51st IEEE Conference on Decision and Control
December 10-13, 2012. Maui, Hawaii, USA

The Cow-Path Game: A Competitive Vehicle Routing Problem

Kevin Spieser and Emilio Frazzoli

Abstract— **This work considers multi-vehicle systems in which self-interested, mobile agents compete to capture a target that has been distributed on a ring. In the scenarios studied, agents face the added difficulty of having minimal sensing capabilities and limited knowledge of where the target is located. We consider strategic algorithms that allow agents to effectively make decisions and plan trajectories in these settings. Specifically, we characterize equilibria strategies for a search game in which two cows compete to find a patch of clover located somewhere on the unit ring. Throughout, we motivate the work using the example of taxi drivers that compete with one another to garner fares in a busy urban landscape.**

I. Introduction

Traditional one-sided search problems involve a mobile agent trying to find a target within a spatial environment. Based on agent attributes, e.g., sensing radius, detection efficiency, and kinematic constraints, as well as other problem specifics, e.g., workspace geometry and target mobility, the agent must devise a search plan that is optimal in some sense, e.g., minimizes the expected capture time. When multiple agents are employed, these plans are often designed in a coordinated manner with a common objective in mind. For example, cooperative strategies may facilitate smaller detection times as compared to protocols in which the agents search independently. Even in the game-theoretic problems considered, competitive tension emerges only between a hider, which can be viewed as a decision-making target, and the team of search agents, in that the search agents act cooperatively, with no preference for which agent, if any, ultimately finds the target.

While the multi-agent, search problems highlighted above address decision-making in a host of pertinent applications, they say little about how the agents should search when they, themselves, are in competition to secure scarce resources. In this paper, we study a new type of multi-agent search game that stresses precisely this type of inter-agent competition. We consider scenarios in which multiple self-interested, search agents compete to capture immobile targets that are distributed in an environment Q. Each search agent possesses a prior probability distribution for where the targets are located, but has limited sensing capabilities and can discover a target only when they are directly over it. The targets are viewed as artifacts of the environment, not as strategic decision-makers. Rather, the search agents are the only decision-makers and the game is played among them, with each agent trying to capture as many targets as possible. To emphasize the stark differences that exist between our framework and cooperative approaches to search, we refer to the problems we consider as *competitive search games*.

We believe the competitive search game framework captures the incentive for strategic decision-making, among mobile agents, at play in a host of applications. For example, in many cities, cab drivers make their living from "street pickups", whereby the cab driver is hailed by prospective passengers in real time [1]. This is in contrast to other operational models in which drivers may have some jobs forwarded to them via centralized call-in centers. In this example, Q is a graph with edges and vertices representing roadways and intersections, respectively. The targets are the potential passengers, which arrive on an ongoing basis according to a spatio-temporal process. The game is played between the cab drivers: each driver trying to maximize their revenue. To operate effectively, drivers must plan routes through the city by accounting for spatial demand patterns as well as other nearby drivers.

As a second example, consider two rival shipwreck-recovery boats searching the outskirts of a jagged coral reef for the remnants of a treasure ship lost at sea. In this case, Q is the subset of $\mathbb{R}^2$ representing the waters surrounding the reef. The lone target is the sunken ship. The game is played between the two recovery boats: each boat trying to discover the wreck, and any associated treasure, first. Given a priori knowledge of where the ship sank, perhaps from crude sonar images, historical maps, word of mouth accounts, etc., each boat must chart a course to search the coastal waters surrounding the reef, while factoring in the presence of a rival salvage boat that harbors similar ambitions.

The scenarios described above both fall within the competitive search game framework we have outlined, yet they differ markedly in terms of workspace geometry, the number of search vehicles, the target arrival process, and the time scales over which searching evolves. Surprisingly, although we believe there are an abundance of practical problems that fit naturally within the competitive search game framework, there have been few associated results, even in the case of just two search agents. It is in this direction that this paper seeks to make an initial contribution.

The remainder of the paper is structured as follows: Sections II and III provide the relevant background and review a well-known, single-agent, line search scenario called the Cow-Path Problem. Section IV considers a variant of this problem, in which the search takes place over a ring. Section V adds a second vehicle to the ring in a competitive

The authors are with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, 02139. {kspieser,frazzoli}@mit.edu
This work was partially supported by the Office of Naval Research under grant N00014-09-1-0751.

978-1-4673-2066-5/12/$31.00 ©2012 IEEE

search game we refer to as the Cow-Path Game. Theoretical results concerning the optimal strategies of searchers are provided in Sections VI and VII. Finally, Section VIII summarizes the prominent ideas of the paper and outlines directions of ongoing and future research.

II. Related Work

The search scenarios we study are closely associated with two fields 1.) search theory or one-sided search and 2.) search games [2]. With regard to the former, [3] has long served as the definitive reference for developing optimal search plans given constraints that limit the resources available for searching. In many cases, the search agent has a finite sensing radius and probability $p > 0$ of failing to detect a target inside its sensing zone. Consequently, a target may go undetected and search is discussed in probabilistic terms.

The Cow-Path Problem, an instance of which involves a nearsighted cow that, starting from the origin, travels at unit speed and wishes to discover a patch of clover, distributed along $\mathbb{R}$ according to a known distribution f, in minimal expected time is introduced in [4] and [5]. A canonical problem in the fields of online algorithms and probabilistic robotic path planning, it has received considerable attention in [6], [7], [8], [9], and [10], in which a variety of conditions an optimal search trajectory must satisfy have emerged. Recently, the problem has received renewed interest in [11], in which costs are associated with the cow turning.

