Noname manuscript No.
(will be inserted by the editor)

Competitive Target Search with Multi-Agent Teams: Symmetric and Asymmetric Communication Constraints

Michael Otte · Michael Kuhlman · Donald Sofge

Received: date / Accepted: date

Abstract We study a search game in which two multi-agent 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 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.

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.

Michael Otte and Donald Sofge
U.S. Naval Research Laboratory
Washington D.C. 20375
Tel.: +202-767-0806
Fax: +202-767-0806
E-mail: michael.otte.ctr@nrl.navy.mil
E-mail: donald.sofge@nrl.navy.mil

Michael Kuhlman
University of Maryland
College Park, MD 20742
E-mail: mkuhlman@umd.edu

Keywords Multi-Agent 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 coordination.

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?

Related work is discussed in Section 2. Nomenclature, a formal statement of assumptions, and the formal game definition appear in Section 3. 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

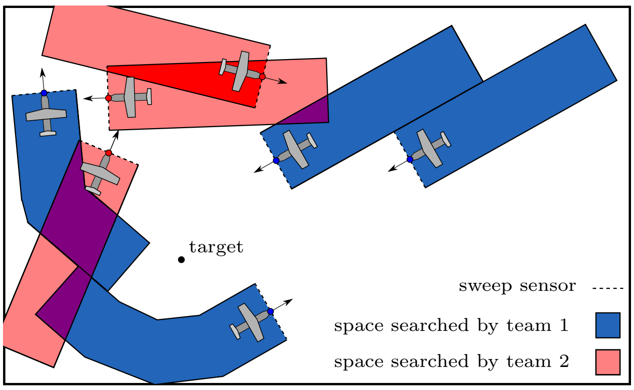

Fig. 1: Four agents (dark blue) compete against three adversaries (light red) to locate a target (black point). Communication enables members of the same team to cooperate (e.g., blue team), lack of communication prohibits cooperation (e.g., red team). The sweep sensors are, by definition, infinitesimally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

equilibrium exists at the point that each team randomizes their distributed searches such that all points are swept at the same expected time. In Section 5 we investigate a game played on $\mathbb{S}^1$ and show that it is ideal. In general, the ideal version of the game provides a useful model that allows us to evaluate how coordination affects game outcome, but is impossible to realize in many environments due to a number of boundary conditions. In Section 6 we extend our results by bounding the performance of non-ideal cases in subsets of $\mathbb{R}^D$ and in $\mathbb{T}^D$. We find that non-ideal games become asymptotically close to ideal as the size of the environment increases toward infinity. Supporting simulations and experimental results appear in Section 7; discussion and conclusions appear in Sections 8 and 9, respectively.

2 Related Work

The target search problem was formalized at least as early as 1956 by Koopman (1956), 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 multi-agent teams is too large to cover here. Extensive surveys of different formulations and approaches can be found in Chung et al. (2011), Waharte and Trigoni (2010), and Dias et al. (2006). Previous work on target search ranges from the purely theoretical to the applied, and has used tools across a variety of fields including: differential equations (Mangel, 1989), graph theory (Trummel and Weisinger, 1986), game theory (Sujit and Ghose, 2004), numerical methods (Bertuccelli and How, 2005; Forsmo et al., 2013), control theory (Flint et al., 2002; Hu et al., 2013), heuristic search (Sato and Royset, 2010), economics (Chandler and Pachter, 2001; Dias, 2004), and biology (Kim et al., 2013; Sydney et al., 2015).

One difference between the current paper and previous work is the scenario that we consider. We assume that an adversarial relationship exists between two different teams of searchers that compete to locate the same target. In contrast, in *cooperative search* a single team of agents tries to locate one or more targets (Vincent and Rubin, 2004; Beard and McLain, 2003) that may be stationary (Hu et al., 2013) or moving (Kim et al., 2013), and a key assumption is that all searchers cooperate with each other.

Pursuit-evasion games assume a different but related scenario in which one agent/team actively tries to avoid capture by another agent/team (Vidal et al., 2002; Gerkey et al., 2005; Noori and Isler, 2013; Kwak and Kim, 2014), leading to an adversarial relationship between the searchers(s) and the target(s). *Capture the flag* (Huang et al., 2015) 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.

Although we consider the general case of the competitive target search game played in subsets of $\mathbb{R}^D$, we also analyze a special 1-dimensional case in the 1-sphere $\mathbb{S}^1$. Our analysis of the game in $\mathbb{S}^1$ shares similarities with *linear search* (Demaine et al., 2006), and *cow path problems* (Zhu and Frazzoli, 2012; Spieser and Frazzoli, 2012). Differences include our extensions to higher dimensional spaces, and that (Demaine et al., 2006) considers an infinite search domain, while we consider a finite search domain. Our higher dimensional extensions build on coverage methods that use lawn-mower sweep patterns (Choset and Pignon, 1998). Our world model shares many of the same assumptions as Choset and Pignon (1998); in particular, an initial uniform prior distribution over target location and perfect sensors. Using sweep patterns for single agent coverage is studied by Choset and Pignon (1998), while Vincent and Rubin (2004) extends these ideas to a single multi-agent team searching for a moving and possibly evading target. Spires and Goldsmith (1998) use the idea of space filling curves to reduce the 2D search problem to a 1D problem.

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 (Beard and McLain, 2003), and/or the ability to communicate between agents was assumed to be dependent on distance (Sujit and Ghose, 2009; Hu et al., 2013), limited by bandwidth (Flint et al., 2002), adversaries (Bhattacharya et al., 2016), other constraints (Hollinger et al., 2015), or impossible (Feinerman et al., 2012).

A preliminary version of this work appeared as a conference paper at the International Workshop on the Algorithmic Foundations of Robotics (WAFR) in 2016 (Otte et al., 2016), and as non-archival supplementary material submitted along with that work. The current paper improves on the archival conference version by including:

– A formal analysis of non-ideal cases in subsets of Euclidean space $\mathbb{R}^D$ and toroid spaces $\mathbb{T}^D$ and proving that such games approach the ideal case, in the limit, as the size of the environment increases without bound.
– An in-depth consideration of the special $\mathbb{S}^1$ case that facilitates visualization of the space of mixed strategies and provides intuition.
– Analysis of the the non-ideal effects of different sweep-sensor shape in $D > 2$ dimensional space.

3 Preliminaries

The search space is denoted X. The multi-agent 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 multi-agent 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

$$S = X \times \varTheta$$

of a single actor includes position X and directional heading $\varTheta$. Let S_{g_i} represent the state space of the i-th agent. The product state space of the team is

$$\mathcal{S}_G = S_1 \times \ldots \times S_n.$$

A particular configuration of the team is denoted $\mathbf{s}_G$, where

$$\mathbf{s}_G \in \mathcal{S}_G.$$

Similarly, for the adversary

$$\mathbf{s}_A \in \mathcal{S}_A = S_1 \times \ldots \times S_m.$$

It is convenient to define the product space of *locations* for each team. Formally,

$$\mathbf{g} = (g_1, \dots, g_n) \in \mathcal{X}_G = X_1 \times \ldots \times X_n$$

and

$$\mathbf{a} = (a_1, \dots, a_m) \in \mathcal{X}_A = X_1 \times \ldots \times X_m$$

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

In general, we consider $X \subset \mathbb{R}^D$ a convex subset of D-dimensional Euclidean space, where $D \geq 2$ and X is bounded. We also consider the D-dimensional torus $X = \mathbb{T}^D$, and the special case when the space is a 1-dimensional sphere, $X = \mathbb{S}^1$. When $X \subset \mathbb{R}^D$ we assume it is "well behaved" such that it is convex, and has a boundary ∂X that can be decomposed into a finite number of locally Lipschitz continuous pieces. Our general formulation assumes X is a continuous[1] 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 initial locations drawn according to $\mathcal{D}_X$. We will usually assume that this sampling is independent and identically distributed (i.i.d.), $\mathcal{D}_X = \mathcal{D}_{X,\mathrm{iid}}$, and it will be clear from context when we use an alternative assumption.

The Lebesgue measure[2] in a D-dimensional space is denoted $\mathscr{L}_D(\cdot)$. Let $\varOmega_X$ and $\varOmega_S$ be the smallest σ-algebras over X and S, respectively. The (extension of) the Lebesgue measure in $\varOmega_S$ is $\mathscr{L}_{\varOmega_S}(\cdot)$. It follows from our assumptions that:

$$\mathbb{P}(\hat{X}) = \int_{\hat{X}} \mathcal{D}_X(x) = \frac{\mathscr{L}_D(\hat{X})}{\mathscr{L}_D(X)}$$

and

$$\mathbb{P}(\hat{S}) = \frac{\mathscr{L}_{\varOmega_S}(\hat{S})}{\mathscr{L}_{\varOmega_S}(S)}$$

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 $(X, \varOmega_X, \mathbb{P})$ and $(S, \varOmega_S, \mathbb{P})$, 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

[1] A discrete formulation can easily 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.

[2] Note that in $\mathbb{T}^D$ this is technically the push-forward extension of the Lebesgue measure that one would naturally assume.

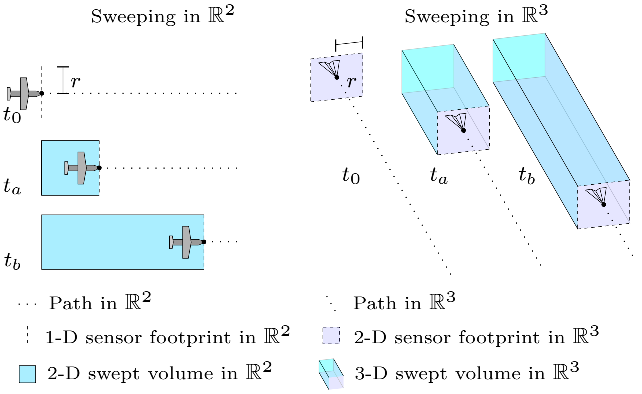

Fig. 2: Examples of sweep sensing in $\mathbb{R}^2$ (Left) and $\mathbb{R}^3$ (Right) for three different times $t_0 < t_a < t_b$. In $\mathbb{R}^2$ the sweep sensor footprint is a 1-D line segment oriented perpendicular to the direction of travel, in $\mathbb{R}^3$ it is a 2-D patch oriented perpendicular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

is much greater than ("$\gg$") r. In "large" search spaces the sweep sensor provides a reasonable idealization of any sensor with finite observation volume[3].

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[4]. We assume that each team is rational and will eventually sweep the entire space.

3.2 Paths, Multipaths, 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 multipath ψ_G is a set of paths containing one path ρ_i per agent in the team G,

$$\psi_G = \bigcup_{g_i \in G} \{\rho_i\}.$$

Let $\hat{\psi}_G$ denote a space covering search multipath. 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:

$$\psi_A = \bigcup_{a_j \in A} \{\rho_j\}$$

$$\bigcup_{s\in\mathbf{s}_A\in\hat{\psi}_A} B(s) \subset X.$$

Let Ψ_G (and Ψ_A) denote the space of all possible search multipaths 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, $\mathcal{C}(G) \times \mathcal{C}(A) = \{0, 1\} \times \{0, 1\}$.

In a (not always realizable) case of perfect coordination, team members 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.

When $X = \mathbb{S}^1$ perfect coordination is only possible if agents begin the game equally spaced around $\mathbb{S}^1$, which is an event with measure 0 if $\mathbf{g}_0$ is drawn according to $\mathcal{D}_{X,\text{iid}}$.

[3]Note that even for sensors with positive measure footprints, $0 < \mathscr{L}_{D-1}(B_r) < \infty$ (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).

[4]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).

In order to investigate perfect coordination when $X = \mathbb{S}^1$ we investigate a variation of the competitive search game in which the starting configuration of a communicating team is still drawn uniformly at random but constrained such that its members' starting locations are equally spaced (fanned-out) around $\mathbb{S}^1$. That is, the initial locations within a particular communicating team T are not i.i.d. (because there must be $\mathscr{L}_1(X)/|T|$ distance between nearby members), but they are distributed uniformly random over $\mathbb{S}^1$. In this case, we say that agent locations are drawn according to the distribution function $\mathcal{D}_{X,\text{fan}}$.

Drawing start locations from $\mathcal{D}_{X,\text{fan}}$ closely approximates "real" behavior when traveling to a particular point in the environment takes much less time than searching the entire environment.

