Competitive Two Team Target Search Game with Communication Symmetry and Asymmetry

Michael Otte[1,3], Michael Kuhlman[2], and Donald Sofge[1]

[1] U.S. Naval Research Lab, Washington D.C. 20375
[2] University of Maryland, College Park, MD 20742
[3] National Research Council RAP Postdoctoral Associate at NRL.

Abstract. We study a search game in which two multiagent teams compete to find a stationary target at an unknown location. Each team plays a mixed strategy over the set of search sweep-patterns allowed from its respective random starting locations. Assuming that communication enables cooperation we find closed-form expressions for the probability of winning the game as a function of team sizes, and vs. the existence or absence of communication within each team. Assuming the target is distributed uniformly at random, an optimal mixed strategy equalizes the expected first-visit time to all points within the search space. The benefits of communication enabled cooperation increase with team size. Simulations and experiments agree well with analytical results.

Keywords: Multiagent System, Competitive Search, Search and Rescue, Search Game.

1 Introduction

We consider the problem of team-vs.-team competitive search, in which two teams of autonomous agents compete to find a stationary target at an unknown location. The game is won by the team of the first agent to locate the target. We are particularly interested in how coordination within each team affects the outcome of the game. We assume that intra-team communication is a prerequisite for coordination, and examine how the expected outcome of the game changes if one or both of the teams lack the ability to communicate—and thus coordinate.

This game models, e.g., an adversarial scenario in which we are searching for a pilot that has crashed in disputed territory, and we want to find the pilot before the adversary does (see Figure 1). Both we and the adversary have multiple autonomous aircraft randomly located throughout the environment to aid in our respective searches (e.g., that were performing unrelated missions prior to the crash), but neither agents nor adversaries have formulated a plan *a priori*. In this paper we answer the questions: How does team size affect game outcome? How beneficial is communication? What is an optimal search strategy?

In Section 4 we derive a closed-form expression for the expected outcome of an "ideal game" in which both teams search at the maximum rate for the entire game. A mixed Nash equilibrium exists at the point that each team randomizes

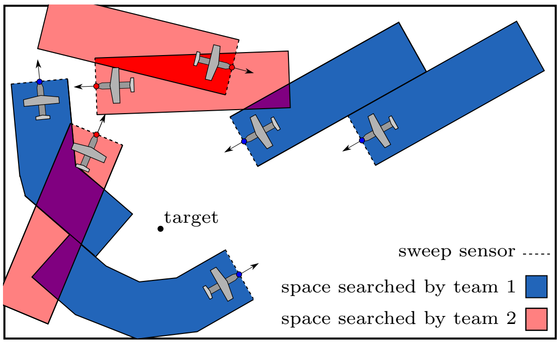

Fig. 1: Four agents (dark blue) compete against three adversaries (light red) to locate a target (black point). Communication enables members of the same team to cooperate (e.g., blue team), lack of communication prohibits cooperation (e.g., red team). The sweep sensors are, by definition, infinitesimally thin in the direction of travel.

their distributed searches such that all points are swept at the same expected time. An ideal game is impossible to realize in many environments due to a number of boundary conditions; however, it provides a useful model that allows us to evaluate how coordination affects game outcome. In Section 4.5 we extend these results by bounding performance in non-ideal cases, and find that non-ideal games become asymptotically close to ideal as the size of the environment increases toward infinity. Related work is discussed in Section 2. Nomenclature, a formal statement of assumptions, and the formal game definition appear in Section 3. Supporting simulations and experimental results appear in Section 5; discussion and conclusions appear in Sections 6 and 7, respectively.

2 Related Work

The target search problem was formalized at least as early as 1957 by [19], who studied aircraft detection of naval vessels in a probabilistic framework. Variations of the problem have been studied in many different communities, resulting in a vast body of related work. Indeed, even the subset of related work involving multiagent teams is too large to cover here. Extensive surveys of different formulations and approaches can be found in [6, 31, 8]. Previous work on target search ranges from the purely theoretical (differential equations [21], graph theory [29], game theory [26], etc.) to the applied (numerical methods [2, 12], control theory [11, 15], heuristic search [23], etc.) and borrows ideas from fields as diverse as economics [4, 9] and biology [18, 28].

One difference between the current paper and previous work is the scenario that we consider, in which two teams compete to locate a target first. In *cooperative search* a single team of agents attempts to locate one or more targets [30, 1] that may be stationary [15] or moving [18], and a key assumption is that all searchers cooperate. In contrast, we assume an adversarial relationship exists between two different teams of searchers.

Closely related *pursuit-evasion* games assume that one agent/team actively tries to avoid capture by another agent/team [13, 22], leading to an adversarial relationship between the searchers(s) and the target(s). *Capture the flag* [17] assumes that one team is attempting to steal a target that is guarded by the other team. Our scenario differs from both pursuit-evasion and capture the flag in that the adversarial relationship is between two different teams of searchers, each individually performing cooperative search for the same target.

Our work shares similarities with 1-dimensional *linear search* [7], and *cow path problems* [32, 24]. Differences include our extensions to higher dimensional spaces, which themselves build on coverage methods that use lawn-mower sweep patterns [5]. Our world model shares many of the same assumptions as [5]. In particular, an initial uniform prior distribution over target location and perfect sensors. Using sweep patterns for single agent coverage is studied by [5], while [30] extends these ideas to a single multiagent team searching for a moving and possibly evading target. Spires and Goldsmith use the idea of space filling curves to reduce the 2D search problem to a 1D problem [25].

Our work explicitly considers how each team's ability to communicate affects the expected outcome of the search game. This allows us to analyze scenarios in which teams have asymmetric communication abilities. A number of previous methods have considered limited communication, but have done so in different ways than those explored here. For example, robots were required to move such that a communication link could be maintained [1], and/or the ability to communicate between agents was assumed to be dependent on distance [27, 15], limited by bandwidth [11], adversaries [3], other constraints [14], or impossible [10].

3 Preliminaries

The search space is denoted X. The multiagent team is denoted G, the adversary team is denoted A, and an arbitrary team is denoted T, i.e., $T \in \{G, A\}$. There are $n = |G|$ agents in the multiagent team, and $m = |A|$ adversaries in the adversary team. The i-th agent is denoted g_i and the j-th adversary a_j, where $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. Both teams search for the same target q. Agents, adversaries, and target are idealized as points, and we abuse our notation by allowing them to indicate their locations in the search space, $g_i, a_j, q \in X$.

