now formalize the solution concept for the differential game defined via Problem 1. Classic solutions to differential games typically rely on local Linear-Quadratic Differential Game (LQDG) approximations to find closed-form Riccati feedback solutions [14]. More recent advanced solvers utilize inference-based strategy alignment for general-sum games [17]. However, whether utilizing local LQDG or inference-based alignment, these methods assume the cost function is in Bolza form. Recall, a Bolza form cost functional consists of a terminal cost and an integral stage cost evaluated over a fixed time horizon. Because our formulation involves an inherently non-Bolza ergodic objective (1) that evaluates the spatial distribution of the entire trajectory holistically, these stage-additive approaches are structurally incompatible. This complexity precludes the existence of analytical recursive feedback laws, necessitating the use of iterative first-order numerical optimization to compute local open-loop equilibrium strategies.

1) *Augmented Lagrangian Formulation & Simultaneous Gradient Descent:* To address the non-Bolza nature of the cost landscape, we formulate the *Augmented Lagrangian* \mathcal{L} as in [18] using the slack variable method [19]:

$$\begin{aligned} \mathcal{L}_i(x_i, u_i, \lambda_i, \mu_i) &= J_i(x_i, u_i) \\ &+ \sum_w \left[\lambda_w c_w(x_i, u_i) + \|c_w(x_i, u_i)\|_2^2 \right] \\ &+ \sum_n \left[\frac{1}{2} \max(0, \mu_n + g_n(x_i, u_i))^2 - \mu_n \right] \end{aligned} \quad (4)$$

where $J_i(\cdot, \cdot)$ is the unconstrained cost for player p_i from (2, 3), $c_w(\cdot, \cdot)$ is the w^{th} equality constraint, $g_n(\cdot, \cdot)$ is the n^{th} inequality constraint. λ_w and μ_n are the lagrange multipliers. To compute open-loop strategies for all players, we employ a simultaneous

gradient descent approach. In each iteration k , both the agent team (\mathcal{A}) and the adversary team (\mathcal{G}) update their set of control sequences $(u_{\mathcal{A}}, u_{\mathcal{G}})$ simultaneously until convergence (i.e. a limit point is reached). This is accomplished by taking steps of size γ in the direction of the negative gradient of their respective Augmented Lagrangian, \mathcal{L}_i , where $p_i \in \mathcal{P}$, until a maximum iteration count k_{max} is reached or the change in cost $\Delta\mathcal{L}_i = \mathcal{L}_i^{(k+1)} - \mathcal{L}_i^{(k)}$ falls below a predefined threshold:

$$u_i^{(k+1)} = u_i^{(k)} - \gamma \nabla_{u_i} \mathcal{L}_i(u_{\mathcal{A}}^{(k)}, u_{\mathcal{G}}^{(k)}) \quad (5)$$

where $\nabla_{u_i} \mathcal{L}$ is the gradient of \mathcal{L} with respect to u_i . Lagrange multiplier updates are omitted for brevity (see [19]). The entire agent team's Augmented Lagrangian, $\mathcal{L}_{\mathcal{A}}$, is defined centrally to enable joint minimization of the ergodic metric (1).

2) Local Nash Equilibrium Characterization: We validate the quality of the control strategies computed via simultaneous gradient descent using (5) by characterizing the game-theoretic properties of the limit points satisfy the conditions for a Local Nash Equilibrium [20].

Definition 1: A control strategy profile $u^* = (u_{\mathcal{A}}^*, u_{\mathcal{G}}^*)$ is a Local Open-Loop Nash Equilibrium if the cost functionals for the agent team \mathcal{A} and adversary team \mathcal{G} satisfy the following conditions, holding the opposing team's strategies fixed:

1) **Stationarity:**

$$\nabla_{u_{\mathcal{A}}} \mathcal{L}_{\mathcal{A}}(u_{\mathcal{A}}^*, u_{\mathcal{G}}^*) = 0, \quad \nabla_{u_{\mathcal{G}}} \mathcal{L}_{\mathcal{G}}(u_{\mathcal{G}}^*, u_{\mathcal{A}}^*) = 0$$

2) **Local Optimality:**

$$\nabla_{u_{\mathcal{A}}}^2 \mathcal{L}_{\mathcal{A}}(u_{\mathcal{A}}^*, u_{\mathcal{G}}^*) \geq 0, \quad \nabla_{u_{\mathcal{G}}}^2 \mathcal{L}_{\mathcal{G}}(u_{\mathcal{G}}^*, u_{\mathcal{A}}^*) \geq 0 \quad \blacksquare$$

To guarantee our solver avoids unstable stationary points, we invoke the following result from non-convex optimization theory [21] then formally characterize the equilibrium properties of the computed strategies.