3.4 Two-Team Target Search Game: General Formulation

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

Given a search space X, a stationary target $q \in X$, a multi-agent 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 multipaths $\psi_G \subset X$ 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, Multipath 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 multipath 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 $\mathcal{D}_{X,\text{iid}}$, or $\mathcal{D}_{X,\text{fan}}$, 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, and the starting locations for members* *of G and A are drawn uniformly at random, the probability team G wins is equal to the ratio of search space it sweeps before the adversary,*

$$\mathbb{P}\left(\omega_{\text{win}} | \psi_G, \psi_A\right) = \frac{\mathscr{L}_D(X_{\text{win}}(\psi_G, \psi_A))}{\mathscr{L}_D(X)}.$$

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 integrating $\mathcal{D}_{\mathbf{a}_0}(\psi_A)$ over $\Psi_A(\mathbf{a}_0)$ for all $\mathbf{a}_0$. Assuming the target and teams are distributed uniformly at random, this is:

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

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:

$$(\Psi_G^*, \Omega_G, \mathbb{P}_G^*) = \underset{((\Psi_G, \Omega_G, \mathbb{P}_G))}{\arg\max} \ \mathbb{P}\left(\omega_{\text{win}}\right).$$

For brevity we denote the expected value of '$\cdot$' over all $\mathcal{S}_G$, $\mathcal{S}_A$, Ψ_G^*, and Ψ_A^* as $\mathbb{E}^*[\cdot]$,

$$\mathbb{E}^*[\cdot] \equiv \mathbb{E}_{\mathcal{S}_G, \mathcal{S}_A, \Psi_G^*, \Psi_A^*}\left[\cdot\right].$$

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 cooperatively 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 $t_{\text{final}} = \min(t_{n,\text{sweep}}, t_{m,\text{sweep}})$ (i.e., for any realization of random start and target locations).

We assume both teams play a mixed strategy defined by a probability distribution over members of a set of strategies. 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 sub-optimal 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).*

We note that even if our assumption of uniform and i.i.d. starting locations were dropped, then an ideal strategy can still be achieved as long as the conditions in Corollary 2 are met.

In an *ideal game* each team plays an optimal mixed strategy over a set of ideal search strategies. (As we will show later, this becomes possible, in the limit, as the size of the environment increases). The following is true by the definition of a Nash equilibrium:

Proposition 4 *Assuming optimal ideal strategies exist for both teams, and the members of each team are able to communicate, 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.

In an ideal game this rate is constant, and integrating over time yields the probability of winning an ideal game in which both teams play optimal mixed strategies.

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \int_0^{t_{\text{final}}} \mathbb{E}^*\left[\frac{d}{dt}\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}\right] dt \tag{2}$$

Details of the derivation of Equation 2 can be found in Appendix A. An expression for the necessary expected value is also derived in Appendix A, and presented in the following Lemma 1.

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

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

The full proof of Lemma 1 appears in Appendix A.

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 $t_{\text{final}} = 1/(c^* \max(n,m))$:

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

Solving this equation yields the following corollary.

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*

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

4.2 Case 2: The Multi-Agent Team G Cannot Communicate but the Adversary Team A Can (Ideal Case)

The starting location of each agent on team G is sampled i.i.d. and uniformly at random. Therefore, a game in which team G cannot communicate but the adversary team A can communicate is equivalent to a situation in which the adversary team A plays n sub-games, one vs. each single-agent sub-team $\{g_i\} \subset G$. The adversary team A wins the overall game if and only if it wins each of the n sub-games. For the adversary to win the overall game, some $a \in A$ must sweep the target before each $\{g_i\} \subset G$. That said, we can restrict our focus to the particular adversary $a_j \in A$ that sweeps the target before all other adversaries $a_{k\neq j} \in A$. Assuming A plays an optimal ideal strategy, then:

- Each $a_j \in A$ sweeps only $1/m$ of the space.
- All $a_j \in A$ sweep mutually disjoint subsets of space.
- The total time required for A to collectively sweep the entire space is $1/(mc^*)$.
- Only a single adversary will sweep the target.

Combining these facts with uniformly random sampling of starting locations and target locations leads to the following theorem (see Appendix B for details):

Theorem 1 *The probability team G wins an ideal game, assuming team G cannot communicate but the adversary team A can communicate, and the adversary team A plays optimal ideal mixed strategies, while each $\{g_i\} \subset G$ plays an optimal ideal single agent mixed strategy, is:*

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

The full proof of Theorem 1 appears in Appendix B.

4.3 Case 3: Team G Can Communicate but the Adversary 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 4 *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:*

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

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 2 *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 decorrelated 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. □

5 Competitive Target Search in $X = \mathbb{S}^1$

We now apply the general results derived in Section 4 to a simple family of competitive target search games played on the 1-dimensional sphere. We design the game such that teams with communication are able to coordinate their starting locations by drawing them from $\mathcal{D}_{X,\text{fan}}$ (teams that cannot communicate use $\mathcal{D}_{X,\text{iid}}$). Start location coordination was not considered in the previous section, but is a practical necessity in $X = \mathbb{S}^1$ due to the topological constraint that moving from one location to another requires passing through all points between them[5]. This game is a reasonable model for scenarios in which searching space requires significantly more time than moving through it. It also allows us to visualize all possible game outcomes as a function of starting locations, which is useful for building intuition. Formally, the game we consider in this section is defined as follows:

Given a search space $X = \mathbb{S}^1$, a stationary target $q \in X$ drawn from $\mathcal{D}_X(q) = \mathcal{D}_{X,\text{iid}}(x)$, and a multi-agent team $G = \{g_1, \ldots, g_n\}$ with communication $\mathcal{C}(G) \in \{0, 1\}$ and initial location drawn from

$$\mathcal{D}_X(g_{i,0}) = \begin{cases} \mathcal{D}_{X,\text{iid}}(x) & \text{if } \mathcal{C}(G) = 0 \\ \mathcal{D}_{X,\text{fan}}(x) & \text{if } \mathcal{C}(G) = 1 \end{cases}$$

and resulting movement chosen along multipath $\psi_G \subset X$, and an adversarial team $A = \{a_1, \ldots, a_m\}$ with communication $\mathcal{C}(A) \in \{0, 1\}$ and initial location drawn from

$$\mathcal{D}_X(a_{j,0}) = \begin{cases} \mathcal{D}_{X,\text{iid}}(x) & \text{if } \mathcal{C}(A) = 0 \\ \mathcal{D}_{X,\text{fan}}(x) & \text{if } \mathcal{C}(A) = 1 \end{cases}$$

and resulting movement chosen along multipath $\psi_A \subset X$; team G if $q \in B(g_i)$ for some $g_i \in \mathbf{s}_G \in \psi_G$ before $q \in B(a_j)$ for some $a_j \in \mathbf{s}_A \in \psi_a$, and otherwise the adversary team A wins.

The symmetry of $\mathbb{S}^1$ combined with the uniformly random positioning of actors and target means that there is an equivalence between strategies such that, with an abuse of notation:

$$\Psi_G(\mathbf{g}_0 + x) = \Psi_G(\mathbf{g}_0) + x. \tag{3}$$

In other words, the set of possible strategies for a team starting at $\mathbf{g}_0$ is identical to the set starting at $\mathbf{g}_0 + x$, except that each component of the multipath is translated by x. Given any optimal conditional mixed strategy $(\Psi^*_G(\mathbf{g}_0), \Omega_G(\mathbf{g}_0), \mathbb{P}^*_{\mathbf{g}_0})$ for some $\mathbf{g}_0$, we can construct an optimal conditional mixed strategy $(\Psi^*_G(\mathbf{g}_0 + x), \Omega_G(\mathbf{g}_0 + x), \mathbb{P}^*_{\mathbf{g}_0+x})$ for any $\mathbf{g}_0 + x$ as:

$$\mathcal{D}^*_{\mathbf{g}_0+x}(\psi_G) = \mathcal{D}^*_{\mathbf{g}_0}(\psi_G + x). \tag{4}$$

Similarly, given any $(\Psi^*_A(\mathbf{a}_0), \Omega_A(\mathbf{a}_0), \mathbb{P}^*_{\mathbf{a}_0})$ for some $\mathbf{a}_0$, the adversary's $(\Psi^*_A(\mathbf{a}_0 + x), \Omega_A(\mathbf{a}_0 + x), \mathbb{P}^*_{\mathbf{a}_0+x})$ can be constructed for any $\mathbf{a}_0 + x$ as:

$$\mathcal{D}^*_{\mathbf{a}_0+x}(\psi_A) = \mathcal{D}^*_{\mathbf{a}_0}(\psi_A + x). \tag{5}$$

Each actor may either move in the forward or backward direction around $\mathbb{S}^1$. Given our formulation and assumptions there is no incentive to back-track once an agent has started to move in a particular direction. Thus, for any $\mathbf{g}_0$ we can assume $\mathcal{D}_{\mathbf{g}_0}(\psi_G) = 0$ for all ψ_G that involve backtracking (and similar for adversaries).

In an uncoordinated team, each agent independently chooses to move forward or backward. Each agent is ignorant of its teammates starting locations, and so it must travel around the entire space to guarantee that all of $\mathbb{S}^1$ is swept. This yields two useful paths (sub-strategies) for each agent. Let $\rho^+_{i,0}$ denote a path where the i-th agent starts at $g_{i,0}$ then moves in the positive direction around $\mathbb{S}^1$ until it again reaches $g_{i,0}$. Let $\rho^-_{i,0}$ denote the opposite path, where the i-th agent starts at $g_{i,0}$ and then moves in the reverse direction until again reaching $g_{i,0}$. All strategies for the team have the form $\bigcup_{g_i \in G}\{\rho_i \in \{\rho^+_{i,0}\ ,\ \rho^-_{i,0}\}\}$ and so the strategy space for an uncoordinated team starting at $\mathbf{g}_0$ is thus:

$$\Psi^{\text{un}}_G(\mathbf{g}_0) = \left\{ \psi_G \mid \psi_G = \bigcup_{g_i \in G} \{\rho_{i,0} \in \{\rho^+_{i,0}\ ,\ \rho^-_{i,0}\}\} \right\}.$$

[5] This formulation is chosen because we are interested in scenarios where communication provides a significant coordination advantage. The topological constraints of $\mathbb{S}^1$ significantly penalize an in-game redistribution from an initial i.i.d. space to an even spacing (that would otherwise be expected to facilitate maximum coordinated search). Although we do not explore it in this paper, an alternative game formulation in $X = \mathbb{S}^1$ would be to have all teams draw their start locations from $\mathcal{D}_{X,\text{fan}}$, regardless of their ability to communicate. In such a case, teams with communication can coordinate by having all members search in the directions expected to be most advantageous given their start locations.

Ideal coordination allows a team to divide the work such that each agent only needs to cover $1/n$ of $\mathbb{S}^1$. All agents start equally spaced around $\mathbb{S}^1$ and move in the same direction (either forward or backward). Given $\mathbf{g}_0$, the strategy space $\Psi_G^{\mathrm{co}}(\mathbf{g}_0)$ available to the team contains two strategies:

$$\Psi_G^{\mathrm{co}}(\mathbf{g}_0) = \{\psi_G^+(\mathbf{g}_0), \psi_G^-(\mathbf{g}_0)\}$$

where $L = \mathscr{L}_1(\mathbb{S}^1)$, and $\psi_G^+(\mathbf{g}_0)$ is the strategy where all agents move forward by distance L/n and $\psi_G^-(\mathbf{g}_0)$ is the strategy where all agents move backward by distance L/n.

The adversary has uncoordinated $\Psi_A^{\mathrm{un}}(\mathbf{a}_0)$ and coordinated $\Psi_A^{\mathrm{co}}(\mathbf{a}_0)$ strategy sets that are analogous to $\Psi_G^{\mathrm{un}}(\mathbf{g}_0)$ and $\Psi_G^{\mathrm{co}}(\mathbf{g}_0)$, respectively.

5.1 Case 1: Both Teams Can Communicate

The communication model is $\mathcal{C}(G) \times \mathcal{C}(A) = 1 \times 1$ and the initial placement of agent and adversaries are drawn according to the distributions $\mathcal{D}_X(g_{i,0}) = \mathcal{D}_{X,\mathrm{fan}}(g_{i,0})$ and $\mathcal{D}_X(a_{j,0}) = \mathcal{D}_{X,\mathrm{fan}}(a_{j,0})$, respectively. The sets of strategies available to the agent and adversary are $\Psi_G^{\mathrm{co}}(\mathbf{g}_0)$ and $\Psi_A^{\mathrm{co}}(\mathbf{a}_0)$. Both team G and the adversary team A are able to coordinate and distribute evenly (but at random offset) around $\mathbb{S}^1$.