The term 'actor' is used to describe a member of the set $G \cup A$. The state space $S = X \times \Theta$ of a single actor includes position X and directional heading Θ. Let S_{g_i} represent the state space of the i-th agent. The product state space of the team is $\mathcal{S}_G = S_1 \times \dots \times S_n$. A particular configuration of the team is denoted $\mathbf{s}_G$, where $\mathbf{s}_G \in \mathcal{S}_G$. Similarly, for the adversary $\mathbf{s}_A \in \mathcal{S}_A = S_1 \times \dots \times S_m$. It is convenient to define the product space of *locations* for each team. Formally, $\mathbf{g} = (g_1, \dots, g_n) \in \mathcal{X}_G = X_1 \times \dots \times X_n$ and $\mathbf{a} = (a_1, \dots, a_m) \in \mathcal{X}_A = X_1 \times \dots \times X_m$ where we continue our abuse of notation that actors denote their own locations.

We use the subscript '0' to denote a starting value. For example, the starting location of g_i is $g_{i,0}$ and the starting configuration of the team is $\mathbf{g}_0$.

3.1 Assumptions

We consider search spaces embedded in D-dimensional Euclidean space, $X \subset \mathbb{R}^D$. We assume X is "well behaved" such that X is bounded, convex, and has a

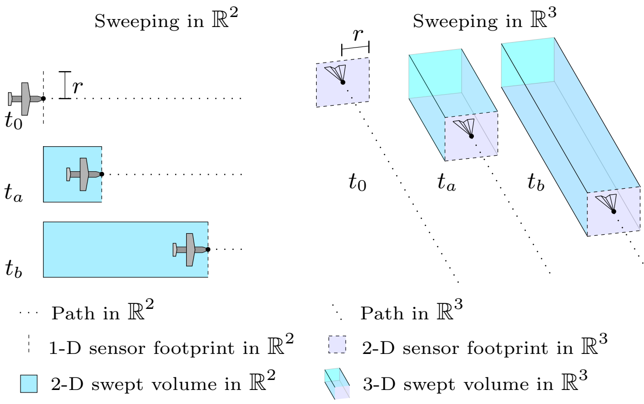

Fig. 2: Examples of sweep sensing in $\mathbb{R}^2$ (Left) and $\mathbb{R}^3$ (Right) for three different times $t_0 < t_a < t_b$. In $\mathbb{R}^2$ the sweep sensor footprint is a 1-D line segment oriented perpendicular to the direction of travel, in $\mathbb{R}^3$ it is a 2-D patch oriented perpendicular to the direction of travel. Swept volume increases as the robot moves forward.

boundary ∂X that can be decomposed into a finite number of locally Lipschitz continuous pieces. Our general formulation assumes X is a continuous[4] space.

The target is assumed to be stationary. Let $\mathcal{D}_X$ be the probability density function for a uniform distribution over X. Agents, adversaries, and targets are idealized as points, and have independent and identically distributed (i.i.d.) initial locations drawn from according to $\mathcal{D}_X$. The Lebesgue measure in $\mathbb{R}^D$ is denoted $\mathscr{L}_D(\cdot)$. Let Ω_X and Ω_S be the smallest σ-algebras over X and S, respectively. The (extension of) the Lebesgue measure in Ω_S is $\mathscr{L}_{\Omega_S}(\cdot)$. Given our assumptions, $\mathbb{P}(\hat{X}) = \int_{\hat{X}} \mathcal{D}_X(x) = \frac{\mathscr{L}_D(\hat{X})}{\mathscr{L}_D(X)}$ and $\mathbb{P}(\hat{S}) = \frac{\mathscr{L}_{\Omega_S}(\hat{S})}{\mathscr{L}_{\Omega_S}(S)}$ for all measurable subspaces $\hat{X} \subset X$ and $\hat{S} \subset S$, respectively, where $\mathbb{P}(\cdot)$ denotes the probability measure, and the integrals are Lebesgue. The probability spaces over starting locations and starting states are defined $(X, \Omega_X, \mathbb{P})$ and $(S, \Omega_S, \mathbb{P})$, respectively.

We assume agents and adversaries use *sweep sensors* with perfect accuracy (see Figure 2). A sweep sensor in $\mathbb{R}^D$ has an infinitesimally thin footprint defined by a subset of a $(D-1)$-dimensional hyperplane oriented perpendicular to the the direction of travel. We denote the sensor footprint B_r, where r refers to the radius of the smallest $(D-1)$-ball that contains the footprint, see Figure 2. Although $\mathscr{L}_D(B_r) = 0$ (e.g., the volume of a 2-dimensional disc is 0 in $\mathbb{R}^3$), the target is detected as the sensor footprint sweeps past it. We assume X is "large" in the sense that the minimum diameter of X is much greater than ("$\gg$") r. In "large" search spaces the sweep sensor provides a reasonable idealization of any sensor with finite observation volume[5].

[4] A discrete formulation can readily be obtained by replacing Lebesgue integrals over continuous spaces with summations over discrete sets, and reasoning about the probability of events directly instead of via probability density.

[5] Note that even for sensors with positive measure footprints, $0 < \mathscr{L}_D(B_r) < \infty$ (e.g., a D-ball instead of a $(D-1)$-ball) nearly all space is searched as the forward boundary of a sensor volume sweeps over it (in contrast to the space that is searched instantaneously at startup due to being within some agent's sensor volume).

We assume that agents and adversaries cannot detect the opposite team or physically interact with each other. This is a reasonable model when the environment is large such that chance encounters are unlikely. Actors are ignorant of their own teammates' locations *a priori* (for example, as if all actors are performing their own individual missions when the search scenario unexpectedly develops). Actors are assumed to sweep at a constant forward velocity v, where $0 < v < \infty$, and to have infinite rotational acceleration so that they are able to change direction instantaneously. Each actor may only change direction at most a countably infinite number of times[6]. We assume that each team is rational and will eventually sweep the entire space.

3.2 Paths, Multi-Paths, and Spaces

Let ρ denote a single actor's search path, $\rho \in S$. Let $B(s)$ denote the set of points in X swept by that actor's sensor when the actor is at a particular point $s \in \rho$. The set of points swept by an actor traversing path ρ is therefore $\bigcup_{s \in \rho} B(s)$. A *space covering* path is denoted $\hat{\rho}$ and has the property that its traversal will cause all points in the search space to be swept, i.e., $X \subset \bigcup_{s \in \hat{\rho}} B(s)$.

A search multi-path ψ_G is a set of paths containing one path ρ_i per agent in the team G, i.e., $\psi_G = \bigcup_{g_i \in G} \{\rho_i\}$. Let $\hat{\psi}_G$ denote a space covering search multi-path. One traversal of $\hat{\psi}_G$ by the members of G sweeps all points in the search space, $\bigcup_{s \in \mathbf{s}_G \in \hat{\psi}_G} B(s) \subset X$. Similarly quantities are defined for the adversary, $\psi_A = \bigcup_{a_j \in A} \{\rho_j\}$ and $\bigcup_{s \in \mathbf{s}_A \in \hat{\psi}_A} B(s) \subset X$. Let Ψ_G (and Ψ_A) denote the space of all possible search multi-paths given a team's state space $\mathcal{S}_G$ (and adversary's state space $\mathcal{S}_A$). Formally, $\Psi_G = \bigcup \{\psi \,|\, \mathcal{S}_G\}$ and $\Psi_A = \bigcup \{\psi \,|\, \mathcal{S}_A\}$.