In [12], the authors consider the problem of finding a target that intermittently emits a signal. The target can be located only when it emits a signal and is inside the sensing zone of a search vehicle. They study the expected search time for a team of vehicles to find the target using periodic search paths. By not employing a prior distribution, the authors effectively assume the target is distributed uniformly over Q. In contrast, in [13], a prior distribution over the target's location is assumed and persistent trajectories are designed such that a patrol vehicle, with limited sensing capabilities, can detect newly arrived incidents in minimal expected time.

Search games, in which a searcher tries to capture a hider in minimal time, while the hider tries to delay capture for as long as possible are considered in [14]. Given the competing objectives, optimal player strategies are best expressed, through the language and formalisms of game theory, using various notions of equilibria. Equilibrium strategies, which are typically highly dependent on the workspace geometry, are reported in [15] for games taking place in an assortment of environments, including line segments, specialized graphs, and compact regions of $\mathbb{R}^2$. In the same work, the authors consider team search games in which multiple search agents scour the environment in an effort to locate the hider. In this respect, the game is played between the team of agents and the lone hider, with individual agents having no preference for who ultimately succeeds in capturing the hider.

We remark that in the multi-agent scenarios mentioned, search agent trajectories are planned in a coordinated or cooperative manner, with consideration for optimizing a mutual objective. In the competitive settings we consider, the search agents must not only explore the workspace to locate targets, but must do so with an awareness for other nearby agents that harbor similar ambitions.

Finally, the competitive search games we consider are closed-loop, in the sense that each agent is, at all times, aware of the position of every other agent [16]. Concurrently, there has been interest in studying open-loop competitive search games, for which agents craft search strategies without knowing the position of any other agent [17].

III. Background

In this section, we present requisite ideas that find application in the analysis and discussion of later sections.

We begin by stating the version of the Cow-Path Problem, or CPP, that we are interested in.

Definition 3.1: ***(The Cow-Path Problem)*** A patch of clover, $\mathcal{T}$, is distributed on the real line according to a known density function $f:\ \mathbb{R} \rightarrow \mathbb{R}^+$. Starting from $x = 0$, a cow C, capable of moving at unit speed, wishes to discover $\mathcal{T}$ in minimal expected time. It is assumed C can change directions instantaneously, but can discover $\mathcal{T}$ only when she is directly over it.

The cow's search plan is a sequence of points, $\{x_i\} \in \mathbb{R}$ satisfying $\cdots < x_4 < x_2 < 0 < x_1 < x_3 < \cdots$, at which to turn around [5]. Intuitively, the points at which C turns are those which, given f, collectively minimize the expected amount of backtracking required before finding $\mathcal{T}$. Exact solutions to the CPP are known for only a handful of target distributions, e.g., rectangular, triangular, Gaussian, though, for general f, dynamic programming-based methods can be used to find solutions of arbitrary accuracy given a sufficiently fine discretization of $\mathbb{R}$ [15].

In our study of competitive search games, we will be interested in search strategies that prove efficient in the face of inter-agent competition. To this end, we make frequent use of the following equilibrium notion from game theory.

Definition 3.2: (***ε-Nash Equilibrium***)[18] Let $\mathcal{G}$ be a game with n players. Let S_i be the strategy set of player i and $S = S_1 \times \cdots \times S_n$ the set of strategy profiles. For $s \in S$, let s_i be the strategy played by player i in s and s_{-i} the strategy profile of all players other than i in s. Let $\mathcal{U}_i(s)$ be the utility of player i under s. A strategy profile $s^* \in S$ is an ε-Nash equilibrium, for $\varepsilon > 0$, if for all $i = 1, \ldots, n$,

$$\mathcal{U}_i\left(s_i^*, s_{-i}^*\right) + \varepsilon \geq \mathcal{U}_i\left(s_i, s_{-i}^*\right) \text{ for all } s_i \in S_i. \qquad (1)$$

In words, (1) says s^* is an ε-Nash equilibrium if no player can unilaterally deviate from s^* and improve their utility by more than ε. We will be interested in cases where ε is small.

IV. The Cow-Path Ring Problem

In the CPP, exploring $\mathbb{R}^-$ and $\mathbb{R}^+$ are mutually exclusive tasks, implying the cow must initially have a contingency plan to turn around at least once. Here, we consider searching along the unit ring, an environment that affords the cow new search possibilities, including being able to conduct an exhaustive search without ever turning around. To this end, let $\mathcal{R}_1$ denote the unit ring and O a point on $\mathcal{R}_1$, henceforth, to be referred to as the origin. We define the Cow-Path Ring Problem (CPRP) as follows:

Definition 4.1: (***The Cow-Path Ring Problem***) A patch of clover, $\mathcal{T}$, is distributed on the unit ring, $\mathcal{R}_1$, according to a known density function $f : \mathcal{R}_1 \to \mathbb{R}$. Starting from O, a cow C, capable of moving at unit speed, wishes to discover $\mathcal{T}$ in minimal expected time. It is assumed C can change directions instantaneously, but is nearsighted and can discover $\mathcal{T}$ only when she is directly over it.

The following coordinate system provides the descriptive power needed to concisely articulate search plans in the CPRP. We refer to a point $x \in \mathcal{R}_1$ by the number in $[0, 2\pi]$ that represents the CCW distance from O to x. At the same time, we may also refer to x by the number in $[-2\pi, 0]$ whose magnitude represents the CW distance from O to x.