For any particular combination of $\mathbf{g}_0$ and $\mathbf{a}_0$ there are only two strategies we need to consider for each team: (1) all agents move forward, (2) all agents move backward. Their combination defines a 2 by 2 space of game possibilities,

$$\Psi_G^{\mathrm{co}}(\mathbf{g}_0) \times \Psi_A^{\mathrm{co}}(\mathbf{a}_0) = \{\psi_G^+(\mathbf{g}_0), \psi_G^-(\mathbf{g}_0)\} \times \{\psi_A^+(\mathbf{a}_0), \psi_A^-(\mathbf{a}_0)\}.$$

Equation 1 simplifies to:

$$\mathbb{P}(\omega_{\mathrm{win}}) = \frac{1}{\mathscr{L}_{\Omega_{S_G}}(\mathcal{S}_G)\mathscr{L}_{\Omega_{S_A}}(\mathcal{S}_A)} \int_{\mathcal{S}_G}\int_{\mathcal{S}_A} \sum_{\psi_G \in \Psi_G^{\mathrm{co}}(\mathbf{g}_0)} \sum_{\psi_A \in \Psi_A^{\mathrm{co}}(\mathbf{a}_0)} \mathcal{D}_{\mathbf{g}_0}(\psi_G)\mathcal{D}_{\mathbf{a}_0}(\psi_A) \frac{\mathscr{L}_D(X_{\mathrm{win}}(\psi_G,\psi_A))}{\mathscr{L}_D(X)}$$

The symmetry of mixed strategies vs. $\mathbb{S}^1$ (equations 3-5) means that the potential benefits of using either $\psi_G^+(\mathbf{g}_0)$ (moving forward) or $\psi_G^-(\mathbf{g}_0)$ (moving backward) are independent of $\mathbf{g}_0$. An optimal strategy exists such that for all $\mathbf{g}_0$ and $\mathbf{a}_0$:

$$\mathcal{D}_{\mathbf{g}_0}(\psi_G^+(\mathbf{g}_0)) = c_{G,+} = \mathcal{D}^*(\psi_G^+)$$

$$\mathcal{D}_{\mathbf{g}_0}(\psi_G^-(\mathbf{g}_0)) = c_{G,-} = \mathcal{D}^*(\psi_G^-)$$

for constants $c_{G,+}$ and $c_{G,-}$. Similarly:

$$\mathcal{D}_{\mathbf{a}_0}(\psi_A^+(\mathbf{a}_0)) = c_{A,+} = \mathcal{D}^*(\psi_A^+)$$

$$\mathcal{D}_{\mathbf{a}_0}(\psi_A^-(\mathbf{a}_0)) = c_{A,-} = \mathcal{D}^*(\psi_A^-)$$

for constants $c_{A,+}$ and $c_{A,-}$ for the adversary. This, combined with the fact that the integral is Lebesgue, and Tonelli's Theorem (and the fact that all quantities are finite and nonnegative) allows us to reorganize our equation for $\mathbb{P}(\omega_{\mathrm{win}})$ as follows:

$$\mathbb{P}(\omega_{\mathrm{win}}) = \sum_{\psi_G \in \{\psi_G^+,\psi_G^-\}} \sum_{\psi_A \in \{\psi_A^+,\psi_A^-\}} \frac{\mathcal{D}^*(\psi_G)\mathcal{D}^*(\psi_A)}{\mathscr{L}_{\Omega_{S_G}}(\mathcal{S}_G)\mathscr{L}_{\Omega_{S_A}}(\mathcal{S}_A)} \int_{\mathbf{g}_0 \in S}\int_{\mathbf{a}_0 \in S} \frac{\mathscr{L}_D(X_{\mathrm{win}}(\psi_G,\psi_A))}{\mathscr{L}_D(X)}$$

From this point on the analysis is most naturally accomplished by considering the expected outcome of a smaller team vs. a larger or equal sized team. In general, either G or A may be the smaller team. However, our derivation assumes that G is the smaller team – with the understanding that the two team's quantities can be swapped to calculate the reverse situation. (Using Tonelli's Theorem, and the fact that all quantities are finite and nonnegative, the integrals can be arranged such that the larger team's integral is on the outside).

The rotational symmetry of $\mathbb{S}^1$ allows us to shift the coordinate system such that the first adversary starts at the origin, without loss of generality. We also rescale the problem such that $\mathscr{L}_D(X) = 1$, without loss of generality. All actors are identical, thus we relabel adversaries 2 through m such that adversaries $a_1, a_2, a_3, \ldots, a_m$ start at $0, 1/m, 2/m, \ldots, (m-1)/m$, and the location $0 = 1$. Similarly, we can relabel agents $g_1, \ldots g_n$ such that agents are spaced $1/k$ distance one after the other around $\mathbb{S}^1$. Note that g_1 is located anywhere on $(0, 1]$ with equal probability given the sampling distribution $\mathcal{D}_{X,\mathrm{fan}}$. On $\mathbb{S}^1$ there are two possible orientations for each agent (forward and backward), but initial orientation has no bearing on the strategies available to each agent given the movement model we are assuming. Therefore, the problem is equivalent to a lower dimensional projection that: (1) ignores rotational components of S, (2) ignores different permutations of agent and adversary order within S, and (3) assumes a fixed adversary starting location in which the first adversary starts at the origin:

$$\mathbb{P}(\omega_{\mathrm{win}}) = \sum_{\psi_G \in \{\psi_G^+,\psi_G^-\}} \sum_{\psi_A \in \{\psi_A^+,\psi_A^-\}} \mathcal{D}^*(\psi_G)\mathcal{D}^*(\psi_A) \int_{g_1 \in X} \mathscr{L}_D(X_{\mathrm{win}}(\psi_G,\psi_A)).$$

We now derive closed-form solutions for each of the integrals that appear in the four terms of this equation. Each term represents a different combination of teams deciding to move forward or backward around $\mathbb{S}^1$.

Given any g_1 and $\psi_G \in \{\psi_G^+, \psi_G^-\}$ and $\psi_A \in \{\psi_A^+, \psi_A^-\}$, it is possible to evaluate $\mathscr{L}_D(X_{\mathrm{win}}(\psi_G,\psi_A))$. For any particular n, m, ψ_G, and ψ_A it is possible to draw an outcome diagram showing how $\mathscr{L}_D(X_{\mathrm{win}}(\psi_G,\psi_A))$ changes over any possible $g_1 \in X$, an example for the scenario where $n = 3$ and $m = 5$ appears in Figure 4. Note that, due to reflective symmetry of $\mathbb{S}^1$ we can reduce our consideration to two cases: (1) both teams move in the same direction and (2) the two teams move in opposite directions. For either case, the distribution

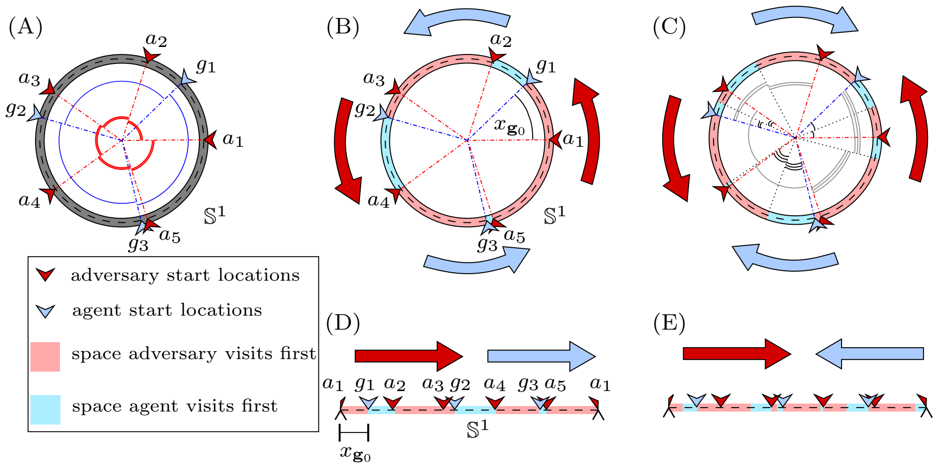

Fig. 3: Two teams, consisting of 3 agents (blue) and 5 adversaries (red), each start equally spaced around $\mathbb{S}^1$ (A). The relative offset between the two teams (first actors) is $x_{\mathbf{g}_0}$. Assuming the adversary moves in the positive direction around $\mathbb{S}^1$, (B) and (C) depict the cases where team G moves in the positive and negative direction around $\mathbb{S}^1$, respectively. Territory visited first by agents/adversaries shaded blue/red, respectively. (D) and (E) show the same scenarios as (B) and (C), respectively, but $\mathbb{S}^1$ is illustrated as a line segment instead of a circle.

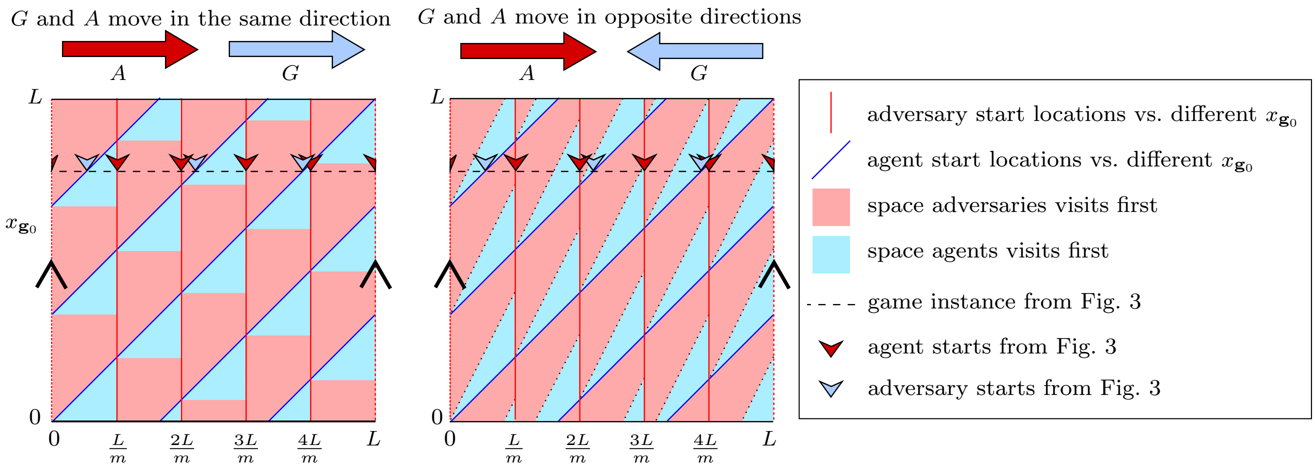

Fig. 4: Territory visited first (by agents/adversaries shaded blue/red, respectively) over the full space of ideal games that can be played between 3 agents and 5 adversaries, assuming both teams play an optimal strategy and the adversary moves in the positive direction around $\mathbb{S}^1$. Left and Right depict the cases where team G moves in the positive and negative direction around $\mathbb{S}^1$, respectively. Each point along the vertical axis represents a different relative offset between the two teams (first actors), i.e., $x_{\mathbf{g}_0} = g_1 - a_1$. For example, the case depicted in detail in Figure 3 is illustrated by the dashed black line. Red and blue lines show the starting locations of each actor vs. $x_{\mathbf{g}_0}$.

of first-reached territory changes vs. the relative offset between the two teams, which is equivalent to g_1.

Figure 3-B,D and Figure 4-Left depict the case where both teams move in the same direction. The territory that will be captured by any particular actor from the smaller team slowly decreases as its starting position approaches that of a member of the larger team. This will always create nm triangles in our outcome space (Figure 4-Left), each with an area of $1/(2m^2)$. Thus,

$\int_{g_1 \in X} \mathscr{L}_D(X_{\text{win}}(\psi_G^+, \psi_A^+)) = \int_{g_1 \in X} \mathscr{L}_D(X_{\text{win}}(\psi_G^-, \psi_A^-)) = \frac{n}{2m}$.