3.3 Communication and Coordination Models

The function $\mathcal{C} : \{G, A\} \to \{0, 1\}$ denotes the communication ability of a team. Communication within a particular team T is either assumed to be perfect $\mathcal{C}(T) = 1$ or nonexistent $\mathcal{C}(T) = 0$. That is, team members can either communicate always or never. Communication enables coordination, which allows the team to find a target more quickly in expectation. When $\mathcal{C}(T) = 1$, the members of T *attempt* to equally divide the effort of searching X such that each $x \in X$ is swept by exactly one agent and each agent travels an equal distance. We investigate the 2 by 2 space of game scenarios this allows, $\mathcal{C}(G) \times \mathcal{C}(A) = \{0, 1\} \times \{0, 1\}$.

3.4 Game Formulation

Given our assumptions, the first team to sweep the target's location wins the game. The family of *competitive team target search games* we consider is defined:

[6]This prevents "cheating" where an agent that continuously rotates through an uncountably infinite number of points is able to use its zero-measure sweep sensor as if it were a volumetric sensor of non-zero-measure (the measure of a countably infinite union of sweep footprints is still 0).

Random team target search games: *Given a search space X, a stationary target $q \in X$, a multiagent team $G = \{g_1, \ldots, g_n\}$, and an adversary team $A = \{a_1, \ldots, a_m\}$; with initial locations drawn i.i.d. from $\mathcal{D}_X(q)$, $\mathcal{D}_X(g_{i,0})$, and $\mathcal{D}_X(g_{i,0})$, respectively; communication $\mathcal{C}(G), \mathcal{C}(A) \in \{0,1\}$; and chosen movement along multi-paths $\psi_G \subset X$ and $\psi_A \subset X$; then team G wins iff $q \in B(g_i)$ for some $g_i \in \mathbf{s}_G \in \psi_G$ before $q \in B(a_j)$ for some $a_j \in \mathbf{s}_A \in \psi_a$.*

3.5 Game Outcomes, Multi-Path Spaces, and Strategies

Let $\Omega_{\text{outcome}} = \{\omega_{\text{lose}}, \omega_{\text{win}}, \omega_{\text{tie}}\}$ denote the space of game outcomes, where ω_{win} is the event that a member of team G finds the target first, ω_{lose} denotes the event that an adversary finds the target first, and ω_{tie} denotes a tie. Given our formulation within a continuous space, ties are a measure 0 set, $\mathbb{P}(\omega_{\text{tie}}) = 0$, that can be ignored for the purposes of analyzing expected performance. In discrete space one could break ties in a number of ways, e.g., by randomly selecting the actor that finds the target first.

Strategies are equivalent to multipaths—any valid multi-path ψ_G that starts at $\mathbf{g}_0$ is a particular search strategy for team G. Let Ψ denote the space of all strategies. Let Ψ_G be a function that maps starting configurations $\mathbf{g}_0$ to the subset of all valid strategies for G that begin at $\mathbf{g}_0$. Let Ω denote the (smallest) σ-algebra over Ψ. Formally, $\Psi_G : \mathcal{S}_G \to \Omega$. The subset of all valid strategies available to G given $\mathbf{g}_0$ is thus denoted $\Psi_G(\mathbf{g}_0)$, where $\Psi_G(\mathbf{g}_0) \subset \Psi$.

A *conditional mixed strategy* is both: (1) conditioned on the event that team G starts at a particular $\mathbf{g}_0$, and (2) mixed such that the particular strategy $\psi_G \in \Psi_G(\mathbf{g}_0)$ used by team G is drawn at random from $\Psi_G(\mathbf{g}_0)$ according to a chosen probability density $\mathcal{D}_{\mathbf{g}_0}(\psi_G)$. By designing $\mathcal{D}_{\mathbf{g}_0}(\psi_G)$ appropriately, it is possible for team G to play any valid conditional mixed strategy given $\mathbf{g}_0$.

Given $\mathcal{D}_{\mathbf{g}_0}(\psi_G)$, a probability measure function $\mathbb{P}_{\mathbf{g}_0}$ can be constructed such that $\int_{\Psi_G(\mathbf{g}_0)} \mathcal{D}_{\mathbf{g}_0}(\psi_G) = 1$ and such that for all subsets $\hat{\Psi} \subset \Psi_G(\mathbf{g}_0)$ we have $\mathbb{P}_{\mathbf{g}_0}(\psi_G \in \hat{\Psi}) = \int_{\hat{\Psi}} \mathcal{D}_{\mathbf{g}_0}(\psi_G)$. A particular conditional mixed strategy (conditioned on $\mathbf{g}_0 \in \mathcal{S}_G$) is thus a probability space that can be represented by the triple $(\Psi_G(\mathbf{g}_0), \Omega_G(\mathbf{g}_0), \mathbb{P}_{\mathbf{g}_0})$, where $\Omega_G(\mathbf{g}_0)$ is the (smallest) σ-algebra over $\Psi_G(\mathbf{g}_0)$.

A *mixed strategy* $(\Psi_G, \Omega_G, \mathbb{P}_G)$ is the set of conditional mixed strategies over all $\mathbf{g}_0 \in \mathcal{S}_G$, where $\Omega_G = \bigcup_{\mathbf{g}_0 \in \mathcal{S}_G} \Omega_G(\mathbf{g}_0)$ and $\mathbb{P}_G(\psi \mid \mathbf{g}_0) = \mathbb{P}_{\mathbf{g}_0}(\psi)$ for all $\mathbf{g}_0 \in \mathcal{S}_G$. *Note that a mixed strategy triple is not a probability space*, per se, because it does not include the probability measure of the starting configurations $\mathbf{g}_0$. That said, when a mixed strategy is combined with such a measure, e.g., the measure implied by $\mathcal{D}_X$, then a probability space is the result. Analogous quantities, $(\Psi_A, \Omega_A, \mathbb{P}_A)$ and $(\Psi_A(\mathbf{a}_0), \Omega_A(\mathbf{a}_0), \mathbb{P}_{\mathbf{a}_0})$, are defined for the adversary.