Lemma 1: (Gradient Descent Avoids Strict Saddles [21]): Let f be a twice continuously differentiable (C^2) function with an L -Lipschitz continuous gradient. Let q^* be a strict saddle point (a stationary point where $\lambda_{\min}(\nabla^2 f(q^*)) < 0$). If the step size satisfies $0 < \gamma < 1/L$, then with random initialization, the probability of converging to q^* is zero: $\mathbb{P}(\lim_{k \rightarrow \infty} q_k = q^*) = 0$.

Remark 1: The applicability of Lemma 1 relies on the cost function being C^2 with a Lipschitz continuous gradient. Our formulation satisfies these conditions as the cost functionals are composed of a finite sum of smooth Fourier basis functions and quadratic penalty terms. Since these components are C^2 and defined on a compact workspace \mathcal{W} , the Hessian is bounded by some finite constant (i.e., $\|\nabla^2 \mathcal{L}\| \leq L < \infty$), satisfying the Lipschitz requirement. We select a step size γ sufficiently small to satisfy the stability bound in Lemma 1.

Remark 2: While Lemma 1 assumes standard gradient descent, it extends to the simultaneous gradient dynamics in (5) because $\nabla_{u_i} \mathcal{L}_i$ is independent of other players. This decouples the simultaneous update into independent steps, preserving the lemma's validity.

Proposition 1: If the simultaneous gradient descent algorithm (5) with random initialization converges to a strategy profile u^* , then u^* constitutes a Local Open-Loop Nash Equilibrium for the differential game defined by (2) and (3).

Proof: Convergence of simultaneous gradient descent implies that the update term in (5) vanishes, satisfying the Stationarity condition of Definition 1. To establish Local Optimality, we first note (via Remark 1) that the cost functionals in (2) and (3) satisfy the Lipschitz gradient continuity condition of Lemma 1. Furthermore, Lemma 1 guarantees that, by

initializing the control trajectories with a continuous random perturbation, gradient descent will not converge to a strict saddle point almost surely. This implies the Hessians at u^* cannot have negative eigenvalues, satisfying the positive semi-definite condition. Satisfying both conditions, u^* is a Local Open-Loop Nash Equilibrium. \blacksquare

Remark 3: While Proposition 1 characterizes limit points conditionally, empirical results demonstrate robust convergence. In $N = 200$ Monte Carlo trials for a 1v1 game, setting $k_{max} = 10000$, $\Delta\mathcal{L}_i \leq 1 \times 10^{-7}$, and $\gamma = 0.1$ achieved convergence to stationary points in 100% of cases.

C. Receding Horizon

Proposition 1 yields an open-loop strategy profile u^* for a fixed time horizon, which is brittle in environments with evolving adversary tactics. As such, we embed the open-loop solver in a Receding Horizon Control (RHC) framework (Alg. 1) approximating a feedback Nash equilibrium [15], providing the reactivity of closed-loop control.

Algorithm 1 Receding Horizon Control Framework

Data: $T, T_H, \delta, \mathbf{R}, \alpha, \beta, \eta, k_{max}, \gamma, \Delta\mathcal{L}_i$