Similar analysis can be conducted for the case where the teams move in the opposite direction (Figure 3-C,E and Figure 4-Right). The division of territory is slightly more complex since a particular agent of the smaller team will take all territory between it and the midpoint between it and the approaching adversary's initial location, and possibly some territory beyond the first adversary — i.e., up to the point where it meets the next adversary. In Figure 4-Right the midpoints are represented by black dotted lines that appear on the boundary of the red and blue regions. This also creates triangular shaped regions that represent territory captured by the small team. Moreover, although the second group of triangles is shaped differently from the first, there is the same number as in the first case, and each triangle has the same area as in the first case. Thus,

$$\int_{g_1 \in X} \mathscr{L}_D(X_{\text{win}}(\psi_G^+, \psi_A^-)) = \frac{n}{2m}$$

and

$$\int_{g_1 \in X} \mathscr{L}_D(X_{\text{win}}(\psi_G^-, \psi_A^+)) = \frac{n}{2m}$$

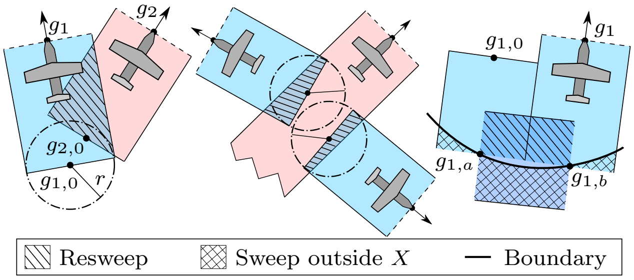

Fig. 5: Left: Two agents start within r of each other and have search paths that immediately overlap, causing some area to be swept by both of them (hashed). Center: Two agents start within $2r$ of each other but their search patterns do not immediately overlap; however, some area must be reswept by some agent later in the search. 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 so

$$\mathbb{P}(\omega_{\text{win}}) = \sum_{\psi_G \in \{\psi_G^+, \psi_G^-\}} \sum_{\psi_A \in \{\psi_A^+, \psi_A^-\}} \mathcal{D}^*(\psi_G)\mathcal{D}^*(\psi_A)\frac{n}{2m}$$

We observe that the mixed strategy distribution that either team chooses is inconsequential with respect to the expected outcome of the game. That is, we can swap probability density between $c_{G,+} = \mathcal{D}^*(\psi_G^+)$ and $c_{G,-} = \mathcal{D}^*(\psi_G^-)$ and it still holds that $\mathbb{P}(\omega_{\text{win}}) = n/(2m)$.

Combining these results with the alternative case where A is smaller than G gives:

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

and

$$\mathbb{P}(\omega_{\text{lose}}) = \begin{cases} 1 - n/(2m) & \text{when } n \leq m \\ m/(2n) & \text{when } n \geq m \end{cases}$$

This implies that, for the scenario considered in this section, an ideal optimal mixed strategy is always realized for the case of perfect coordination. It also means that the techniques used in Section 4 to extend case 1 to cases 2 and 3 (that is, assuming that the adversary must win n sub-games played vs. a single agent team) can also be used to extend these results to other communication situations in $\mathbb{S}^1$. Such extensions leverage the fact that $\mathcal{D}_{X,\text{fan}}$ and $\mathcal{D}_{X,\text{iid}}$ are equivalent for single agent teams.

6 Extensions to Non-Ideal Games, $\mathbb{T}^D$, and Subsets of $\mathbb{R}^D$

The realization of an ideal case requires that optimal mixed strategies exist such that Equation 2 holds. This idealization is possibly broken by:

I The startup locations of the robots.
II The environment.
III The sensor model (this only applies when $D > 2$).

In $\mathbb{T}^D$ and large convex subsets of $\mathbb{R}^D$ we are able to bound the departure from the ideal case. We find that for $D = 2$ the actual performance approaches the ideal case prediction as the sensor radius of the robot decreases toward 0 (or, in an equivalent interpretation, as the size of the environment increases without bound relative to the sensor radius). For $D > 2$ a similar result holds whenever the shape of the sensor footprint forms a tiling over the search space.

We can bound the individual effects of I, II, and III, respectively (i.e., assuming the others can be ignored). The general case where I, II, and III simultaneously occur can be found by combining these results. Full proofs are presented in Appendix C (in Appendices C.1, C.2, and C.3 we bound the individual effects of I, II, and III, respectively). However, now we summarize the major results now in the following Sections 6.1, 6.2, and 6.3, respectively.

6.1 Summary of Results Considering the Effects of Non-Ideal Start Locations

We now summarize the major results regarding the effects of non-ideal start locations. Let t_{startup} be an upper bound on the cumulative time the team spends performing non-ideal movement during a search. Let $\tilde{t}_{\text{startup}}$ be an analogous quantity for the adversary team. For convenience we also define

$$\hat{t}_{\text{final}} = \min(t_{\text{startup}} + \frac{1}{nc^*}, \frac{1}{mc^*})$$

and

$$\tilde{t}_{\text{final}} = \min(\tilde{t}_{\text{startup}} + \frac{1}{mc^*}, \frac{1}{nc^*}).$$

Theorem 3 *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, the probability we win is bounded as follows:*

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

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

The Full proof of Theorem 3 appears in Appendix C.1.

We observe that, by Theorem 3, in any $\mathbb{T}^D$ or subset of $\mathbb{R}^D$ we have the following limiting behavior:

$$\lim_{r \to 0} \frac{t_{\text{startup}}}{t_{\text{final}}} = 0.$$

That is, startup effects tend to vanish, in the limit, as sensor range shrinks toward 0. Or, in other words, relative startup inefficiencies vanish as environments get larger vs. sensor range.

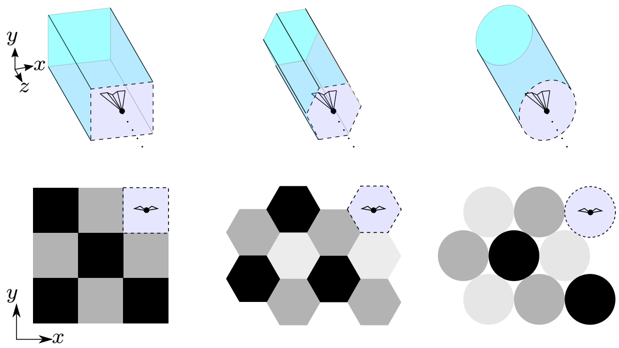

Fig. 6: Square and hexagonal sensors allow a tiling when projected along the axis of travel (along z) to the plane at $z = 0$, Left and Center. Discs do not permit a tiling (Right).

6.2 Summary of Results Considering Turns and Other Non-Ideal Boundary Effects

In this section we present bounds derived for the case of turns and other non-ideal boundary effects. t_{startup}, $\tilde{t}_{\text{startup}}$, $\hat{t}_{\text{final}}$, and $\tilde{t}_{\text{final}}$ have similar meanings as in Section 6.1; however, the non-ideal effects they account for in this subsection are related to boundary effects that occur near the edge of the search space (see Figure 5-Right for an example).

Theorem 4 *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 boundary conditions of re-sweep and sweep outside the search space, the probability we win is bounded as follows:*

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

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

The full proof of Theorem 4 is presented in Appendix C.2.

6.3 Summary of Results Considering Sensor Models when $D > 2$

In more than two dimensions it is possible to have sweep sensors that are non-tessellating (e.g., discs, see Figure 6-Right). The rate that *new* territory is swept depends on how much overlap is required with previous search passes. A detailed explanation is provided in Appendix C.3; but as a result, the sweep rate c^* in Equation 2 can no longer be assumed constant, and will change up to a finite[6] number of times K, where K is a function of the geometries of the sensor footprint and the search space and the numbers of agents and adversaries.

We break our analysis into K different intervals depending on the number of times that either team changes their sweep rate (changes happen whenever either team exhausts the amount of space that it can sweep at a particular rate). Let the k-th time interval starts at time t_k and end at time t_{k+1}. By the beginning of the k-th time interval the adversary has already swept F_k proportion of the total search space. During the k-th interval the rate team G sweeps new (for us) territory is determined by $nc_{G,k}$ and the rate the adversary sweeps new (for them) territory is determined by $mc_{A,k}$.

The total fraction of area that has been swept by the adversary prior to the start of the k-th interval is given by

$$F_k = 1 - \sum_{k=1}^{K-1} \int_{t_k}^{t_{k+1}} tmc_{A,k}\ dt$$

As a corollary of Lemma 1 we get:

Corollary 5 *the probability we win, assuming an optimal mixed strategy exists for both teams, and both teams play an optimal mixed strategy, when both teams communicate, and the game is ideal except for sensor footprint is*

$$\mathbb{P}(\omega_{\text{win}}) = \sum_{k=1}^{K-1} \int_{t_k}^{t_{k+1}} (F_k - tmc_{A,k})\, nc_{G,k}\ dt.$$

[6] To guarantee that this number is finite, we require that the sensor footprint contains some convex subset of space. In other words, degenerate sensors of measure zero or fractal-like geometry might produce an infinite number of different sweep rates.

The full analysis building up to Corollary 5 is presented in Appendix C.3.

These these type-III effects are a consequence of geometry. Unlike type-I and II sub-optimal effects, they are not reduced as the size of the environment is increased.

7 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. Simulations use a continuous space representation (in contrast to the discretized representation used in the simulations that appeared in a preliminary version of this work Otte et al. (2016)).

Multipaths are selected from a library of predefined sweep patterns, such that each pattern: forms a cycle, sweeps the entire space, and is designed to minimize sweep overlap between different parts of the search. An example of such a multipath library appears in Figure 9. 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. Next, each agent/adversary moves to its start point and searches along its allotted sub-path.

Agents are assigned to the multipath such that the maximum distance any particular agent needs to travel is approximately minimized (to within ϵ). In particular, we discretize the multipath into steps of length ϵ (where $\epsilon = 0.1$ km in our simulations involving 10 and 100 km width search spaces) and then solve the bottleneck bipartite graph matching for each of the $\hat{\ell}/(n\epsilon)$ possible matching vs. start locations using a modified version of Hopcroft and Karp (1971) and picking the best one. Figure 9 depicts the mean first-visit time for points in the environment, assuming a five agent team, for two different search areas both with and without communication.

Simulations are run in the Julia language using various sized search spaces. Selected results comparing predictions based on the ideal case vs. the average results from Monte Carlo simulations are presented in Figure 7.

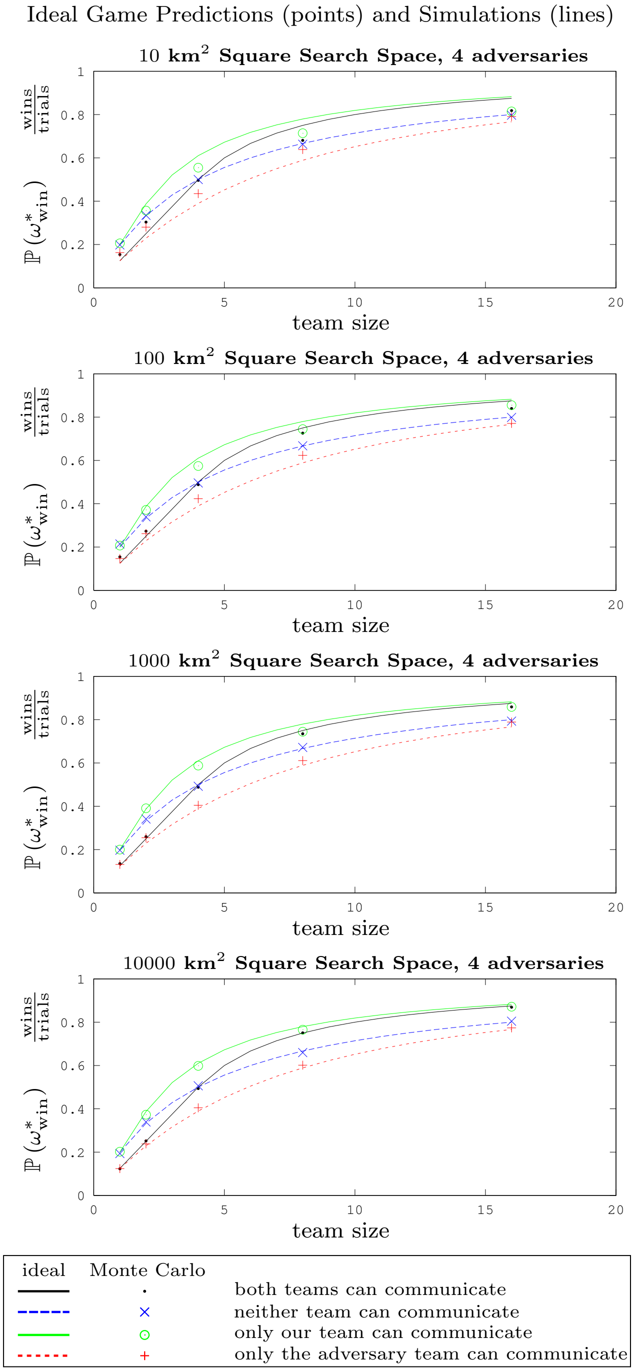

Fig. 7: The probability of winning an ideal game (lines) and observed win ratios given simulations (points) for various sized teams (horizontal axis) playing against 4 adversaries. Different line and point colors show the four various communication scenarios. Top to bottom sub-figures show small to large search spaces. **Note that the ideal predictions more accurately match results from Monte Carlo simulations as the size of the search space increases.**