Given our assumption that the two teams cannot detect each other, one team's mixed strategy is necessarily independent of the other team's starting location. Let $t(\psi, x)$ denote the earliest time at which a team following ψ sweeps location $x \in X$. Given ψ_G and ψ_A, and a target at q (with location unknown to either team), team G wins if and only if $t(\psi_G, q) < t(\psi_A, q)$. Let $X_{\text{win}}(\psi_G, \psi_A) \subset X$

denote the subset of the search space where $t(\psi_G, x) < t(\psi_A, x)$.

$$X_{\text{win}}(\psi_G, \psi_A) = \{x \in X \mid t(\psi_G, x) < t(\psi_A, x)\}.$$

Team G wins if and only if $q \in X_{\text{win}}$. When G plays ψ_G and A plays ψ_A, we get:

Proposition 1. *Assuming the target is located uniformly at random in X, the probability team G wins is equal to the ratio of search space it sweeps before the adversary,* $\mathbb{P}(\omega_{\text{win}}|\psi_G, \psi_A) = \frac{\mathscr{L}_D(X_{\text{win}}(\psi_G,\psi_A))}{\mathscr{L}_D(X)}$.

The probability team G wins in a particular search space while playing a particular adversary is calculated by integrating $\mathcal{D}_{\mathbf{g}_0}(\psi_G)$ over $\Psi_G(\mathbf{g}_0)$ for all $\mathbf{g}_0$ and $\mathcal{D}_{\mathbf{a}_0}(\psi_A)$ over $\Psi_A(\mathbf{a}_0)$ for all $\mathbf{a}_0$. Assuming the target and teams are distributed uniformly at random, this is calculated:

$$\mathbb{P}(\omega_{\text{win}}) = \frac{1}{\mathscr{L}_{\Omega_{\mathcal{S}_G}}(\mathcal{S}_G)} \int_{\mathcal{S}_G} \frac{1}{\mathscr{L}_{\Omega_{\mathcal{S}_A}}(\mathcal{S}_A)} \int_{\mathcal{S}_A} \int_{\Psi_G(\mathbf{g}_0)} \mathcal{D}_{\mathbf{g}_0}(\psi_G) \int_{\Psi_A(\mathbf{a}_0)} \mathcal{D}_{\mathbf{a}_0}(\psi_A) \frac{\mathscr{L}_D(X_{\text{win}}(\psi_G, \psi_A))}{\mathscr{L}_D(X)} \quad (1)$$

where the Lebesgue integrals are respectively over all $\mathbf{g}_0 \in \mathcal{S}_G$, all $\mathbf{a}_0 \in \mathcal{S}_A$, all $\psi_G \in \Psi_G(\mathbf{g}_0)$ and all $\psi_A \in \Psi_A(\mathbf{a}_0)$.

We use "$*$" to denote quantities related to optimality. An optimal mixed strategy is defined: $(\Psi_G^*, \Omega_G, \mathbb{P}_G^*) = \arg\max_{((\Psi_G, \Omega_G, \mathbb{P}_G))} \mathbb{P}(\omega_{\text{win}})$.

4 Optimal Strategies for Ideal Games

Let X_{swept} denote the space team G has swept (X_{swept} is different from X_{win} in that X_{swept} may include space that has also been swept by the adversary). The instantaneous rate team G sweeps new space is given by: $\frac{d}{dt}[\mathscr{L}_D(X_{\text{swept}})]$. The optimal instantaneous rate at which an agent sweeps new space can be expressed as the agent's velocity multiplied by the $(D-1)$-dimensional hypervolume of the sensor footprint: $v\mathscr{L}_{D-1}(B_r)$. Given our assumptions, we have the following:

Proposition 2. *The optimal instantaneous normalized rate that a single agent sweeps new space is:* $c^* = v\frac{\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$.

4.1 Both Teams Can Communicate (ideal case)

The optimal instantaneous normalized rate (c^*) occurs when there is no sensor overlap between agents. Building on Proposition 2 we get:

Corollary 1. *The optimal instantaneous normalized rate that n agents can co-operatively sweep new space is:* $\frac{d^*}{dt}[\frac{\mathscr{L}_D(X_{\text{swept}})}{\mathscr{L}_D(X)}] = nv\frac{\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)} = nc^*$.

In an "ideal" cooperative search we assume that the team can maintain the optimal rate of sweep for the entire duration of search. The time required for an ideal search with n agents is $t_{n,\text{sweep}} = 1/(nc^*)$. The game is guaranteed to end by time $t_{\text{final}} = \min(t_{n,\text{sweep}}, t_{m,\text{sweep}})$.

We observe that any bias or predictability by a particular team (e.g., a mixed strategy that leads to a subset of the environment being swept sooner or later in expectation, over all possible starting locations) could be exploited by the opposing team. This observation leads to the following proposition.

Proposition 3. *A mixed strategy that causes some portion of the environment to be swept sooner or later, in expectation, over the set of all strategies and distributions of agent and adversary starting locations is a suboptimal strategy.*

As a corollary of proposition 3 we have the following:

Corollary 2. *If an optimal ideal mixed strategy $(\Psi_G^*, \Omega_G, \mathbb{P}_G^*)$ exists for a team G, then in that strategy the first sweep time for any point $x \in X$ is distributed uniformly at random between 0 and $t_{n,\text{sweep}}$ (over the space of all possible starting configurations).*

In an *ideal game* each team plays an optimal mixed strategy over a set of ideal search strategies. The following is true by the definition of a Nash equilibrium:

Proposition 4. *Assuming optimal ideal strategies exist for both teams, a mixed strategy Nash equilibrium exists when both teams play an optimal mixed strategy.*

At such a Nash equilibrium, the first sweep time of any point x by one team is completely decorrelated from the first sweep time of x by the other team (over the space of all possible actor starting locations).

Let $X_{\text{new}}(t)$ be the space that has not yet been swept by either team by time t, and $\frac{d}{dt}[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}]$ be the instantaneous normalized rate team G sweeps this unswept space at time t. We note that, given a particular ψ_G and ψ_A,

$$\frac{\mathscr{L}_D(X_{\text{win}}(\psi_G, \psi_A))}{\mathscr{L}_D(X)} = \int_0^{t_{\text{final}}} \frac{d}{dt}\left[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}\right] dt,$$

where $t_{\text{final}} = \min(\frac{1}{nc^*}, \frac{1}{mc^*})$. Thus, Equation 1 can be reformulated for the Nash equilibrium of an ideal game with cooperation within both teams as:

$$\mathbb{P}\left(\omega_{\text{win}}^*\right) = \frac{1}{\mathscr{L}_{\Omega_{S_G}}(\mathcal{S}_G)} \int_{\mathcal{S}_G} \frac{1}{\mathscr{L}_{\Omega_{S_A}}(\mathcal{S}_A)} \int_{\mathcal{S}_A} \int_{\Psi_G^*(\mathbf{g}_0)} \int_{\Psi_A^*(\mathbf{g}_0)} \mathcal{D}_{\mathbf{g}_0}(\psi_G)\mathcal{D}_{\mathbf{a}_0}(\psi_A) \int_0^{t_{\text{final}}} \frac{d}{dt}\left[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}\right] dt$$

where integrals are Lebesgue. Using the independence of the two team's optimal mixed strategies, i.e., $\mathcal{D}_{\mathbf{g}_0}(\psi_G)$ and $\mathcal{D}_{\mathbf{a}_0}(\psi_A)$ for all $\mathbf{g}_0$ and $\mathbf{a}_0$ yields:

$$\mathbb{P}\left(\omega_{\text{win}}^*\right) = \int_0^{t_{\text{final}}} \frac{1}{\mathscr{L}_{\Omega_{S_G}}(\mathcal{S}_G)\mathscr{L}_{\Omega_{S_A}}(\mathcal{S}_A)} \int_{\mathcal{S}_G} \int_{\mathcal{S}_A} \int_{\Psi_G^*(\mathbf{g}_0)} \int_{\Psi_A^*(\mathbf{g}_0)} \mathcal{D}_{\mathbf{g}_0}(\psi_G)\mathcal{D}_{\mathbf{a}_0}(\psi_A) \frac{d}{dt}\left[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}\right] dt$$