Input: Initial State $x_i(0) \forall p_i \in \mathcal{P}$

```

1 for  $t = 0, \dots, T$  do
2    $k = 0$ ;
3   Assign  $\{x_{g, evade}, x_{a, pursue}\} \forall p_i \in \mathcal{P}$ ;
4   Initialize  $u_{\mathcal{A}}, u_{\mathcal{G}}$  and Lagrangians (4);
5   while  $k \leq k_{max}$  and not converged do
6      $u_{\mathcal{A}}^{(k+1)} = u_{\mathcal{A}}^{(k)} - \gamma \nabla_{u_{\mathcal{A}}} \mathcal{L}_{\mathcal{A}}(u_{\mathcal{A}}^{(k)}, u_{\mathcal{G}}^{(k)})$ ;
7      $u_{\mathcal{G}}^{(k+1)} = u_{\mathcal{G}}^{(k)} - \gamma \nabla_{u_{\mathcal{G}}} \mathcal{L}_{\mathcal{G}}(u_{\mathcal{A}}^{(k)}, u_{\mathcal{G}}^{(k)}) \forall g \in \mathcal{G}$ ;
8      $k = k + 1$ ;
9   end
10  Execute  $u_i^{(k)} \forall p_i \in \mathcal{P}$ ;
11 end
```

V. EXPERIMENTAL RESULTS AND DISCUSSION

We now validate our framework via numerical simulations, analyzing the impact of communication on coordination and the trade-off between capture and no-capture strategies.

A. Experimental Setup & Metrics

We evaluate our framework in scenarios with three agents, $A = 3$, and two adversaries, $G = 2$. We conducted 100 randomized trials across 10 distinct information maps, initializing players such that the adversaries begin "behind" the agents. Players follow double integrator dynamics within a bounded workspace $\mathcal{W} = [0, 100]^2$. Model parameters are set as follows: horizon $T = 500$, receding horizon $T_H = 100$, capture radius $\delta = 1.95$, pursuit objective weight $\eta = 10.0$, and control penalty $\mathbf{R} = 0.001 \times \mathbf{I}$ where \mathbf{I} is an identity matrix of the same dimension as \mathcal{U} . The solver parameters are $k_{max} = 1000$, $\Delta\mathcal{L}_i \leq 1 \times 10^{-7}$, and $\gamma = 0.1$. To quantify performance, we define two primary metrics:

- **Spectral Ergodic Error (\mathcal{E}):** Measures search completeness; lower values imply effective coverage
- **Time-To-Capture (T_c):** Measures adversary pressure; lower values imply effective pursuit




Fig. 3. Qualitative comparison of emergent team behaviors across communication topologies. Circles are the starting positions and triangles are ending positions. (Left) *Passive Sensing* causes naive reaction to adversaries and inefficient boundary hugging. (Middle) *Agent Coordination* allows pursued agents to evade while others search, though adversaries pursue redundantly and overlap. (Right) *Perfect Information* eliminates adversary overlap.

To evaluate the search-survival trade-off, we varied weights α and β in a 1v1 scenario ($A = G = 1$). Prioritizing search ($\alpha = 42.25$, $\beta = 2.75$) yields lower ergodicity ($\mathcal{E} = 0.03$) but rapid capture ($T_c = 0.33$). Conversely, prioritizing evasion ($\alpha = 20.25$, $\beta = 4.75$) ensures indefinite survival ($T_c = 1.0$) at the expense of search ($\mathcal{E} = 0.26$). We select $\alpha = 20.25$ and $\beta = 2.75$ for subsequent trials to balance survival ($T_c = 0.96$) with efficient search ($\mathcal{E} = 0.09$).

Our framework avoids combinatorial explosion because the optimization is structurally decoupled (see Remark 2). While our simulation utilizes a centralized implementation for evaluating game dynamics — resulting in a 9.39 s step time when $A = 3$ — this ergodic minimization bottleneck can be resolved in physical deployment via decentralized Fourier coefficient sharing [22]. Runtimes can be further managed by tuning planning hyperparameters (T_h , k_{max}) to balance algorithmic efficacy with real-time requirements.

B. Impact of Team Coordination

In contested environments, robust operation heavily relies on effective team coordination, a capability that fundamentally alters optimal player strategies [8]. To capture this, we evaluate three communication topologies that dictate coordination by restricting how agents and adversaries select their target opponents in equations (2) and (3), implemented at each receding horizon step (Algorithm 1, Line 3):

- 1) **Passive Sensing (Baseline):** Agents evade the nearest adversary ($x_{g, evade}$), and adversaries pursue the nearest agent ($x_{a, pursue}$), with no information exchange.
- 2) **Agent Coordination (Agent Comm.):** Agents pool information to identify which specific agents are being actively pursued. Pursued agents evade, while unthreatened agents focus exclusively on search ($\beta = 0$).
- 3) **Perfect Information (Full Comm.):** Adversaries also coordinate their target assignments ($x_{a, pursue}$) to ensure no two adversaries pursue the same agent.