The mixed platform combines Asctec Pelican quadrotor UAVs that have onboard Odroid single board computers with simulated agents that run on a laptop using the Ubuntu 14.04 operating system. The quadrotors receive position measurements from a Vicon motion capture system and runs ETH-Zurich modular sensor fusion framework for state estimation (Lynen et al.,

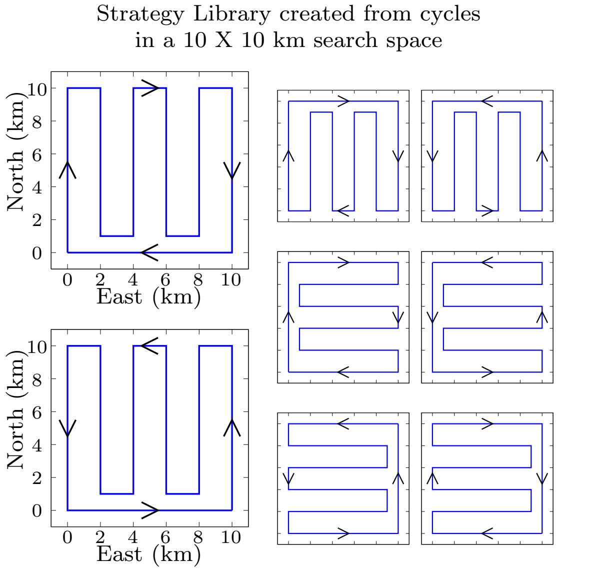

Fig. 8: An example sweep trajectory library with 8 different cycles in a 10 by 10 km space. Black arrows indicate the direction a cycle is traveled by the team. This library assumes robots have sweep sensors with radius 1 km.

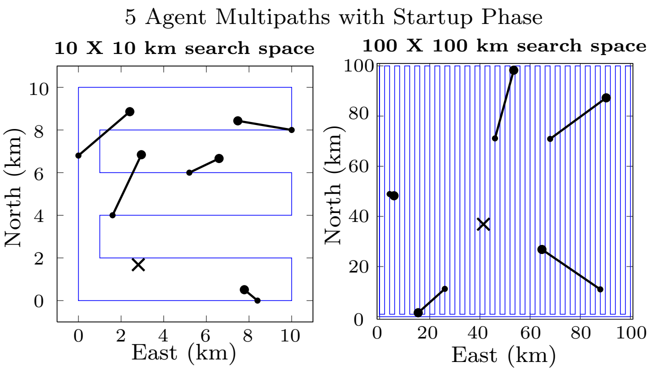

Fig. 9: 5 agent multipaths in 10 by 10 and 100 by 100 km environments (left and right, respectively). Robots start at the large black dots and then join the cyclic multipath (blue/gray lines) at equally spaced locations during the startup phase (black lines). The target location is marked with a black 'x'. Sweep sensor radius is 1 km. The lower left corners of the two search spaces are both located at $(0, 0)$. Agents travel clockwise in both examples.

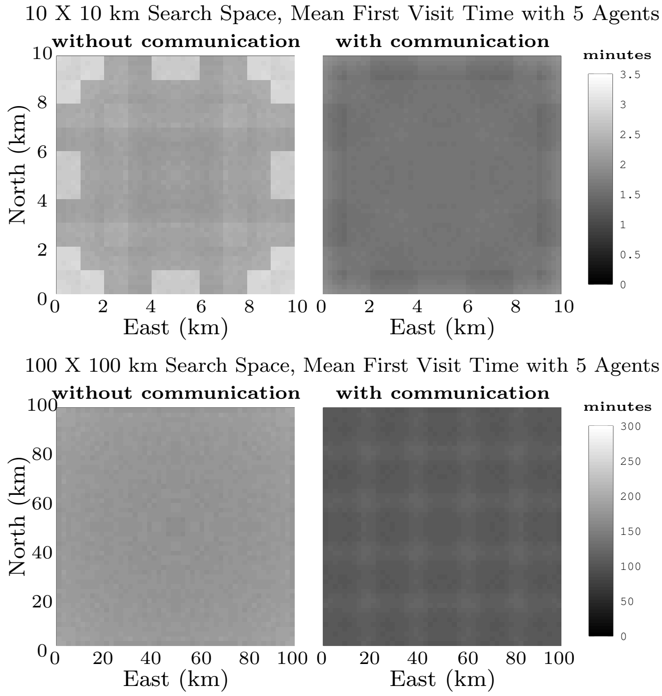

Fig. 10: First visit times of locations in the environment, averaged over repeated trials. A team of 5 agents, each with a 1 km radius sweep sensor, is deployed in a 10 by 10 km search area (top) and a 100 by 100 km search area (bottom). The speed of all agents is 1 km per minute. These figures were generated using Monte Carlo simulations in which the color at *each pixel* in the 50 X 50 pixel images represents the average over 500 random trials in which the target was placed at a point drawn uniformly at random within that pixel and each agent was placed at a point drawn uniformly at random in the entire search space. Each of the four sub-figures represents the outcomes of 1.25 million trials.

2013). 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 (Huang et al., 2010) for intervehicle communications.

The mixed platform experiments use a discrete grid environment where movement is allowed along the cardinal directions. 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 starting locations 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. An additional trial was aborted by our safety system when the distance between two actors (in this case the virtual agent and the adversary) fell below a safety threshold of 1 meter. Results from experiments with the mixed platform appear in Figure 11.

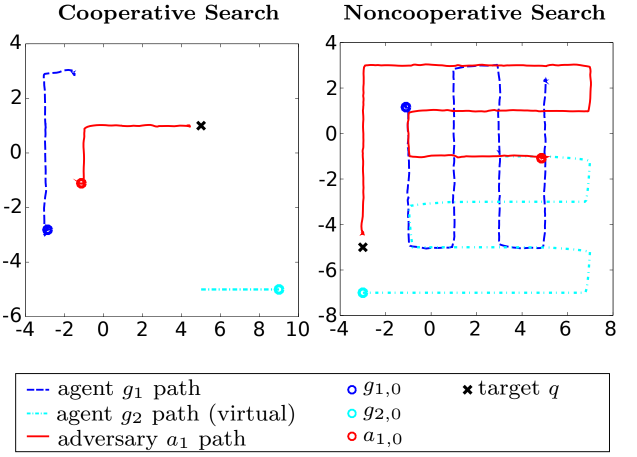

Fig. 11: Mixed platform experiments. A two agent team (consisting of a Pelican Quadrotor, blue, and a simulated agent, light-blue) vs. an adversary (Pelican Quadrotor, red) to find a target (black). Left and Right: Examples of paths when agents do and do not collaborate, respectively.

Table 1: Game outcomes of repeated trials.

Cooperation		No Cooperation	
Trial	Winner	Trial	Winner
1	G	1	A
2	G	2	G
3	G	3	A
4	A	4	G
5	A	5	A

8 Discussion

8.1 Using the Ideal Case as a Model for Non-Ideal Cases

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. This is easy to observe in Figures 7 and 10. In general, departure from the ideal case is due to the boundary issues involving bad startup locations, environmental geometry, and non-tiling sensor footprints. The bounds derived in Sections C.1 vs. C.2 suggest that the non-ideal effects cause by environmental geometry will usually be greater than those caused by conflicting start locations. Those in C.1 are proportional to nr, while those in C.2 are proportional to the size of the search space boundary. Thus, although both effects tend toward zero, in the limit, as the size of the environment increases, the relative importance of turns vs. start location conflicts increases without bound.

In Figure 7 we also observe that the average experimental performance of all communication scenarios appears to approach that of the "no teams communicate" scenario. This makes sense given that the outcome of the "no teams communicate" scenario is essentially determined by drawing a single actor uniformly at random from the set of all actors (agents and adversaries), and awarding the win to the team of the actor that is drawn. Likewise, if the environment is so small that the target will probably be found during the startup phase, then the outcome of the game is determined mostly by the actors' randomly determined starting locations than by any collaborative action enabled by communication.

8.2 Communication vs. Game Outcome

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 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.

8.3 Communication as a Prerequisite for Cooperation

Our assumption that communication is a prerequisite for cooperation only makes sense in scenarios in which a search was not anticipated or planned in advance, and in which case neither team has *a priori* information about the search space geometry, starting locations, or team size. Teams are created out of whatever agents happen to be nearby, e.g., performing unrelated missions.

If, on the other hand, all agents in the team had advanced knowledge that a search would be required within a particular space, cooperation could also be achieved by having the team agree on a search strategy *a priori*. This is known as using a "locker-room agreement." This would mitigate the effect of communication loss on team cooperation. However, it would also require that the team be equipped with a library of strategies (space covering multipaths for different sized search-spaces) and some way of determining which strategy to use on-the-fly (since the size of the search space, and each agent's initial position within it, would not be predictable).

8.4 Thoughts Regarding Competitive Target Search and the Stealth Assumption

The version of the two-team target search game that we consider assumes that neither team can observe the other, which is consistent with a stealth scenario. However, because the teams cannot observe each other, there is no back-and-forth interaction in which one team adopts its own strategy in response to moves made by the other team, etc. Instead, each team makes a single move that consists of selecting the multipath its agents will follow. One might reasonably ask "what are the interesting aspects of this game"? We now discuss what we find interesting about the scenario we consider.

Any known *pure* strategy can easily be thwarted by the adversary. Even though the adversary cannot directly observe our agents, if the adversary knows we are playing a particular pure strategy, then it can sweep search slightly ahead of $\min(n, m)$ of our agents — thwarting those $\min(n, m)$ agent's abilities to win. For this reason, any optimal strategy must be a mixed strategy; something that has major ramifications for the motion planning and coordination of our multi-agent team.

The expressions describing the probability of winning an ideal game are (perhaps unexpectedly) simple. This simplicity helps to provide intuition about the relative importance of team size and a team's ability to communicate. The bounds that we derive show that the ideal game itself has a convenient relationship to the non-ideal game: the ideal game exists as the limit of any non-ideal game, as the size of the environment increases relative to sensor radius. Thus, the intuition provided by the simple equations for the ideal game becomes more accurate as the size of the search space increases.

The particular stealth two-team target search game that we study belongs to a much larger family of competitive target search games. We hope that our formalization of this game will inspire other to consider this and other related scenarios.

8.5 A Note About Spiral Search Strategies

Spiral strategies work well when the search area is a disc (or close to a disc). If an environment is approximately disc-shaped, then a near-optimal spiral strategy is as follows: create a closed path by connecting the two arms of a double spiral (in particular, two equally spaced involutes of a circle) with short path segments, i.e., one near the edge of the search space and the other near the origin in the center of the spiral. The resulting closed path is homomorphic to a circle. An agent can move away from the center along one of the arms and return to the center along the other. Starting at any point along this closed path, a single agent will eventually sweep the entire search space and end up back where it started. A team of n agents begins the search by spacing out equally along the closed path (and with a random offset from the origin of the spiral); next, all agents move in the same direction around the closed path. In disc-like environments, this strategy is close-to-ideal—the strategy becomes ideal in the limit, as the size of the environment increases.

In environments that are *not* disc-like we believe that any spiral strategy will perform poorly for the following reasons: (1) If a single spiral is used (as was described above), then a large amount of area will be swept that is outside of the search space. (2) Alternatively, if multiple spirals are used to cover the space, then any covering will require significant overlap between different spirals (and some of the spirals will also sweep a large amount of area outside of the search space).

9 Summary and Conclusions

We study the effects of cooperation on multi-agent two-team competitive search games, a class of games in which two multi-agent 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.

We leverage this result to bound the expected game outcome in the presence of boundary effects including: sensor overlap at starting locations, turning at the search space boundary, and non-tiling sensor footprints (tiling is only an issue when $D > 2$). In general, the (non-tiling) transient boundary effects vanish, in the limit, as the size of the search space increases toward infinity.