We observe that the quantity inside the outermost integral describes the expected value of $\frac{d}{dt}[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}]$ over all $\mathcal{S}_G$, $\mathcal{S}_A$, Ψ_G^*, and Ψ_A^*. For brevity we denote the expected value of '$\cdot$' over all $\mathcal{S}_G$, $\mathcal{S}_A$, Ψ_G^*, and Ψ_A^* as $\mathbb{E}^*[\cdot]$, i.e., $\mathbb{E}^*[\cdot] \equiv \mathbb{E}_{\mathcal{S}_G,\mathcal{S}_A,\Psi_G^*,\Psi_A^*}[\cdot]$. Thus, formally,

$$\mathbb{E}^*\left[\frac{d}{dt}\mathscr{L}_D(X_{\text{new}}(t))\right] = \frac{1}{\mathscr{L}_{\Omega_{S_G}}(\mathcal{S}_G)\mathscr{L}_{\Omega_{S_A}}(\mathcal{S}_A)} \int_{\mathcal{S}_G} \int_{\mathcal{S}_A} \int_{\Psi_G(\mathbf{g}_0)} \int_{\Psi_A(\mathbf{a}_0)} \mathcal{D}_{\mathbf{g}_0,}(\psi_G)\mathcal{D}_{\mathbf{a}_0}(\psi_A) \frac{d}{dt}\left[\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}\right]$$

Lemma 1. *Assuming optimal ideal mixed strategies exist and both teams play an optimal ideal mixed strategy,*

$$\mathbb{E}^*\left[\frac{d}{dt}\mathscr{L}_D(X_{\text{new}}(t))\right] = (1 - tmc^*)\, nc^* \tag{2}$$

Proof. At time t the adversary (operating according to its own ideal optimal strategy) has swept $tmv\frac{\mathscr{L}_D(B_r)}{\mathscr{L}_D(X)}$ portion of the entire search space. The interplay between the mixed ideal optimal strategies for each team forces the expected instantaneous overlap between teams to be uncorrelated. Thus, for all $t \in [0, t_{\text{final}}]$, the instantaneous expected rate team G sweeps $\mathscr{L}_D(X_{\text{new}})$ is discounted by a factor of $1 - tmv\frac{\mathscr{L}_D(B_r)}{\mathscr{L}_D(X)}$ vs. $\frac{d^*}{dt}\mathscr{L}_D(X_{\text{swept}})$.

$$\mathbb{E}^*\left[\frac{d}{dt}\mathscr{L}_D(X_{\text{new}}(t))\right] = \left(1 - tmv\frac{\mathscr{L}_D(B_r)}{\mathscr{L}_D(X)}\right)\frac{d^*}{dt}\left[\frac{\mathscr{L}_D(X_{\text{swept}})}{\mathscr{L}_D(X)}\right].$$

Substitution with Proposition 2 and Corollary 1 yields the desired result. □

In other words, team G covers new territory at a rate that decreases, in expectation, proportionally to the proportion of space the adversary has covered up to time t. Substituting this result back into the previous equations and noting that $t_{\text{final}} = 1/(c^* \max(n, m))$:

$$\mathbb{P}(\omega^*_{\text{win}}) = \int_0^{1/(c^* \max(n,m))} (1 - tmc^*)\, nc^*\, dt.$$

Solving this equation yields the following theorem.

Corollary 3. *The probability team G wins an ideal game assuming both G and A are able to communicate and play optimal ideal mixed strategies is*

$$\mathbb{P}(\omega^*_{\text{win}}) = \begin{cases} n/(2m) & \text{when } n \leq m \\ 1 - m/(2n) & \text{when } n \geq m \end{cases}.$$

4.2 Case 2: multiagent team G cannot communicate but the adversary team A can (ideal case)

Given that the starting location of each agent on team G is sampled i.i.d. a game in which team G cannot communicate is equivalent to the situation in which the adversary team A plays n sub-games, one vs. each member $\{g_i\} \subset G$, and A wins the overall game if and only if it wins all n sub-games. Because the target location is identical for each of these n sub-games, the n games are not indipendent (as was erroniously assumed in a preliminary version of this work). To win, the adversary must sweep the target before the particular agent of G that happens to sweep the target first among all members of G. This can be calculated by reformulating Equation 1 to integrate over the distribution of the smallest of n first sweep times:

$$\mathbb{P}(\omega^*_{\text{lose}}) = \frac{m - (m-1)^{n+1}m^{-n}}{n+1}$$

Combining $\mathbb{P}(\omega^*_{\text{tie}}) = 0$ with Corollary 3 we get:

Corollary 4. *The probability team G wins an ideal game, assuming team G cannot communicate but the adversary team A can, and the adversary team A plays optimal ideal mixed strategies, while each $\{g_i\} \subset G$ individually plays an optimal ideal mixed strategy, is:* $\mathbb{P}(\omega^*_{\text{win}}) = 1 - \mathbb{P}(\omega^*_{\text{lose}})$.

4.3 Case 3: Team G can communicate but the adversary's team A cannot (ideal case)

This case is complementary to the previous one, due to symmetry and the fact that $\mathbb{P}\left(\omega^*_{\text{tie}}\right) = 0$. We swap n and m and also ω_{lose} and ω_{win} from the results in the previous section to get:

Corollary 5. *The probability team G wins an ideal game, assuming team G can communicate but the adversary team A cannot, and team G plays an optimal ideal mixed strategy, while each of the adversary team's individual uncoordinated sub-teams $\{a_j\} \subset A$ for $1 \leq j \leq m$ plays an optimal ideal mixed strategy, is:*

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \frac{n - (n-1)^{m+1} n^{-m}}{m+1}$$

We also note that $\mathbb{P}\left(\omega^*_{\text{lose}}\right) = 1 - \mathbb{P}\left(\omega^*_{\text{win}}\right)$.