We now consider how to formally represent search plans in the CPRP. Given the close association to the CPP, it comes as little surprise that, in the CPRP, search plans are again naturally specified by an initial heading, followed by a sequence of turn-around points. However, because $\mathcal{R}_1$ is a closed curve, C now has the option of maintaining her heading and sweeping $\mathcal{R}_1$ until she discovers $\mathcal{T}$. Indeed, it is easy to see that once a contiguous segment of length $\ell \geq \pi$ has been explored along $\mathcal{R}_1$, sweeping becomes the optimal online contingency for finding $\mathcal{T}$.

With these ideas in mind, we represent a search plan in the CPRP by a sequence $s = \{x_i\}_{i=0}^n$, where $x_i \in \mathcal{R}_1$ is the coordinate of the i-th turn-around point. Because it proves notationally advantageous for the analysis to follow, we assume $x_0 = 0$, i.e., the origin. Semantically, s stipulates the search for $\mathcal{T}$ is to evolve by exploring the following path (delineated for the case where $x_1 > 0$ and n is even):

$$0 \xrightarrow{ccw} x_1 \xrightarrow{cw} x_2 \xrightarrow{ccw} x_3 \xrightarrow{cw} \cdots \xrightarrow{cw} x_n \xrightarrow{ccw} (2\pi + x_n). \quad (2)$$

In words, (2) states that, starting from the origin, C travels in the CCW direction toward x_1. Should she reach x_1 having not found $\mathcal{T}$, she reverses direction and travels toward x_2. Should she reach x_2, still having not found $\mathcal{T}$, she again reverses direction and travels toward x_3, and so on. Now, should she reach x_n, with $\mathcal{T}$ still proving elusive, she reverses direction one last time and then sweeps $\mathcal{R}_1$ until $\mathcal{T}$ is found. For the time being, we leave open the possibility that $n \to \infty$, and remark that there exist distributions in the CPP for which C's optimal search plan budgets a countably infinite number of turns. The pertinent geometric features of the CPRP are illustrated in Figure 1.

Since we are interested in optimal search plans, we restrict our focus to those s that satisfy the following conditions:

1) $x_i \cdot x_{i+1} \leq 0$, for $0 \leq i \leq n-1$,
2) $|x_i| < |x_{i+2}|$, for $1 \leq i \leq n-2$, and
3) $|x_i| + |x_{i+1}| \leq \pi$, for $1 \leq i \leq n-1$.

Taken together, conditions 1 and 2 ensure C alternate between moving in the CW and CCW directions when searching unexplored regions. Violation of either condition would introduce the possibility of reversing direction in a previously explored region — a wasteful policy. Condition 3, in conjunction with the first condition, ensures the cow does not waste time backtracking over previously explored territory to access unexplored regions can be more readily reached by maintaining the course and continuing to travel in the current direction.

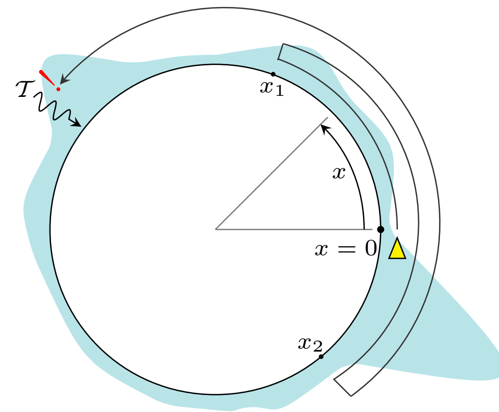

Fig. 1. The Cow-Path Ring Problem. The target density $f(x)$ as a function of radial position, $x \in \mathcal{R}_1$, is shown in blue. In this instance, the target, $\mathcal{T}$, is located along the North-West portion of the ring. In the search plan shown, the cow (yellow triangle) travels in the CCW direction toward x_1, where, having not found $\mathcal{T}$, she reverses direction, and travels in the CW direction toward x_2. Upon reaching x_2, having still not found $\mathcal{T}$, she again reverses direction and continues with her search in the CCW direction. The site at which $\mathcal{T}$ is found is illustrated with a red exclamation mark.

Assuming the well-formedness conditions above, we are now in a position to consider the discovery time, or cost, of search plan s, which we denote as $c(s)$. Recognizing $c(s)$ is a random variable, and following the approach in [5], we have (again assuming n is even and, thus, $x_n < 0$):

$$\mathbb{E}[c(s)] = \int_{y=0}^{2\pi + x_n} x f(x)\, dx + \int_{x=0}^{|x_n|} x f(-x)\, dx + 2\sum_{i=1}^{n} |x_i| \left(1 - F(x_i, x_{i-1})\right), \quad (3)$$

where we have extensively employed our coordinate system and for any $y, z \in [-2\pi, 2\pi]$ such that $yz \leq 0$, we define $F(y, z)$ to be the probability that $\mathcal{T}$ is located along the segment of $\mathcal{R}_1$ with endpoints y and z, i.e., $F(y, z) = \left| \int_y^z f(x)\, dx \right|$.

As mentioned, because $\mathcal{R}_1$ is a closed curve, a viable search plan is to sweep $\mathcal{R}_1$ in the CCW direction. Because this strategy is somewhat unique, we denote it explicitly by $s_{ccw} := \{0\}$. Similarly, we denote sweeping $\mathcal{R}_1$ in the CW direction by $s_{cw} := \{0, 0\}$, i.e., C immediately turns around. For target distribution f, let $s^* = \{0, x_1^*, x_2^*, \ldots, x_{n^*}^*\}$ denote a search plan with minimal expected cost. Recalling that C moves at unit speed, we immediately have the following upperbound on the search time:

$$\mathbb{E}[c(s^*)] \leq \min\left(\mathbb{E}[c(s_{ccw})], \mathbb{E}[c(s_{cw})]\right) < 2\pi. \quad (4)$$

Before addressing the optimal number of turns to schedule in a search plan, we introduce machinery to better describe the maximum number of turns C could be forced to make under search plan s. We note that this worst-case statistic is deterministic, although the number of turns is, in general, a random variable. We let $|s|$ denote the number of turn-around points in s. According to this definition, $|s^*| = n^*$. Finally, let s_n^* denote an optimal search strategy containing n turn-around points. For a large class of target densities, we have the following result.