When $D = 2$ (and also for tiling sensor footprints when $D > 2$) 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. We note this expectation is performed over the space of all multipaths starting at all possible starting configurations of the team (and with respect to a probability density function of our choosing that characterizes the likelihood a particular multipath is used given each possible starting location). A Nash equilibrium exists for an ideal game.

The chances of winning the search game increases vs. team size, and also increases 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 and experimental results on a mixed platform with 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 number N0001416WX01272, "Mobile Autonomous Navy Teams for Information Search and Surveillance (MANTISS)," and grant number N0001416WX01271, "Autonomous Multi-Agent Search and Rescue in Unpredictable Contested Environments (AMASR)." 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

Beard RW, McLain TW (2003) 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

Bertuccelli LF, How JP (2005) Robust uav search for environments with imprecise probability maps. In: Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC '05. 44th IEEE Conference on, pp 5680–5685, DOI 10.1109/CDC.2005.1583068

Bhattacharya S, Khanafer A, Başar T (2016) A Double-Sided Jamming Game with Resource Constraints, Springer International Publishing, pp 209–227

Chandler P, Pachter M (2001) Hierarchical control for autonomous teams. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, pp 632–642

Choset H, Pignon P (1998) Coverage path planning: The boustrophedon cellular decomposition. In: Field and Service Robotics, Springer, pp 203–209

Chung TH, Hollinger GA, Isler V (2011) Search and pursuit-evasion in mobile robotics. Autonomous robots 31(4):299–316

Demaine ED, Fekete SP, Gal S (2006) Online searching with turn cost. Theoretical Computer Science 361(2):342–355

Dias MB (2004) Traderbots: A new paradigm for robust and efficient multirobot coordination in dynamic environments. PhD thesis, Carnegie Mellon University Pittsburgh

Dias MB, Zlot R, Kalra N, Stentz A (2006) Market-based multirobot coordination: A survey and analysis. Proceedings of the IEEE 94(7):1257–1270

Feinerman O, Korman A, Lotker Z, Sereni JS (2012) Collaborative search on the plane without communication. In: Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing, ACM, New York, NY, USA, PODC '12, pp 77–86, DOI 10.1145/2332432.2332444, URL http://doi.acm.org/10.1145/2332432.2332444

Flint M, Polycarpou M, Fernandez-Gaucherand E (2002) 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

Forsmo EJ, Grotli EI, Fossen TI, Johansen TA (2013) Optimal search mission with unmanned aerial vehicles using mixed integer linear programming. In: Unmanned Aircraft Systems (ICUAS), 2013 International Conference on, pp 253–259, DOI 10.1109/ICUAS.2013.6564697

Gerkey BP, Thrun S, Gordon G (2005) Parallel stochastic hill-climbing with small teams. In: Multi-Robot Systems. From Swarms to Intelligent Automata Volume III, Springer, pp 65–77

Hollinger GA, Yerramalli S, Singh S, Mitra U, Sukhatme GS (2015) Distributed data fusion for multirobot search. Robotics, IEEE Transactions on 31(1):55–66

Hopcroft JE, Karp RM (1971) A n5/2 algorithm for maximum matchings in bipartite. In: Switching and Automata Theory, 1971., 12th Annual Symposium on, IEEE, pp 122–125

Hu J, Xie L, Lum KY, Xu J (2013) Multiagent information fusion and cooperative control in target search. IEEE Transactions on Control Systems Technology 21(4):1223–1235, DOI 10.1109/TCST.2012.2198650

Huang AS, Olson E, Moore DC (2010) Lcm: Lightweight communications and marshalling. In: Intelligent robots and systems (IROS), 2010 IEEE/RSJ international conference on, IEEE, pp 4057–4062

Huang H, Ding J, Zhang W, Tomlin CJ (2015) Automation-assisted capture-the-flag: A differential game approach. IEEE Transactions on Control Systems Technology 23(3):1014–1028

Kim MH, Baik H, Lee S (2013) Response threshold model based uav search planning and task allocation. Journal of Intelligent & Robotic Systems 75(3):625–640, DOI 10.1007/s10846-013-9887-6

Koopman B (1956) The theory of search. ii. target detection. Operations research 4(5):503–531

Kwak DJ, Kim HJ (2014) Policy improvements for probabilistic pursuit-evasion game. J Intell Robotics Syst 74(3-4):709–724, DOI 10.1007/s10846-013-9857-z, URL http://dx.doi.org/10.1007/s10846-013-9857-z

Lynen S, Achtelik MW, Weiss S, Chli M, Siegwart R (2013) 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

Mangel M (1989) Marcel Dekker, New York

Noori N, Isler V (2013) Lion and man with visibility in monotone polygons. The International Journal of Robotics Research p 0278364913498291

Otte M, Kuhlman M, Sofge D (2016) Competitive two team target search game with communication symmetry and asymmetry. In: International Workshop on the Algorithmic Foundations of Robotics (WAFR), San Francisco, USA

Sato H, Royset JO (2010) Path optimization for the resource-constrained searcher. Naval Research Logistics (NRL) 57(5):422–440

Spieser K, Frazzoli E (2012) The cow-path game: A competitive vehicle routing problem. In: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, pp 6513–6520

Spires SV, Goldsmith SY (1998) Exhaustive geographic search with mobile robots along space-filling curves. In: Collective robotics, Springer, pp 1–12

Sujit PB, Ghose D (2004) Multiple agent search of an unknown environment using game theoretical models. In: American Control Conference, 2004. Proceedings of the 2004, vol 6, pp 5564–5569 vol.6

Sujit PB, Ghose D (2009) Negotiation schemes for multi-agent cooperative search. Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering 223(6):791–813

Sydney N, Paley DA, Sofge D (2015) Physics-inspired motion planning for information-theoretic target detection using multiple aerial robots. Autonomous Robots pp 1–11

Trummel K, Weisinger J (1986) Technical notethe complexity of the optimal searcher path problem. Operations Research 34(2):324–327

Vidal R, Shakernia O, Kim HJ, Shim DH, Sastry S (2002) Probabilistic pursuit-evasion games: theory, implementation, and experimental evaluation. IEEE Transactions on Robotics and Automation 18(5):662–669, DOI 10.1109/TRA.2002.804040

Vincent P, Rubin I (2004) A framework and analysis for cooperative search using uav swarms. In: Proceedings of the 2004 ACM Symposium on Applied Computing, ACM, New York, NY, USA, SAC '04, pp 79–86, DOI 10.1145/967900.967919, URL http://doi.acm.org/10.1145/967900.967919

Waharte S, Trigoni N (2010) Supporting search and rescue operations with uavs. In: Emerging Security Technologies (EST), 2010 International Conference on, pp 142–147

Zhu M, Frazzoli E (2012) On competitive search games for multiple vehicles. In: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, pp 5798–5803, DOI 10.1109/CDC.2012.6426371

A Details of Proofs Leading up to Lemma 1

We now present technical details leading up to the proof of Lemma 1.

Given particular multipath strategies ψ_G and ψ_A for our team and the adversary, respectively, we can compute the ratio of space our team visits first by integrating the normalized rate new (to anybody) territory is swept by our team:

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

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

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

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:

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

We observe that the quantity inside the outermost integral describes the expected value of $\frac{d}{dt}[\frac{\mathscr{L}_D(X_{\mathrm{new}}(t))}{\mathscr{L}_D(X)}]$ over all $\mathcal{S}_G$, $\mathcal{S}_A$, Ψ^*_G, and Ψ^*_A. Recall that the expected value of '$\cdot$' over all $\mathcal{S}_G$, $\mathcal{S}_A$, Ψ^*_G, and Ψ^*_A as $\mathbb{E}^*[\cdot] \equiv \mathbb{E}_{\mathcal{S}_G,\mathcal{S}_A,\Psi^*_G,\Psi^*_A}[\cdot]$. Thus, formally,

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

and

$$\mathbb{P}\left(\omega^*_{\mathrm{win}}\right) = \int_0^{t_{\mathrm{final}}} \mathbb{E}^*\left[\frac{d}{dt}\frac{\mathscr{L}_D(X_{\mathrm{new}}(t))}{\mathscr{L}_D(X)}\right] dt$$

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

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

Proof At time t the adversary (operating according to its own ideal optimal strategy) has swept $tmv\frac{\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$ 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 $1 - tmv\frac{\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$ vs. $\frac{d^*}{dt}\frac{\mathscr{L}_D(X_{\text{new}}(t))}{\mathscr{L}_D(X)}$.

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

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

B Details of Analysis Leading to Theorem 1

We now present details of our analysis leading up to and including the proof of Theorem 1.

Proposition 5 *Given a uniform distribution of starting locations for all $a \in A$ and target locations q, and assuming an ideal mixed strategy is played by A, then the distribution of times at which that target-sweeping adversary a_j sweeps the target is uniform on the interval $[0, 1/(mc^*)]$.*

Let t_A be a realization of a uniformly random sample from $[0, 1/(mc^*)]$. Given our assumptions:

$$\mathbb{P}\left(t_A < t\right) = \frac{t}{1/(mc^*)}.$$

where $0 \le t \le 1/(mc^*)$. Note that the probability density of t_A on $[0, 1/(mc^*)]$ is mc^*, formally:

$$f_t(t) = \begin{cases} mc^* & \text{if } 0 \le t \le 1/(mc^*) \\ 0 & \text{otherwise} \end{cases}. \tag{6}$$

Team G does not communicate and can be viewed as a confederation of n independent single-agent sub-teams $\{g_1\}, \dots \{g_n\}$. Each subteam $\{g_i\}$ plays a single agent ideal mixed strategy such that g_i sweeps (new to g_i) space at the rate c^* and so g_i requires $1/c^*$ time to sweep the entire space by itself. This leads to the single-agent team counterpart to Proposition 5:

Proposition 6 *Given a uniform distribution of starting locations for g_i and target locations q, and assuming a single agent ideal mixed strategy is played by $\{g_i\}$, then the distribution of times at which agent g_i sweeps the target is evenly distributed on the interval $[0, 1/c^*]$.*

Let $t_{g,i}$ be a realization of a uniformly random sample from $[0, 1/c^*]$. Given our assumptions on start locations, and assuming a single-agent ideal mixed strategy is played by $\{g_i\}$, the probability g_i sweeps the target before time t is:

$$\mathbb{P}\left(t_{g,i} < t\right) = \frac{t}{1/c^*}$$

where $0 \le t \le 1/c^*$.

The target is swept by team G as soon as it is swept by *any* agent $g_i \in G$, thus team G essentially gets to draw n i.i.d. uniformly random samples $t_{g,1}, \dots t_{g,n}$ from $[0, 1/c^*]$ and play the best (smallest) of these vs. the adversary team A's single draw from $[0, 1/(mc^*)]$. Let $t_G = \min_i(t_{g,i})$ denote a realization of the smallest out of n values sampled uniformly at random and i.i.d. from $[0, 1/c^*]$.

The distribution of t_G can be determined directly using order statistics. The probability that at least one of the n (uncommunicating) subteams $\{g_1\}, \dots \{g_n\}$ sweeps the target before time t is:

$$\mathbb{P}\left(t_G < t\right) = 1 - \left(1 - \frac{t}{1/c^*}\right)^n \tag{7}$$

Team G wins whenever $t_G < t_A$. We observe that $\mathbb{P}\left(\omega^*_{\text{tie}}\right) = \mathbb{P}\left(t_G = t_A\right) = 0$ and given our assumptions:

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \mathbb{P}\left(t_G < t_A\right) = 1 - \mathbb{P}\left(t_G > t_A\right) = 1 - \mathbb{P}\left(\omega^*_{\text{lose}}\right).$$

Theorem 1 *The probability team G wins an ideal game, assuming team G cannot communicate but the adversary team A can communicate, and the adversary team A plays optimal ideal mixed strategies, while each $\{g_i\} \subset G$ plays an optimal ideal single agent mixed strategy, is:*

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

Proof We can compute $\mathbb{P}\left(\omega^*_{\text{win}}\right)$ using the Law of Total Probability:

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \int_{-\infty}^{\infty} \mathbb{P}\left(t_G < t_A | t_A = t\right) f_t(t) dt.$$

We note that $f_t(t) = 0$ for all $t < 0$ and for all $t > 1/(mc^*)$ because the game has not started yet and the adversary will have already won, respectively. Substituting Equation 6 and 7 we get:

$$\mathbb{P}\left(\omega^*_{\text{win}}\right) = \int_0^{1/(mc^*)} \left(1 - \left(1 - \frac{t_A}{1/c^*}\right)^n\right) mc^* dt$$

and performing the integration yields:

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

The final result is obtained using algebra. □

C Details of Proofs Pertaining to Non-Ideal Cases

In Sections C.1, C.2, and C.3 we bound the individual non-ideal effects of type I, II, and III, respectively that were discussed in Section 6. The each effect is analyzed assuming the others can be ignored. The general case where I, II, and III simultaneously occur can be found by combining these results. Here we focus on the case where both teams can communicate. Bounds on cases where one or the other team cannot communicate can be found by extending these results as was done for the ideal case[7] in Section 4.