4.4 Case 4: Neither team can communicate (ideal case)

The case when neither team has communication must be analyzed separately, but is somewhat trivial.

Theorem 1. *The probability team G wins an ideal game, assuming no team can communicate but all actors individually play an optimal ideal mixed strategy, is:*

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \frac{n}{n+m} \qquad \mathbb{P}\left(\omega^*_{\text{lose}}\right) = \frac{m}{n+m}.$$

Proof. Our assumption of uniformly random i.i.d. starting locations of actors and target, combined with the fact that optimal mixed strategies decorrelate the expected sweep time of any particular point x, means that, in expectation, each actor has a $1/(n+m)$ chance of being the agent with the least amount of travel (i.e., time) required to sweep q. The probability that team G finds the target before A can be calculated as the ratio of agents to total actors. □

4.5 Extensions to non-ideal games

The realization of an ideal game requires that an optimal mixed strategy exists such that Equation 2 holds. In practice, this idealization is often broken by both the startup locations of the actors and the boundary of the search space (see Figure 3). However, it is possible to modify the equations for an ideal game to obtain bounds on $\mathbb{P}\left(\omega_{\text{win}}\right)$. This is accomplished by breaking the multiagent search into two mutually exclusive phases: (1) a phase containing all portions of the search wherein the agents of G *are not* able to perform an ideal search and (2) a separate phase containing all other (ideal) portions of the search. Let t_{startup} denote the time required for the non-ideal portion of search. We assume that the adversaries in A are able to maintain an ideal search rate for the entire game, which produces a lower bound on $\mathbb{P}\left(\omega_{\text{win}}\right)$. Reversing the roles of adversaries

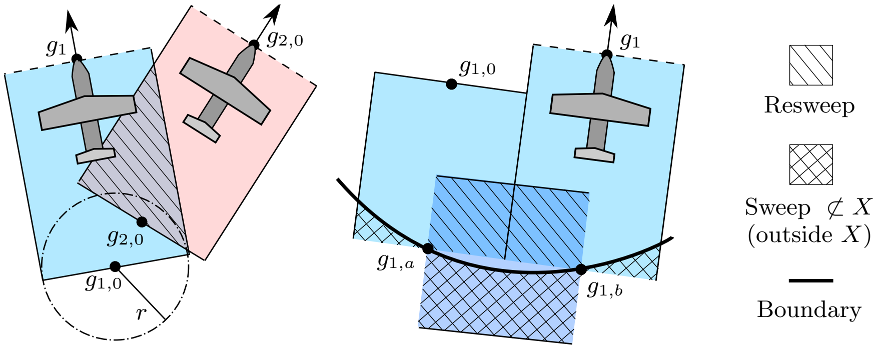

Fig. 3: Left: Two agents start within r of each other, causing some area to be swept by both of them (hashed). Right: a single robot turns near the boundary of the search space as it moves from $g_{1,0} \to g_{1,a} \to g_{1,b} \to g_1$; this causes some space to be swept multiple times (hashed) and space outside the search space to be swept (criss-crossed).

and agents provides an upper bound on $\mathbb{P}(\omega_{\text{win}})$; let $\tilde{t}_{\text{startup}}$ be the time it takes A to perform non-ideal search in this case. Recall that $c* = v\frac{\mathscr{L}_D(B_r)}{\mathscr{L}_D(X)}$. Also let $\hat{t}_{\text{final}} = \min(t_{\text{startup}} + \frac{1}{nc^*}, \frac{1}{mc^*})$ and $\tilde{t}_{\text{final}} = \min(\tilde{t}_{\text{startup}} + \frac{1}{mc^*}, \frac{1}{nc^*})$.

Theorem 2. *Assuming that both teams can communicate, and an optimal mixed strategy exists for both teams, and that both teams play an optimal mixed strategy, and that the game is ideal in every sense except for starting locations and boundary effects, the probability team G wins is bounded as follows:*

$$\left[\int_{t_{\text{startup}}}^{\hat{t}_{\text{final}}} (1 - tmc^*)\, nc^*\ dt\right] - t_{\text{startup}} mc^* \leq \mathbb{P}(\omega_{\text{win}}) \leq 1 - \left[\int_{\tilde{t}_{\text{startup}}}^{\tilde{t}_{\text{final}}} (1 - tnc^*)\, mc^*\ dt\right] - \tilde{t}_{\text{startup}} nc^*.$$

Proof. (Sketch) Non-ideal effects become increasingly detrimental to G's probability of winning the game as they occur earlier and earlier in the game. Thus, it is possible to construct a scenario that is even worse than a worst-case non-ideal search (in terms of team G's probability of winning the game) by: (1) assuming that all negative ramifications of a non-ideal search happen at the beginning of the search for team G, instead of whenever they actually occur, and (2) assuming the adversary team A is allowed to realize an ideal sweep rate for the entire game. The length of the non-ideal startup phase for the worse-than-worst case can be bounded as follows: $t_{\text{startup}} < c_1 r \mathscr{L}_{D-1}(\partial X)$, where c_1 is a dimensionally dependent constant, r is sweep radius, and $\mathscr{L}_{D-1}(\partial X)$ is the surface area of the search space boundary. □

Theorem 2 shows that the proportion of time spent dealing with non-ideal startup and boundary approaches zero as environments get larger vs. sensor range, $\lim_{r \to 0} \frac{t_{\text{startup}}}{t_{\text{final}}} = 0$. In other words, the ideal equations model the non-ideal case more-and-more accurately as the size of the environment increases.

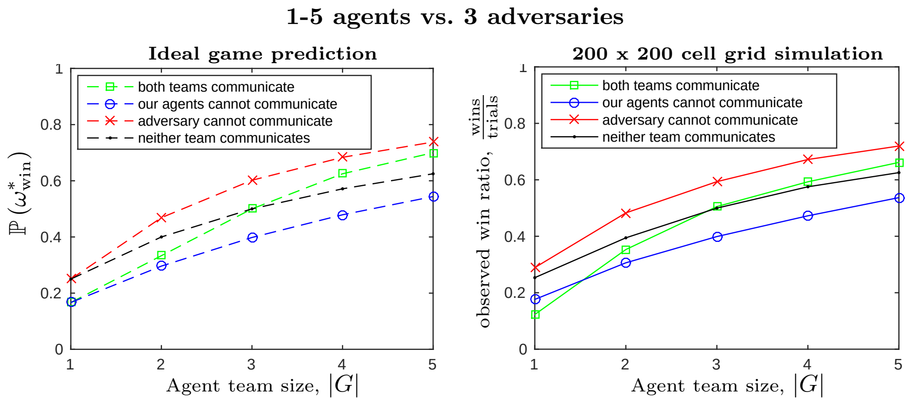

Fig. 4: Comparison of results predicted by analysis of the ideal case (left) vs. those from a simulation experiment (right) for games with 1 to 5 Agents vs. 3 Adversaries. Vertical axes measure the probability team G wins vs. its observed win ratio in repeated random trials performed in simulation, respectively. Colors denote which teams may or may not communicate. Each datapoint represents the mean result over 10^4 trials.