Proposition 1: For a CPRP with bounded target density $f : \mathcal{R}_1 \to [a,b]$, $0 < a \le b$, n^* is finite. □

Before proving the above result, we comment on the applicability of the assumption that $f(x) \in [a,b]$ for $x \in \mathcal{R}_1$. In many applications for which the exact position of the target is unknown, it is reasonable to associate some non-zero probability to $\mathcal{T}$ residing in any interval of finite length. Moreover, it is also the case that for many applications, it is only over intervals of finite length that one can associate positive probability to finding $\mathcal{T}$. For the class of problems over which both of the preceding conditions are true, the assumption that $f : \mathcal{R}_1 \to [a,b]$, is reasonable.

Proof: Assume, to obtain a contradiction, that for a given $f : \mathcal{R}_1 \to [a,b]$ the optimal number of turn-around points is infinite. Let s^* denote an optimal search plan. From earlier discussion, s^* must satisfy $|x_i| + |x_{i+1}| \le \pi$ for all $i \in \mathbb{Z}_{\ge 0}$. Therefore, for any $n \in \mathbb{Z}_{>0}$, it must be the case that $|x_n| + |x_{n+1}| \le \pi$. Additionally, it must also be the case that $|x_{2n}| + |x_{2n+1}| \le \pi$. Therefore, with finite probability $p \ge a\pi$, the target will not have been found upon reaching x_{2n+1}. It follows that the expected search time satisfies $\mathbb{E}[c(s^*)] \ge pnx_n$, and, for sufficiently large n, exceeds $\min(\mathbb{E}[c(s_{cw})], \mathbb{E}[c(s_{ccw})]) = 2\pi$. The latter point contradicts the assumption that it is optimal to turn around an infinite number of times. ■

The next result formally captures the intuitive idea that search plans that employ more than n^* turn-around points can recover optimal performance.

Proposition 2: For target density $f : \mathcal{R}_1 \to [a,b]$, $0 < a \le b$, and $n > n^*$, $\mathbb{E}[c(s_n^*)]$ is arbitrarily close to $\mathbb{E}[c(s^*)]$. □

Proof: Since n^* is, by definition, the optimal number of turn-around points, the need to specify a search plan with one or more additional turn-around point(s) offers no advantage to C, i.e., $\mathbb{E}[c(s)] \ge \mathbb{E}[c(s^*)]$, for all $n \ge n^*$. In the case where $n = n^* + 1$, a sensible strategy to determine s_n^* is to nullify the extra turn-around point so that it stands only an insignificantly small chance of ever being required. That is, the target is virtually guaranteed to be found before having to act on the last turn around point. To this end, for small $\varepsilon > 0$, consider the search plan $s_{n^*+1} = \{s_{n^*}^*, x_{n^*+1}\}$, where

$$x_{n^*+1} = -2\pi \cdot \operatorname{sgn}(x_{n^*}) + x_{n^*} + \varepsilon \cdot \operatorname{sgn}(x_{n^*}). \tag{5}$$

The extra turn-around point in $s_{n^*+1}^*$ mandates that, should it prove necessary, C turn for the last time just before she would have finished sweeping $\mathcal{R}_1$ in $s_{n^*}^*$. Consequently, the expected discovery time changes only in those cases where $\mathcal{T}$ lies within distance ε to one side of x_{n^*}. Because $f(x) \in [a,b]$ for all $x \in \mathcal{R}_1$, $\lim_{\varepsilon \to 0^+} \mathbb{E}[c(s_{n^*+1})] = \mathbb{E}[c(s_{n^*}^*)]$, such that s_{n^*+1} discovers $\mathcal{T}$ in expected time arbitrarily close to that in s^*. For $n \ge n^* + 2$, we can again make the extra turn-around points have arbitrarily small chance of being acted upon by taking for $k = 0, 1, \ldots, n - n^*$,

$$x_{n^*+k} = \begin{cases} x_{n^*} - 0.5^k \cdot \varepsilon \cdot \operatorname{sgn}(x_{n^*}) & , k \text{ even} \\ -2\pi \cdot \operatorname{sgn}(x_{n^*}) + x_{n^*} + 0.5^k \cdot \varepsilon \cdot \operatorname{sgn}(x_{n^*}) & , k \text{ odd,} \end{cases} \tag{6}$$

such that $\mathbb{E}[c(s_{n*+k})] \to \mathbb{E}[c(s^*)]$ as $\varepsilon \to 0^+$. ■

The above discussion suggests the following practical approach to determine n^*. Staring from $n = 1$, compute an s_n^* that minimizes (3). If, for s_n^*, $|x_{n-1}| + |x_n| \approx 2\pi$, then an appropriate value to take for n^* is $n - 1$, otherwise let $n = n + 1$ and repeat.

We conclude that, left to her own devices, C's optimal search plan is to turn around a finite number of times at select location on $\mathcal{R}_1$. Moreover, the exact number of turns and the specific locations in question can be determined using a nonlinear optimization package in conjunction with the simple algorithm described. In the next section, we add a second hungry cow to the ring and study the behavior that emerges when the cows compete to find the patch of clover.