[7] That is, using the fact that G winning a game with n communicating agents vs. m non-communicating adversaries is equivalent to G winning each of the m independent subgames involving each different adversary (and vice versa). This is how the results from Section 4.1 were extended in Sections 4.2 and 4.3.

Boundary effects caused by non-ideal startup locations and environmental geometry both cause the team to perform more than the ideal amount of sweeping (see Figure 5). We use the same basic proof technique for their analysis in Sections C.1 and C.2.

The analytical technique used in Sections C.1 and C.2 relies on the fact that search-hindering non-ideal effects are expected to be increasingly detrimental to G's probability of winning the game the earlier in the game they occur. This happens because the ratio $\frac{d\mathscr{L}_D(X_{\text{swept}})}{d\mathscr{L}_D(X_{\text{new}})}$ decreases vs. time. It is possible to construct a scenario that is even worse than a worst-case non-ideal search, in term's of team G's probability of winning the game, by: (first) 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; (second) Assuming the adversary team A sweeps at the ideal-game rate for the entire game. The relative length of the non-ideal startup phase in this modified scenario is guaranteed to be *worse than the worst-case*[8], and is bounded by a dimensionally dependent constant length of time.

A road-map of our analytical technique, including the construction of the worse than worst-case scenario we use, is now presented:

1. We break the space X into two non-overlapping sets, X_{ideal} and $X_{startup}$, depending on if the sweep search through it is "ideal" (it is swept at a time when the team is sweeping new space at the maximum rate).
   (a) Let all search space that is not reswept be combined in set X_{ideal}.
   (b) Let all search space that is ever reswept (plus, when relevant, all non-search space that is swept) be combined in set $X_{startup}$. Multiple copies of each reswept portion of space are included in $X_{startup}$ — if a subset of space is swept i different times, then i different copies of that subset are included. For the following discussion each copy is considered distinct such that each contributes its own volume to $\mathscr{L}_D(X_{startup})$. Thus, by construction, $\mathscr{L}_D(X_{startup}) + \mathscr{L}_D(X_{ideal}) \geq \mathscr{L}_D(X)$.
2. We consider (a worse than worst-case scenario) where the time costs, but not the target detection benefits, of sweeping $X_{startup}$ are incurred prior to performing an ideal search through X_{ideal}.
   (a) We design a multipath that is guaranteed to sweep $X_{startup}$, including all duplicate copies of reswept space, and derive an upper bound t_{startup} on the time required for the team to travel this multipath.
   (b) We assume team G begins the game by moving along a path sufficient to sweep $X_{startup}$ — but not actually performing search as it moves along this path (e.g., with its detection sensor turned off). For each duration of non-ideal behavior that occurs in a normal scenario, this essentially shifts an equivalent duration of non-ideal behavior to the beginning of the search without providing any target detection benefits.
   (c) However, after accounting for the penalty we receive for not searching before t_{startup}, we assume that search through X_{ideal} happens at the ideal rate. In other words, for the purposes of deriving performance bounds, we essentially pay an up-front performance penalty "with interest" to move each piece of non-ideal search such that it happens before t_{startup}.
   (d) Finally, we account for actually searching $X_{startup}$ (since we moved through it without searching before t_{startup}).
      i. Our original search would have already completed by this point; thus, we can assume that any rate of sweep for the second pass over unique elements of $X_{startup}$ and our scenario will still be worse than the original (and provide a valid performance bound).
      ii. Thus, it is permissible to assume the ideal search rate in this phase (for convenience) without destroying the worse than worst-case bound.

We now apply this technique in Sections C.1 and C.2.

[8] In other words, the scenario we consider is provably worse than a worst-case scenario. Thus, the bounds we derive are outside bounds on a worst-case scenario; and as a result, they are also outside bounds on the actual scenario. Note that we choose to use the worse than worst-case scenario because it is straightforward to analyze (unlike the worst-case scenario and the actual scenario).

C.1 Non-Ideal Starting Locations

If two robots on the same team start closer than $2r$, then either some nonzero measure subset of space will be swept by both of them (see Figure 5-Left), or some nonzero measure subset of space will be swept by them and during some other point in the search (see Figure 5-Center). Such an event occurs with nonzero probability given an assumption of uniform random i.i.d. start locations.

The ill-effects of non-ideal start locations can be bounded by considering the following worse than worst-case scenario: all n team members start at exactly the same point $\mathbf{g}_0 = (x_{1,0}, \ldots, x_{n,0})$, and then begin the game by moving (*without* actually searching) to the closest configuration $\mathbf{g}_0^*$ at which type-I effects would not have occurred if the robots had started at $\mathbf{g}_0^*$ in the first place (see Figure 13). We assume that the entire team waits to start searching until all robots have reached their coordinate of $\mathbf{g}_0^*$, and that this requires t_{startup} time.

The probability the adversary wins before t_{startup} is $\frac{t_{\text{startup}} m v \mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$. After, t_{startup} our expected search rate is that of the ideal case such that $\frac{d\mathscr{L}_D(X_{\text{new}}(t))}{dt}$ decreases vs. $\frac{d\mathscr{L}_D(X_{\text{swept}}(t))}{dt}$ by the usual factor of $1 - \frac{tmv\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$ after t_{startup}.

Recalling that $c^* = v\frac{\mathscr{L}_{D-1}(B_r)}{\mathscr{L}_D(X)}$, the resulting worse than worst-case bound on the probability we win the game is:

$$\mathbb{P}(\omega_{\text{win}}) \geq \left[\int_{t_{\text{startup}}}^{\hat{t}_{\text{final}}} (1 - tmc^*)\, nc^* \, dt\right] - t_{\text{startup}} mc^* \tag{8}$$

where

$$\hat{t}_{\text{final}} = \min(t_{\text{startup}} + \frac{1}{nc^*}, \frac{1}{mc^*}).$$

As discussed above, this bound accounts only for inefficiencies in search caused by non-ideal search locations.

In toroid $\mathbb{T}^D$ environments the furthest distance that any individual member of the team must move during the startup phase is $2r(n-1)$ (see Figure 13) and so $t_{\text{startup}} < 2r(n-1)/v$.

A similar bound exists for convex subsets of $\mathbb{R}^D$ as long as other boundary effects can be ignored, and $\mathcal{W} > 2rn$, where $\mathcal{W}$ is the maximum distance between any two points in X along a geodesic. In other words, $\mathcal{W}$ is the maximum width of the environment. We assume other boundary effects can be ignored, but address them directly in Sections C.2 and C.3. We now limit our consideration to $\mathcal{W} > 2rn$, i.e., wide environments. In thin environments, such that $\mathcal{W} \not> 2rn$, moving

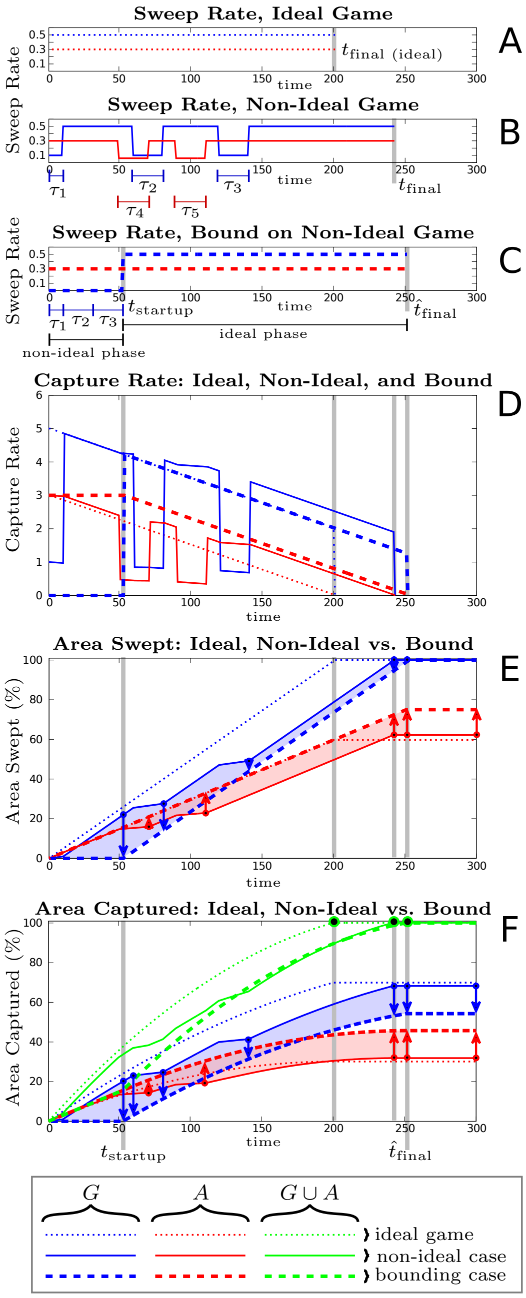

Fig. 12: Intuition for the worse than worst-case bound: an example in which $|G| = 5$ and $|A| = 5$. (A) Ideal sweep rates. (B) Sweep rates in a non-ideal case in which G has three periods of non-ideal sweep (τ_1, τ_2, τ_3) and A has two (τ_4, τ_5). (C) The corresponding worse than worst-case bound, in which periods of non-ideal sweep have been moved to the beginning of search for G and eliminated for A. (D) Territory capture rate for both teams in all three cases. (E) Total area swept in all three cases; by construction G sweeps less space in the bounding case than in the non-ideal case and A sweeps more (shaded regions and arrows show difference). (F) Total area captured in all three scenarios; by construction G captures less space in the bounding case than in the non-ideal case and A captures more (shaded regions and arrows show difference). Games end once all territory is captured (green).

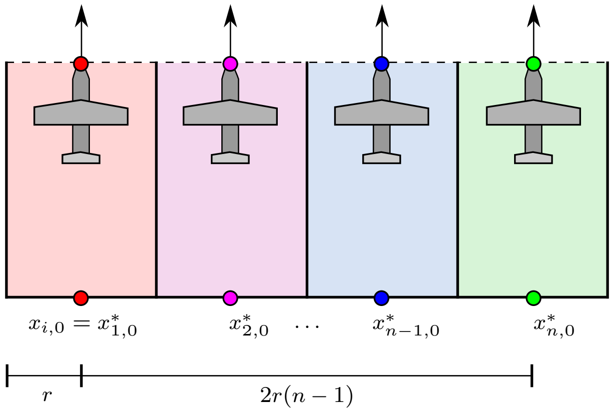

Fig. 13: An ideal strategy could be played if all robots start at $\mathbf{g}^*_0 = (x^*_{1,0}, \ldots, x^*_{n,0})$. Different colors represent area swept by different robots. In the worst-case, all robots all start at the same position, $\mathbf{g}_0 = (x_{1,0}, \ldots, x_{n,0}) = (x_{1,0}, \ldots, x_{1,0})$, and movement from $\mathbf{g}_0$ to $\mathbf{g}^*_0$ requires each robot to move no further than $2r(n-1)$, which can be accomplished in time $t_{\text{startup}} = 2r(n-1)/v$.

to an ideal start location may require time on the same order as sweeping the environment.

An upper bound on the probability we win is found by swapping the roles played by team G and the adversary. We let $\tilde{t}_{\text{startup}} < 2rm/v$ denote the startup time required for an adversary that must compete with an ideal version of team G (for achieving this other bound), and define

$$\tilde{t}_{\text{final}} = \min(\tilde{t}_{\text{startup}} + \frac{1}{mc^*}, \frac{1}{nc^*}).$$

The preceding discussion leads to the following theorem:

Theorem 3 *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, the probability we win is bounded as follows:*

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

and

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

C.2 Turns and Other Non-Ideal Boundary Effects

Re-sweeping of previously swept space occurs near the boundary of the search space due to the necessity of turning (Figure 5-Right). Moreover, sweeping all of the space within the search space sometimes requires sweeping some portion of space outside the search space (Figure 5-Right). Finally, a third boundary effect occurs due to the fact that, in general, the search spaces cannot be covered by an integer number of sweep passes (this also happens in $\mathbb{T}^D$, see Figure 14). All of these effects can be analyzed simultaneously.