5 Simulations and Experiments

We compare the results derived in the previous section to repeated trials of search and rescue in contested environments performed both in simulation and on a mixed platform of real and virtual agents. For these simulations and experiments we assume a discrete grid environment where movement is allowed along the cardinal directions (note that this contrasts with the more general continuous space formulation assumed in previous sections).

Multipaths are selected from a library of predefined sweep patterns, such that each pattern forms a cycle, sweeps the entire space, and are designed to minimize sweep overlap between different parts of the search. If an agent/adversary cannot communicate with its team then it moves to the nearest point on a randomly selected cycle and then follows it. If an agent (resp. adversary) can communicate with its team then all team members agree on a cycle, divide the path into n (resp. m) contiguous sub-paths, and then allocate one sub-path per team member such that the cumulative distance traveled by the team to their start locations is minimized. Next, each agent/adversary moves to its start point and then searches its allotted sub-path. Unlike the ideal case, the probability of a tie is non-zero; ties are broken by a coin toss weighted proportionally to the number of {team members vs. all actors} that simultaneously discover the target.

Simulations are run in the Julia language using a 200 by 200 meter search space composed of 1 by 1 meter grid cells. B is an L-∞ ball of radius 5 meters. Locations of actors and target are determined uniformly at random over the set of grid cells. Selected results comparing predictions based on the ideal case vs. the average results from simulation of 4×10^5 trials (10^4 trials per datapoint) are

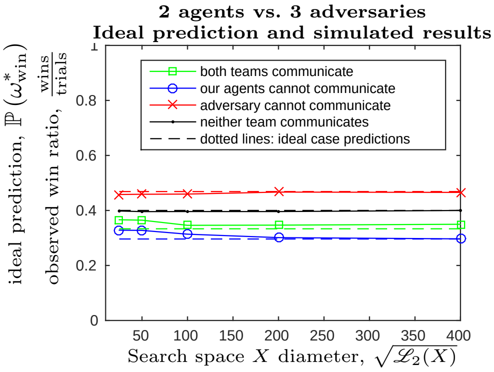

Fig. 5: The ideal case (dashed lines) more accurately predicts the mean result of Monte Carlo simulations (solid lines) as $\mathscr{L}_2(X)$ increases. These particular results are for games with 2 agents vs. 3 adversaries over various $\sqrt{\mathscr{L}_2(X)}$. Colors represent which teams may or may not communicate. Each datapoint represents the mean result over 10^4 trials.

presented in Figures 4 and 5, each datapoint in the simulation curves represents the mean results over 10^4 repeated trials.

The mixed platform combines Asctec Pelican quadrotor UAVs that have onboard Odroid single board computers with simulated agents that run on a laptop (Ubuntu 14.04). The quadrotors receive position measurements from a Vicon motion capture system and runs ETH-Zurich modular sensor fusion framework for state estimation [20]. Robot Operating System (ROS) is used on all computers for local interprocess communications and NRL's *Puppeteer* framework is used for coordination of all vehicles, which uses Lightweight Communications and Marshalling [16] for intervehicle communications. Grid cells are 2 by 2 meters, and the contested search space is 12 by 12 meters. We use a virtual target sensor such that an actor discovers the target if their locations are closer than 1 meter. All actors fly at an altitude of 2 meters, which corresponds to a field of view of approximately 60° when searching for ground targets with a downward facing camera. A random number generator is used to determine start location of the real actors as well as the virtual actors and the target. We perform repeated trials for a two agent team (consisting of one Asctec Pelican and one virtual agent) vs. an adversary (Asctec Pelican). We perform 10 successful trials: 5 trials for the case where team G can communicate and 5 for the case where it cannot. Results from experiments with the mixed platform appear in Figure 6.

6 Discussion

Our analysis, simulations, and experiments show that using the ideal case to predict $\mathbb{P}(\omega_{\text{win}})$ works reasonably well, and provides a more accurate prediction as the size of the environment increases. We also show that the relative effects of non-ideal startup locations and boundary conditions vanish, in the limit, as the size of the environment increases.

With respect to communication symmetry vs. asymmetry, our results verify the intuition that team G benefits from a situation in which G can communicate and team A cannot. More interesting is the result that moving from a scenario where both G and A can communicate to a scenario where neither G nor A can

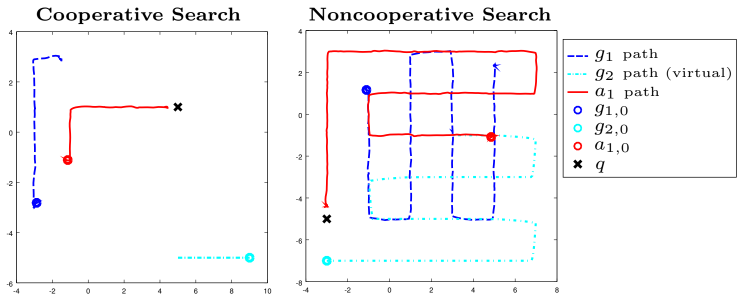

Cooperation

Trial	Winner
1	G
2	G
3	G
4	A
5	A

No Cooperation

Trial	Winner
1	A
2	G
3	A
4	G
5	A

Fig. 6: Mixed platform experiments. A two agent team (consisting of a Pelican Quadrotor, blue, and a simulated agent, light-blue) vs. an adversary (Pelican Quadrotor, red) to find a target (black). Left and Center: Examples of paths when agents do and do not collaborate, respectively. Right: Game outcomes of repeated trials.

communicate benefits G only if $n < m$. The advantages of performing a coordinated search vs. uncoordinated search increase vs. team size. Uncommunicating larger teams will outperform uncommunicating smaller teams, in general.

7 Summary and Conclusions

We study the effects of cooperation on multiagent two-team competitive search games, a class of games in which two multiagent teams compete to locate a stationary target placed at an unknown location. Given an assumption that communication is required for coordination, this enables us to analyze how communication symmetry and asymmetry between teams affects the outcome of the game. For the case involving perfect finite sweep sensors, random initial placement of actors/target, and non-observability of the other team's movements, we find closed-form solutions for the probability of winning an "ideal game" in which transient boundary effects are ignored.

A team maximizes its chances of winning by playing a mixed strategy such that all points are eventually swept, the expected time a point is (first) swept is identical for all points, and there is as little search overlap as possible. A Nash equilibrium exists for an ideal game.