V. The Cow-Path Game

In the CPRP, it is not immediately obvious, for general f, at which points C should turn to find $\mathcal{T}$ in minimal expected time. In this section, we compound matters by adding a second cow to $\mathcal{R}_1$. We continue to assume $\mathcal{R}_1$ contains a single patch of clover. Each cow must now devise a search plan to find $\mathcal{T}$ based not only on their position and the spatial profile of f, but also the knowledge that a rival cow is present on $\mathcal{R}_1$. This latter point necessitates a re-examining of bovine logic. We assume that, although either cow would like to discover $\mathcal{T}$ quickly, it is more important that, given the scarcity of clover, that they be the first to find it. Since many elements of this scenario draw heavily from the CPRP, which itself is largely inspired by the CPP, we refer to it as the Cow-Path Game (CPG), formally defined as follows.

Definition 5.1: (***The Cow-Path Game***) A patch of clover, $\mathcal{T}$, is distributed on $\mathcal{R}_1$ according to a known density function $f : \mathcal{R}_1 \to \mathbb{R}^+$. For $i = 1, 2$, C_i is a rational cow, capable of moving at unit speed, that, starting from point $x_i(0) \in \mathcal{R}_1$ with initial heading $\phi_i(0) \in \{CW, CCW\}$, wishes to discover $\mathcal{T}$ first. It is assumed C_i can reverse directions instantaneously, but can discover $\mathcal{T}$ (if not already found) only when she is directly over it. Finally, it is assumed C_i is aware of her rival's position at all times.

The CPG is a game, not only in the vernacular, but also in the game-theoretic sense, because each cow wants to find $\mathcal{T}$ before the other. We emphasize this added level of strategic decision-making by henceforth referring to search plans in the CPG as *search strategies*. To understand how C_i should search in the CPG, let S_i denote the set of C_i's search strategies and $S = S_1 \times S_2$ the set of search strategy profiles. For brevity, we will henceforth refer to S simply as the set of search profiles. For $s \in S$, define the *landclaim* of C_i as

$$X_i(s) := \{x \in \mathcal{R}_1 : C_i \text{ is the first cow to visit } x \text{ under } s\}. \tag{7}$$

The utility C_i derives from s may then be expressed as

$$\mathcal{U}_i(s) = \mathbb{P}(C_i \text{ finds } \mathcal{T} \text{ under } s) = \int_{x \in X_i(s)} f(x)\, dx. \tag{8}$$

At time t, let $x_i(t)$ denote the position of C_i on $\mathcal{R}_1$. Additionally, let $\bar{x}_{i,t} : [0,t] \to \mathcal{R}_1$ denote the trajectory travelled by C_i up to and including time t, such that $x_{i,t}(\tau) = x_i(\tau)$ for $\tau \in [0,t]$. Given our information model, for any $t \geq 0$, C_i knows $x_{-i}(\tau)$ for all $0 \leq \tau \leq t$. Assuming C_i remembers the past position of C_{-i}, it follows that C_i knows $\bar{x}_{-i,t}$ for all times $t \geq 0$. Consequently, s_i will, in general, be a function of both $\bar{x}_{i,t}$ and $\bar{x}_{-i,t}$, such that upon letting $SC = \{\text{straight, turn}\}$ denote the set of local steering commands, we have that $s_i : (\bar{x}_{i,t}, \bar{x}_{-i,t}) \mapsto \{\text{straight, turn}\}$.

Before proceeding, it is worth mentioning why we have chosen to study competitive search in a ring environment. The reason is that we feel it is the most basic topology that affords non-trivial solutions. Namely, if the two cows were competing to find $\mathcal{T}$ on a line segment, the equilibrium solutions are easy to characterize: each cow travels toward their rival and, upon meeting, reverses direction to cover any previously unexplored territory.

To gain an appreciation for the types of complex decisions C_1 and C_2 face in the CPG, consider the scenario in Figure 2. In particular, focus on the segment of $\mathcal{R}_1$ described by $-\pi/4 \leq x \leq 0$, on which f has a global maximum. We are temped to ask, should C_1 explore this region immediately, or is she better off to "set it aside" and return to it later? The answer, of course, is highly contingent on C_{-i}'s strategy. For example, C_{-i}, could threaten, and occasionally follow through with, raids into territories that C_i values highly, and may have had her own aspirations of searching first. In the next section, we consider a restricted version of the CPG more amenable to a first analysis of the underlying strategic dynamics at play.

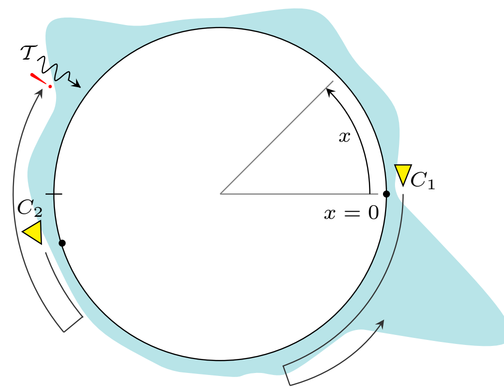

Fig. 2. An instance of the CPG illustrating the initial positions and initial headings of cows C_1 and C_2. Portions of each cow's trajectory, including the location of each cow's first turn-around point, are shown as gray paths with arrowheads. In this instance of the game, $\mathcal{T}$ is located along the North-West portion of the ring and the site at which $\mathcal{T}$ is found, in this case by C_2, is shown with a red exclamation mark.

VI. THE 1-TURN COW-PATH GAME

In this section, we consider the 1-Turn Cow-Path Game (1T-CPG) in which each cow is allowed to turn at most once.