We assume that the team starts at locations that do not suffer from the startup effects discussed in Section C.1. We

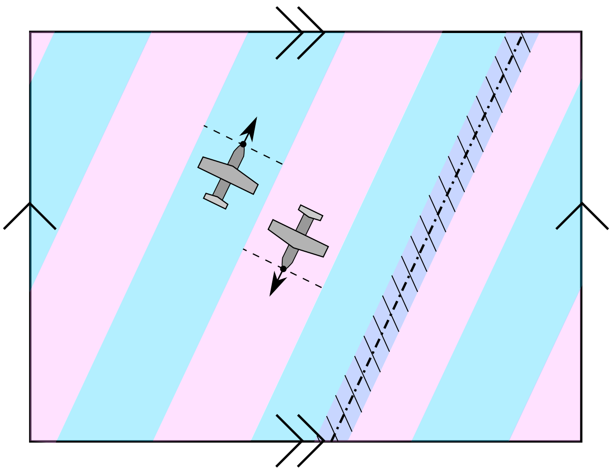

Fig. 14: Two robots perform a nearly ideal search over a $\mathbb{T}^2$ space. Different colors represent the different subspaces swept by either robot. The space cannot be divided into an integer number of non-overlapping wraps around the space and so some space is swept more than once (hashed). The manifold at the center of the reswept space is marked with a dot-dash line.

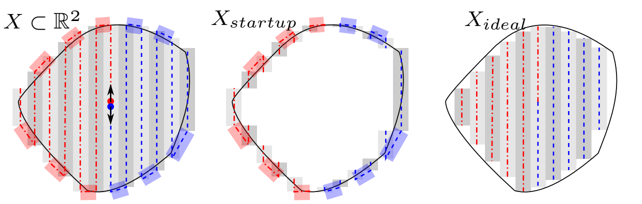

Fig. 15: (A) A nearly ideal two robot search strategy to cover the convex space. Red and blue represent different robots. Space boundary is black, red and blue dashed lines depict the multipath, shades of gray depict search sweeps, light blue and light red depict sweeps that re-sweep some area. The strategy can be broken into an non-ideal part over $X_{startup}$ and an ideal part over X_{ideal}.

also assume that the projection of the sensor footprint along the direction of movement permits a tiling over $D - 1$ space, see Figure 6 (we investigate the case where this is not true in Section C.3).

Given these assumptions, the space swept in any $\mathbb{T}^D$ search space or subset of $\mathbb{R}^D$ can be decomposed into two different non-overlapping subsets: the first ($X_{startup}$) contains all space involved in turning induced resweep as well as all non-search area that is swept, while the second (X_{ideal}) contains the remaining search space (that involved an ideal search). An example appears in Figure 15.

Similar to the analysis in Section C.1, we derive t_{startup} an upper bound on the time required to sweep $X_{startup}$ (including any necessary resweeps). And as in Section C.1, the worse than worst-case scenario we use for our analysis requires team G to wait for $X_{startup}$ time before beginning an ideal search (during which time the adversary searches at the ideal rate).

The first step to find t_{startup} is to bound the length of a path that is guaranteed to cover $X_{startup}$ (see Figure 16). Consider the largest hypercube that is contained within the robot's sensor footprint and that is aligned with the direction of forward movement. We call this hypercube $\hat{C}_r$. Let the boundary of the search space be denoted ∂X. Note, in toroid spaces the only boundary effects are related to the last pass, and so for $\mathbb{T}^D$, ∂X represents the manifold located between the paths taken during the first and last sweeps (See Figure 14).

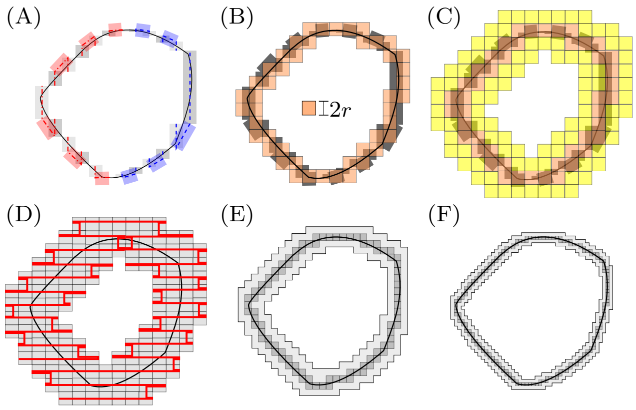

Fig. 16: It is possible to design a strategy for sweeping $X_{startup}$, the non-ideal part of a search. (A) The non ideal part of a search (red, blue, gray) and the search space boundary (black). (B) The search space boundary is covered with grid cells that have side lengths twice the minimum sensor radius. (C) Adding additional grids within r of the original grid cells is sufficient to cover $X_{startup}$. (D) We can design a strategy that sweeps the latter gridded space (there are N_{cover} grid cells) in time $< c_2 6 r N_{cover} \mathscr{L}_{D-1}(\partial X)$ where c_2 depends on the dimensionality of the search space. (E-F) as the radius of the sensor decreases relative to the size of search space the portion of time spent sweeping $X_{startup}$ decreases toward 0.

Since ∂X is essentially a lower-dimensional manifold embedded in our D-dimensional search space, it is possible to cover ∂X with N_{cover} hypercubes (each of dimension D), where

$$N_{cover} < c_2 r \mathscr{L}_{D-1}(\partial X)$$

here c_2 is a constant that counts the maximum number of tiled hypercubes required to cover a non-axis aligned line segment of length $\sqrt{D}$, and where $\sqrt{D}$ is the maximum distance between any two points in the same unit grid cell. Figure 16-A-B-C depicts such a covering.

It is important to note that c_2 is a dimensionally dependent constant. It is possible to construct a tour upper bounded by $\hat{\ell}$ that covers all grid cells within c_2 of ∂X (and thus covers $X_{startup}$). We note that $\hat{\ell}$ is also, by design, longer than the cumulative length traveled along the subset of the original multipath involved in the boundary effects we are investigating in this section. We calculate $\hat{\ell}$ by considering a naive tour of the hypercube covering of the search space boundary (see Figure 16-Bottom). For each hypercube, it is possible to construct such a covering by traveling a distance that is at most three times a side length (i.e., 6 times r_c) as follows: $2r_c$ to reach the center of the nearest face, $2r_c$ to reach the opposite face and thus sweep the entire cube, and $2r_c$ to exit through the center of any other face (in most cases it will be much less than this since multiple cubes can be swept without changing direction). Thus, $\hat{\ell} < 6 N_{cover} r_c$ and so

$$\hat{\ell} < 6 c_2 r N_{cover} \mathscr{L}_D(\partial X).$$

The time required to perform the startup phases is upper bounded by

$$t_{\text{startup}} < \hat{\ell}/v$$

This bound implicitly assumes a worst-case situation in which all boundary effects must be dealt with by a single agent. Thus, this bound on t_{startup} may be up to n times too large (i.e., in the best-case boundary problems are divided evenly between agents). We can now proceed as in the previous subsection (using our new definition of t_{startup} in Equation 8). This discussion leads to the following theorem:

Theorem 4 *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 boundary conditions of re-sweep and sweep outside the search space, the probability we win is bounded as follows:*

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

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

We note that $\frac{N_{cover} \mathscr{L}_D(\hat{C}_r)}{\mathscr{L}_D(X)} \rightarrow 0$, in the limit, as the size of the environment increases relative to the sensor radius.

The non-ideal effects from this section can be combined with those from the previous subsection by simply combining the startup phases used for analysis into a single startup phase of combined duration.

C.3 Sensor Model when $D > 2$

A sensor sweep footprint that forms a space tiling in $D-1$ dimensions is required for ideal search, See Figure 6. Non-tiling footprints require neighboring sweep passes to overlap in order to sweep the full space. When $D > 2$ the vast majority of sensor models will not permit a sweep footprint that forms a space tiling in $D-1$ dimensions. While an appropriately chosen sensor footprint, such as the $D-1$ dimensional L_∞-ball, does permit such a covering, other common symmetrical coverings such as the $D-1$ dimensional L_2-ball do not. We now investigate the effects of what happens when a non-tiling sensor is used.

Assume that, other than the tiling of the search sensor, the search is otherwise ideal (i.e., we are ignoring the startup and boundary effects that were addressed in Sections C.1 and C.2). Search necessarily happens in a number of different separate phases that are characterized by different $\frac{d\mathscr{L}_D(X_{\text{swept}})}{dt}$ (sweep rates of space we have not yet swept) and thus different $\frac{d\mathscr{L}_D(X_{\text{new}})}{dt}$ (sweep rates of space that has not been swept by ether team).

For example, we could perform search in two meta-phases represented by the gray discs and red circles in Figure 17. During the first phase, search happens at the ideal rate due to the fact that each pass (gray disc) covers terrain that we have never visited before. However, after some time, sweeping any new space will necessarily require re-sweeping some previously swept terrain (red circles). Thus, in the second phase, $\frac{d\mathscr{L}_D(X_{\text{swept}})}{dt}$ and $\frac{d\mathscr{L}_D(X_{\text{new}})}{dt}$ are substantially reduced.

We could alternatively sweep the entire space much more quickly using a multipath defined by sensor discs inscribed

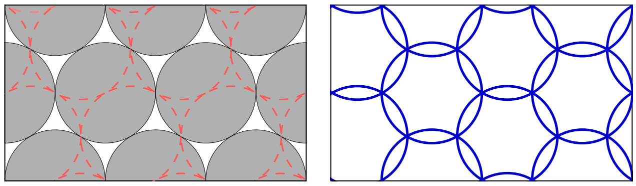

Fig. 17: Two different sweep patterns that use the same sensor footprint, projected along the direction of travel to remove depth (as in Figure 6-bottom). This particular (projected) *patch* could be tessellated arbitrarily many times in the horizontal and/or vertical directions. Left: Sensor footprints of the robot appear (gray, dashed red) for a two phase search strategy that covers the space using an ideal search with 12 passes per patch (6 gray, 6 red). Right: the same space can be swept using only 8 passes per patch (blue); although this requires sweeping the entire space more quickly, it requires searching below the ideal search rate after the first 1/3 of search.

by the blue hexagons in Figure 17. The price we pay for a quicker overall search turns out to be an earlier decrease in $\frac{d\mathscr{L}_D(X_{\text{swept}})}{dt}$ and $\frac{d\mathscr{L}_D(X_{\text{new}})}{dt}$.

The fact that we can control $\frac{d\mathscr{L}_D(X_{\text{swept}})}{dt}$ via choosing how different sweep passes overlap adds a significant amount of complexity to our analysis. Each time either team's search rate changes, we must use a slightly different version of Equation 2. This can be accomplished by breaking our analysis into K different intervals depending on the number of times that either team changes their sweep rate.

Let the k-th time interval start at time t_k and end at time t_{k+1}. By the beginning of the k-th time interval the adversary has already swept F_k proportion of the total search space. During the k-th interval the rate team G sweeps new (for us) territory is determined by $nc_{G,k}$ and the rate the adversary sweeps new (for them) territory is determined by $mc_{A,k}$.

The total fraction of area that has been swept by the adversary prior to the start of the k-th interval is given by

$$F_k = 1 - \sum_{k=1}^{K-1} \int_{t_k}^{t_{k+1}} tmc_{A,k} \; dt$$

As a corollary of Lemma 1 we get:

Corollary 5 *The probability we win, assuming an optimal mixed strategy exists for both teams, and both teams play an optimal mixed strategy, when both teams communicate, and the game is ideal except for sensor footprint is*

$$\mathbb{P}\left(\omega_{\text{win}}\right) = \sum_{k=1}^{K-1} \int_{t_k}^{t_{k+1}} \left(F_k - tmc_{A,k}\right) nc_{G,k} \; dt.$$

We note that these effects do not vanish as the size of the environment increases relative to the sensor footprint.

Corollary 5 works for the case that both teams communicate. In general, extending these results to cases where one team cannot communicate is more involved than how the ideal results from Section 4.1 were extended in Sections 4.2 and 4.3. Each of the sub-games vs. a single (i.e., non-communicating) agent must be analyzed separately to account for the fact that each agent will independently choose when and how much its own sweeps will overlap between different passes through the environment.