The chances of winning the search game increase vs. team size, and also increase if the team is able to communicate. Moving from a situation in which both teams can communicate to a situation where neither team can communicate will benefit the smaller team and hinder the larger team (this effect becomes stronger as the difference between the two teams' sizes increases).

Monte Carlo simulations over random start locations and experimental results on a platform with AscTec Pelican quadrotor UAVs validate that the observed outcomes of non-ideal games are predicted reasonably well by equations derived for the ideal case, and that these predictions become more accurate as the size of the search space increases.

Acknowledgments

We would like to thank Colin Ward, Corbin Wilhelmi, and Cyrus Vorwald for their help in facilitating the mixed platform experiments.

This work was performed at the Naval Research Laboratory and was funded by the Office of Naval Research under grant numbers N0001416WX01271 and N0001416WX01272. The views, positions and conclusions expressed herein reflect only the authors opinions and expressly do not reflect those of the Office of Naval Research, nor those of the Naval Research Laboratory.

References

1. Beard, R.W., McLain, T.W.: Multiple uav cooperative search under collision avoidance and limited range communication constraints. In: Decision and Control, 2003. Proceedings. 42nd IEEE Conference on. vol. 1, pp. 25–30 Vol.1 (Dec 2003)
2. Bertuccelli, L.F., How, J.P.: Robust UAV search for environments with imprecise probability maps. In: IEEE Conf. on Decision and Control and the European Control Conferenc. pp. 5680–5685 (Dec 2005)
3. Bhattacharya, S., Khanafer, A., Başar, T.: A Double-Sided Jamming Game with Resource Constraints, pp. 209–227. Springer International Publishing (2016)
4. Chandler, P., Pachter, M.: Hierarchical control for autonomous teams. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference. pp. 632–642 (2001)
5. Choset, H., Pignon, P.: Coverage path planning: The boustrophedon cellular decomposition. In: Field and Service Robotics. pp. 203–209. Springer (1998)
6. Chung, T.H., Hollinger, G.A., Isler, V.: Search and pursuit-evasion in mobile robotics. Auton. Robots 31(4), 299–316 (2011)
7. Demaine, E.D., Fekete, S.P., Gal, S.: Online searching with turn cost. Theoretical Computer Science 361(2), 342–355 (2006)
8. 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 (July 2006)
9. Dias, M.B.: Traderbots: A new paradigm for robust and efficient multirobot coordination in dynamic environments. Ph.D. thesis, Carnegie Mellon University Pittsburgh (2004)
10. Feinerman, O., Korman, A., Lotker, Z., Sereni, J.S.: Collaborative search on the plane without communication. In: Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing. pp. 77–86. PODC '12, ACM, New York, NY, USA (2012), `http://doi.acm.org/10.1145/2332432.2332444`
11. Flint, M., Polycarpou, M., Fernandez-Gaucherand, E.: Cooperative control for multiple autonomous uav's searching for targets. In: Decision and Control, 2002, Proceedings of the 41st IEEE Conference on. vol. 3, pp. 2823–2828 vol.3 (Dec 2002)
12. Forsmo, E.J., Grotli, E.I., Fossen, T.I., Johansen, T.A.: Optimal search mission with unmanned aerial vehicles using mixed integer linear programming. In: Unmanned Aircraft Systems (ICUAS), 2013 International Conference on. pp. 253–259 (May 2013)
13. Gerkey, B.P., Thrun, S., Gordon, G.: Parallel stochastic hill-climbing with small teams. In: Multi-Robot Systems. From Swarms to Intelligent Automata Volume III, pp. 65–77. Springer (2005)

14. Hollinger, G.A., Yerramalli, S., Singh, S., Mitra, U., Sukhatme, G.S.: Distributed data fusion for multirobot search 31(1), 55–66 (2015)
15. Hu, J., Xie, L., Lum, K.Y., Xu, J.: Multiagent information fusion and cooperative control in target search 21(4), 1223–1235 (July 2013)
16. Huang, A.S., Olson, E., Moore, D.C.: Lcm: Lightweight communications and marshalling. In: Intelligent robots and systems (IROS), 2010 IEEE/RSJ international conference on. pp. 4057–4062. IEEE (2010)
17. Huang, H., Ding, J., Zhang, W., Tomlin, C.J.: Automation-assisted capture-the-flag: A differential game approach. IEEE Transactions on Control Systems Technology 23(3), 1014–1028 (2015)
18. Kim, M.H., Baik, H., Lee, S.: Response threshold model based uav search planning and task allocation. Journal of Intelligent & Robotic Systems 75(3), 625–640 (2013)
19. Koopman, B.: The theory of search. ii. target detection. Operations Research 4(5), 503–531 (1956)
20. Lynen, S., Achtelik, M.W., Weiss, S., Chli, M., Siegwart, R.: A robust and modular multi-sensor fusion approach applied to mav navigation. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 3923–3929 (Nov 2013)
21. Mangel, M.: Marcel Dekker, New York (1989)
22. Noori, N., Isler, V.: Lion and man with visibility in monotone polygons. The International Journal of Robotics Research p. 0278364913498291 (2013)
23. Sato, H., Royset, J.O.: Path optimization for the resource-constrained searcher. Naval Research Logistics 57(5), 422–440 (2010)
24. Spieser, K., Frazzoli, E.: The cow-path game: A competitive vehicle routing problem. In: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on. pp. 6513–6520 (Dec 2012)
25. Spires, S.V., Goldsmith, S.Y.: Exhaustive geographic search with mobile robots along space-filling curves. In: Collective robotics, pp. 1–12. Springer (1998)
26. Sujit, P.B., Ghose, D.: Multiple agent search of an unknown environment using game theoretical models. In: American Control Conf. vol. 6, pp. 5564–5569 vol.6 (June 2004)
27. Sujit, P.B., Ghose, D.: Negotiation schemes for multi-agent cooperative search. Proc. of the Institution of Mech. Engineers, Part G: J. of Aero. Eng. 223(6), 791–813 (2009)
28. Sydney, N., Paley, D.A., Sofge, D.: Physics-inspired motion planning for information-theoretic target detection using multiple aerial robots. Auton. Robots pp. 1–11 (2015)
29. Trummel, K., Weisinger, J.: Technical notethe complexity of the optimal searcher path problem. Operations Research 34(2), 324–327 (1986)
30. Vincent, P., Rubin, I.: A framework and analysis for cooperative search using UAV swarms. In: ACM Symposium on Applied Computing. pp. 79–86. SAC '04, ACM, New York, NY, USA (2004), `http://doi.acm.org/10.1145/967900.967919`
31. Waharte, S., Trigoni, N.: Supporting search and rescue operations with UAVs. In: Int. Conf. on Emerging Security Technologies. pp. 142–147 (Sept 2010)
32. Zhu, M., Frazzoli, E.: On competitive search games for multiple vehicles. In: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on. pp. 5798–5803 (Dec 2012)