Fig. 3 illustrates the qualitative impact of these topologies. In passive sensing (Fig. 3a), lacking intention information, multiple agents react simultaneously to the same adversary and flee to the boundary of the workspace. In contrast, agent coordination (Fig. 3b) enables emergent role allocation:




Fig. 4. Impact of communication topologies on search performance (Ergodic Error) compared to baseline. Agent communication yields significant improvement ($\Delta\mathcal{E} < 0$), but shows slightly degraded search performance when adversaries can communicate as well.

threatened agents evade, unburdened agents exploit the free space to search, and adversaries redundantly pursue the same agent. Perfect information (Fig. 3c) avoids this redundancy by efficiently assigning an agent to at most one adversary.

We compare the performance differential of each topology with the baseline, reporting $\Delta\mathcal{E}$ (Fig. 4) and ΔT_c (Fig. 5), where negative values indicate improvement compared to baseline. Agent coordination significantly improves ergodicity, while perfect information slightly degrades search compared to agent coordination. Fig. 5 reveals that additional coordination does not impact ΔT_c implying that improvements in ergodicity are not solely driven by extended survival times, highlighting the search efficiency gains derived from coordinated role allocation and challenging the correlation suggested by sensitivity analysis (Fig. 2).

C. Attrition vs. Area Denial

To investigate whether capture is truly the optimal disruptive strategy, we analyze the fundamental trade-off between Attrition (capture) and Area Denial (no capture). Defensive game theory suggests that optimal strategies shift based on

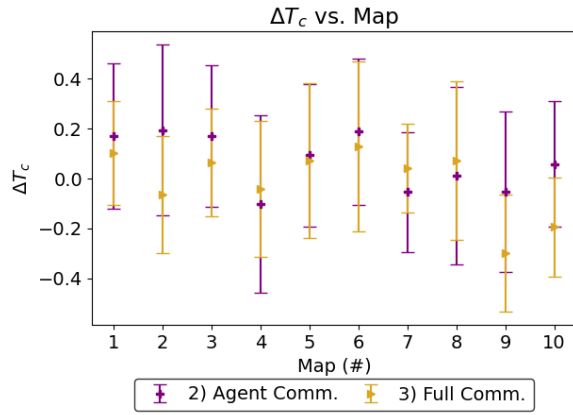


Fig. 5. Impact of coordination on survival (Time-to-Capture) compared to baseline. Varying communication topology does not significantly affect time-to-capture. Thus, performance gains in Fig. 4 stem from maximizing safe agents' utility rather than extending the survival of targeted ones.




Fig. 6. Strategic trade-off between Attrition (Capture) and Area Denial (No-Capture). Ergodic error difference ($\Delta\mathcal{E} = \mathcal{E}_{attrition} - \mathcal{E}_{denial}$) reveals that in resource-constrained scenarios ($A > G$), Area Denial is a superior ($\Delta\mathcal{E} < 0$), as a single adversary can continuously disrupt multiple searching agents more effectively than capturing one.

relative team size [3]. We examine this relationship by simulating scenarios where the adversary operates at a numerical deficit, parity, or advantage ($A = 3$ and $G = \{2, 3, 4\}$) using the perfect information topology to isolate strategic intent. Crucially, disabling capture reveals that continuous threats naturally displace agents from information peaks. This area denial emerges naturally from equations (2) and (3) without cost function modifications.

Fig. 6 presents the difference in ergodic performance $\Delta\mathcal{E}$ between capture and no-capture scenarios. We observe that in resource-constrained scenarios where the adversary is outnumbered ($A > G$), emergent Area Denial is significantly more effective at degrading search performance than Attrition. Under the Attrition mechanic, a successful capture removes both the agent and the guarding adversary from the analysis, inadvertently freeing the remaining agents to search uncontested. In contrast, under the no-capture mechanic, a single adversary can persistently disrupt multiple agents and prevent them from settling on high-value information peaks.