Definition 6.1: ***(1-Turn Cow-Path Game)*** The 1-Turn Cow-Path Game is a special case of the CPG, defined in Definition 5.1, with the following amendment: for $i = 1,2$, C_i can turn around at most once.

In analyzing the 1T-CPG, it proves advantageous to work in terms of turn-around times, rather than turn-around points. We remark that the latter can be determined from the former, and vice versa, using the appropriate initial conditions and recalling that cows are assumed to move with unit speed. To this end, let $s_i = t$ denote the strategy in which at time t, C_i turns (irrespective of C_{-i}'s behavior) and, after turning, sweeps $\mathcal{R}_1$. Then $\mathcal{U}_i(t_i, t_{-i})$ is the probability C_i captures $\mathcal{T}$ given C_i and C_{-i} turn at times t_i and t_{-i}, respectively. Given the cows travel with unit speed and $\mathcal{R}_1$ has circumference 2π, we assume $t_i \in [0, 2\pi]$ and adopt the convention that $t_i = 2\pi$ implies C_i forgoes turning during the search.

In the interest of characterizing equilibria strategies, consider the search profile, $s = (2\pi, 2\pi)$, in which neither cow turns. In the event neither cow can unilaterally deviate and improve their utility by more than ε, s is an ε-equilibrium. If, however, one of the cows, say C_i, has an incentive to deviate, then it must be that C_i would prefer to turn at some time $t \in [0, 2\pi)$. Moreover, if C_i could guarantee that she turns first, then her optimal turning time is

$$t_i^1 = \arg\max_{t_i \in [0,2\pi]} \mathcal{U}_i\Big(t_i, \arg\max_{t_i \leq t_{-i} \leq 2\pi} \mathcal{U}_{-i}(t_i, t_{-i})\Big), \tag{9}$$

and C_{-i}'s best response is to turn at time

$$t_{-i}^1 = \arg\max_{t_i^1 \leq t_{-i} \leq 2\pi} \mathcal{U}_{-i}\left(t_i^1, t_{-i}\right). \tag{10}$$

Note that since the game is zero-sum, i.e., $\mathcal{U}_1(s) + \mathcal{U}_2(s) = 1$ for $s \in S$, (9) is reminiscent of C_i utilizing a maximin strategy for which C_{-i} is restricted in the actions she can choose from. Although (9) ensures C_i is doing the very best she can as the designated first mover, it does not guarantee C_{-i} is willing to resign herself to turning second. To explore the issue further, for $t \in [0, 2\pi]$, define the strategy $s_i = \sim t$ as

$$s_i \Rightarrow \begin{cases} C_i \text{ best responds to } C_{-i} & , C_{-i} \text{ turns by time } t \\ C_i \text{ turns around at } t & , \text{otherwise.} \end{cases} \tag{11}$$

Returning to C_i and C_{-i}, which turn at (9) and (10), respectively, it is reasonable to ask if C_{-i} would prefer to turn before C_i. If the answer is no, then we have reached an ε-equilibrium. Otherwise, assuming C_i adopts strategy $\sim t_i^1$, C_{-i}'s best strategy is to turn around at time

$$t_{-i}^2 = \arg\max_{0 \leq t_{-i} \leq t_i^1} \mathcal{U}_{-i}\Big(\arg\max_{t_{-i} \leq t_i \leq 2\pi} \mathcal{U}_i(t_i, t_{-i}), t_{-i}\Big), \tag{12}$$

and C_i to respond by turning at a time given the appropriate modification to (10). We may then ponder if C_i is content turning second in $(\sim t_i^1, \sim t_{-i}^2)$, or if she would prefer to, once again, turn first, and so on, and for each occurrence in which one cow elects to turn before the other, record the time at which the associated turn occurs. Let Ψ be the sequence of all such times that emerge from this procession of one-upmanship. In Algorithm 1, we capture the turn-around times and strategies the cows adopt while jockeying for position throughout the experiment we have described. The important issue of termination is addressed in the discussion to follow.

Algorithm 1 : solve for ε-equilibrium search strategies

1: i = index of any cow that favors turning under $s = (2\pi, 2\pi)$
2: $\Psi = \{\emptyset\}$
3: $s_i = \sim t_i : t_i = \arg\max_{t_i \in [0,2\pi]} \mathcal{U}_i\big(t_i, \arg\max_{t_i \le t_{-i} \le 2\pi} \mathcal{U}_{-i}(t_i, t_{-i})\big)$
4: $s_{-i} = \sim t_{-i} : t_{-i} = \arg\max_{t_i \le t_{-i} \le 2\pi} \mathcal{U}_{-i}(t_i, t_{-i})$
5: $\Psi \leftarrow \{t_i\}$
6: **while** no ε-equilibrium established **do**
7:   **if** $\mathcal{U}_{-i}(t_i, t_{-i}) + \varepsilon \ge \max_{0 \le t_{-i} \le t_i} \mathcal{U}_{-i}(\arg\max_{t_{-i} \le t' \le 2\pi} \mathcal{U}_i(t', t_{-i}), t_{-i})$ **then**
8:     break
9:   **else**
10:     $t_{-i} \leftarrow \arg\max_{0 \le t_{-i} \le t_i} \mathcal{U}_{-i}(\arg\max_{t_{-i} \le t' \le 2\pi} \mathcal{U}_i(t', t_{-i}), t_{-i})$
11:     $s_{-i} \leftarrow \sim t_{-i}$
12:     $\Psi \leftarrow \{\Psi, t_{-i}\}$
13:     $i \leftarrow (-i)$
14:   **end if**
15: **end while**