VI. CONCLUSION

This work establishes the initial unified game-theoretic framework for ergodic search in contested environments,

systematically characterizing how communication topologies and team size dictate optimal strategies. Future work will extend this framework to decentralized solvers and incorporate belief-space planning to account for sensor uncertainty.

REFERENCES

- [1] H. Choset, "Coverage for robotics—A survey of recent results," *Ann. Math. Artif. Intell.*, vol. 31, nos. 1–4, pp. 113–126, Oct. 2001.
- [2] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robot. Auto. Syst.*, vol. 61, no. 12, pp. 1258–1276, Dec. 2013.
- [3] T. H. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics: A survey," *Auto. Robots*, vol. 31, no. 4, pp. 299–316, Nov. 2011.
- [4] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Phys. D, Nonlinear Phenomena*, vol. 240, nos. 4–5, pp. 432–442, 2010.
- [5] R. Isaacs, *Differential Games: A Mathematical Theory With Applications to Warfare and Pursuit, Control and Optimization*. Hoboken, NJ, USA: Wiley, 1965.
- [6] D. Maity, "Optimal intermittent sensing for Pursuit-Evasion games," *IEEE Control Syst. Lett.*, vol. 7, pp. 3000–3005, 2023.
- [7] S. K. Singh, P. V. Reddy, and B. Vundurthy, "Study of multiple target defense differential games using receding horizon-based switching strategies," *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 4, pp. 1403–1419, Jul. 2022.
- [8] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. Philadelphia, PA, USA: SIAM, 1998.
- [9] S. McCammon and G. A. Hollinger, "Topological path planning for autonomous information gathering," *Auto. Robots*, vol. 45, no. 6, pp. 821–842, Sep. 2021.
- [10] P. Tokekar, J. V. Hook, D. Mulla, and V. Isler, "Sensor planning for a symbiotic UAV and UGV system for precision agriculture," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1498–1511, Dec. 2016.
- [11] B. Shirose, A. M. Johnson, B. Vundurthy, H. Choset, and M. Travers, "GESCE: Graph-based ergodic search in cluttered environments," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Oct. 2024, pp. 7611–7616.
- [12] A. Xu, B. Vundurthy, G. Gutow, I. Abraham, J. Schneider, and H. Choset, "Measure preserving flows for ergodic search in convoluted environments," 2024, *arXiv:2409.09164*.
- [13] M. Chen, Z. Zhou, and C. J. Tomlin, "Multiplayer reach-avoid games via pairwise outcomes," *IEEE Trans. Autom. Control*, vol. 62, no. 3, pp. 1451–1457, Mar. 2017.
- [14] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, "Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 1475–1481.
- [15] M. Wang, Z. Wang, J. Talbot, J. C. Gerdes, and M. Schwager, "Game-theoretic planning for self-driving cars in multivehicle competitive scenarios," *IEEE Trans. Robot.*, vol. 37, no. 4, pp. 1313–1325, Aug. 2021.
- [16] 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 *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2023, pp. 844–849.
- [17] L. Peters, D. Fridovich-Keil, C. J. Tomlin, and Z. N. Sunberg, "Inference-based strategy alignment for general-sum differential games," in *Proc. Int. Joint Conf. Auto. Agents Multiagent Syst.*, May 2020, pp. 1037–1045.
- [18] C. Lerch, D. Dong, and I. Abraham, "Safety-critical ergodic exploration in cluttered environments via control barrier functions," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Jun. 2023, pp. 10205–10211.
- [19] J. Nocedal and S. J. Wright, *Numerical Optimization*. Cham, Switzerland: Springer, 2006.
- [20] L. J. Ratliff, S. A. Burden, and S. S. Sastry, "On the characterization of local Nash equilibria in continuous games," *IEEE Trans. Autom. Control*, vol. 61, no. 8, pp. 2301–2307, Aug. 2016.
- [21] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, "Gradient descent only converges to minimizers," in *Proc. 29th Annu. Conf. Learn. Theory*, June 2016, pp. 1246–1257.
- [22] I. Abraham and T. D. Murphey, "Decentralized ergodic control: Distribution-driven sensing and exploration for multiagent systems," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 2987–2994, Oct. 2018.