A few remarks are in order. First, sequence Ψ to emerge from the algorithm is non-increasing. Second, the difficulty in evaluating the required maximization operations is alleviated somewhat by the circular geometry of $\mathcal{R}_1$ and the fact both cows travel at the same speed. For example, assume, as in Figure 2, that C_1 and C_2 are initially heading toward one another. If $t_1 = t_2$, then the land claims X_1 and X_2 from (7) may be readily calculated from symmetry. However, should C_1 unilaterally deviate and turn, instead, at time $t_1 + \Delta t$, then C_i's other frontier is ultimately eroded by $2\Delta t$, making the land claims and utilities easy to calculate. In the event the cows are chasing each other (e.g., both cows initially have a CW heading) the strategy of the cow that turns second is simple: turn the instant before meeting the other cow.

We now describe some properties of Algorithm 1 pertaining to termination and which, if any, of the cows may have an incentive to deviate unilaterally from the strategy profiles prescribed throughout the algorithm.

Proposition 3: Let a, b, c be three consecutive times in Ψ. For search strategies $s_1 = \sim a$ and $s_2 = \sim b$, C_1 is the only cow with a unilateral incentive to deviate from the search profile $s = (s_1, s_2)$. Moreover, C_1's only profitable deviations involve preemptively turning before C_2. □

Proof: To begin, note the logic of Algorithm 1 implies $a \ge b \ge c$. Now consider the strategies $s_1 = \sim a$, $s_2 = \sim b$, and the search profile $s = (s_1, s_2)$. Since c immediately follows b in the non-increasing sequence Ψ, the clause on line 7 of Algorithm 1 must fail for s, indicating C_1 can obtain a utility increase (of more than ε) by deviating from $s_1 = \sim a$ in favor of $s_1 = \sim c$; for which C_1 turns first in $(\sim c, \sim b)$. To establish that C_1 must turn before C_2 to improve her utility, note that in s, C_1 is already best responding to C_2 turning at time b, implying the absence of any profitable unilateral deviations in which C_1 remains the second cow to turn.

Now consider C_2, the first cow to turn under s. Using logic similar to that employed in the previous paragraph, we conclude from b immediately following a in Ψ, that C_2 prefers turning first at or before time a, rather than responding to C_1 turning first at a. Therefore, C_2 has no incentive to deviate from s to a strategy in which she responds to C_1 turning first at time a. Furthermore, since adoption of the strategy $s_2 = \sim b$ was selected using the assignment in lines 10 and 11 of Algorithm 1, C_2 selects her turn-around time optimally over $[0, a]$, implying there are no profitable deviations that involve turning first in $[0, a]$. We conclude that C_1 is the only cow with a unilateral incentive to deviate from s, and any profitable deviations involve C_1 turning before C_2. ■

From Proposition 3, the strategies s_1 and s_2 iteratively assigned to the cows in Algorithm 1, are always such that it is only the cow that turns second in (s_1, s_2) that, by preferring to turn first, has an incentive to deviate. This realization begs the question, can this succession of one-upmanship continue indefinitely? The following proposition asserts that, for a large class of target densities, the answer is no.

Proposition 4: Let f be a bounded target density with $f(x) \le M$ for $x \in \mathcal{R}_1$ and finite $M > 0$. For any combination of initial cow positions $x_i(0)$ and initial cow headings $\phi_i(0)$, $i = 1, 2$, the sequence Ψ is finite. □

Proof: Assume, to obtain a contradiction, that there exist initial cow positions and headings, $x_i(0)$ and $\phi_i(0)$, $i = 1, 2$, such that Ψ is infinite. Let $\Psi = \{t_1, t_2, \ldots\}$. From Algorithm 1, Ψ is a non-increasing sequence. Moreover, because the cows cannot turn before time zero, Ψ is non-negative. It follows Ψ must approach a limiting value $\upsilon \ge 0$, and for any $\delta > 0$, there exists a sufficiently large $n_o(\delta) \in \mathbb{N}$ such that $0 \le t_n - t_{n+1} \le \delta$ for all $n \ge n_o(\delta)$.

For $\delta > 0$, let $a \ge b \ge c$ be three consecutive elements of Ψ such that $0 \le a - b \le \delta$ and $0 \le b - c \le \delta$. From the assumption on Ψ, such times are guaranteed to exist. Consider the following search profiles: $s_1 = (\sim a, \sim b)$, $s_2 = (a, \sim b)$, $s_3 = (\sim c, \sim b)$, and $s_4 = (\sim c, b)$. In s_1, C_1 best responds to C_2 turning first at time b. Since $a \ge b$, $\mathcal{U}_1(s_1) \ge \mathcal{U}_1(s_2)$, because the option of turning at time a is considered when forming the best response. Moreover, because c immediately follows b in Ψ, it must be the case that $\mathcal{U}_1(s_3) > \mathcal{U}_1(s_1) + \varepsilon$, and, subsequently, that $\mathcal{U}_1(s_3) > \mathcal{U}_1(s_2) + \varepsilon$. Now consider the search profile $s_4 = (\sim c, b)$, in which C_2 responds to C_1 turning first at time c, by turning at time b. However, because C_2 best responds to C_1 turning first at time c in s_3 we have $\mathcal{U}_2(s_3) \ge \mathcal{U}_2(s_4)$. Since the game is zero-sum, the inequality chain implies

$$\mathcal{U}_2(s_2) > \mathcal{U}_2(s_4) + \varepsilon. \tag{13}$$

The inequality in (13) indicates the difference in utility C_2 sees between s_2 and s_4 is more than ε. However, because f

is bounded and $0 \leq a-b \leq \delta$ and $0 \leq b-c \leq \delta$, we have

$$|\mathcal{U}_2(s_2) - \mathcal{U}_2(s_4)| \leq \Big| \int_{X_2(s_2)} f(x)\,dx - \int_{X_2(s_4)} f(x)\,dx \Big| \quad (14)$$

$$\leq \int_{X_2'} f(x)\,dx \;\leq\; 4M\delta, \quad (15)$$

where $X_2' = X_2(s_2) \ominus X_2(s_4)$ is the symmetric difference of $X_2(s_2)$ and $X_2(s_4)$; i.e., the collection of all elements that are in $X_2(s_2)$ or $X_2(s_4)$, but not both. Thus, by choosing δ such that $0 < \delta \leq \varepsilon/4M$, we have, from (15), that $|\mathcal{U}_2(s_2) - \mathcal{U}_2(s_4)| \leq \varepsilon$, which contradicts (13), thereby refuting the initial assumption, and establishing Ψ is finite. ■

We remark that in many applications, the assumption that f is bounded over $\mathcal{R}_1$ is quite reasonable. Namely, it is rarely the case that given a continuum, e.g, $\mathcal{R}_1$, of possible locations where $\mathcal{T}$ might be located, that one can be so bold as to assign positive probability to finding $\mathcal{T}$ at a specific point. A much more realistic model for searching applications is to assume, as we have done, that $\mathcal{T}$'s location is described by a continuous (and therefore bounded) density function. Combining Propositions 3 and 4 gives the following result.

Theorem 1: Let f be a bounded target density. The 1-Turn Cow-Path Game has an ε-Nash equilibrium in $S_1 \times S_2$. □

Proof: From Proposition 4, Algorithm 1 terminates with Ψ finite. Let the search strategies of C_1 and C_2 that emerge from Algorithm 1 be s_1 and s_2, respectively. Let $s = (s_1, s_2) \in S_1 \times S_2$ be the associated search profile. Let $i \in \{1,2\}$ be the index of the cow that turns first in s. From Proposition 3, C_{-i} is the only cow that could have an incentive to unilaterally deviate in s. However, if C_{-i} had an incentive to deviate from s, Ψ must have cardinality at least one greater than its actual value, a contradiction. Therefore, neither C_1 nor C_2 has an incentive to unilaterally deviate from s, implying s is an ε-Nash equilibrium. ■

VII. MULTI-TURN COW-PATH GAMES

In the preceding analysis, we assumed C_1 and C_2 may turn at most once. This is a rather severe limitation to impose on the hungry cows. In this section, we briefly describe how the 1T-CPG serves as a foundation for the study of CPGs where C_1 and C_2 may turn up to a finite number of times.

To begin, let $\mathcal{G}_1 = \mathcal{G}_{CPG}(f, x_i(0), \phi_i(0), n_i)_{i=1,2}$ denote the CPG with density f in which C_i has initial conditions $(x_i(0), \phi_i(0))$ and may turn up to $n_i > 0$ times. Let s be the search profile in which C_j turns first at a time denoted $t_j \in [0, 2\pi)$ in $\mathcal{G}_1$. Also, let $X(s,t_j) \subseteq \mathcal{R}_1$ denote the set of all points visited by a cow over $[0,t_j]$ in s. The decisions the cows face in the remainder of the game, i.e., the game unfolding for $t > t_i$, are precisely those captured by $\mathcal{G}_2 = \mathcal{G}_{CPG}(\tilde{f}, x_i(t_j), \phi_i(t_j), \tilde{n}_i)_{i=1,2}$, where

$$\tilde{f}(x) = \begin{cases} 0, & \text{if } x \in X(s,t_j) \\ f(x), & \text{otherwise} \end{cases}, \text{ for } x \in \mathcal{R}_1, \text{ and} \quad (16)$$

$$\tilde{n}_i = \begin{cases} n_i - 1, & \text{if } i = j \\ n_i, & \text{otherwise.} \end{cases} \quad (17)$$

A remark is in order as it relates to $\tilde{f}$ in $\mathcal{G}_2$. Equation (16) implies $\tilde{f}$ is deficient, i.e., $\int_{\mathcal{R}_1} \tilde{f}(x)dx < 1$, if, by time t_j, the cows have visited a subset of $\mathcal{R}_1$ on which the target may have been located. However, with respect to maximizing (8), we see that what is important is the value of $\int_{X_i} f(x)\,dx$, where X_i is the land claim acquired, on an on-going basis, throughout the game. From this perspective, we can think of the cows as gathering density throughout the search, which is well-defined even in the case of deficient target densities. Consequently, when we refer to a game using the notation above, it is with this tacit assumption in place.

With a dynamic notational system in place, we now address when to turn in a multi-turn game. Starting from $\mathcal{G}_1$, the optimal time for C_i to turn first is given by

$$t_i = \arg \max_{t_i \in [0,2\pi]} \{h(t_i)\}, \text{ with} \quad (18)$$

$$h(t_i) = \int_{X_i(t_i)} f(x)\,dx + \mathcal{G}^{*,i}_{CPG}\left(\tilde{f}, x_j(t_i), \phi_j(t_i), \tilde{n}_j\right)_{j=1,2}, \quad (19)$$

where $X_i(t_i)$ is the land claim of C_i acquired in $[0,t_i]$ and $\mathcal{G}^{*,i}_{CPG}\left(\tilde{f}, x_j(t_i), \phi_j(t_i), \tilde{n}_j\right)_{j=1,2}$ is the optimal utility C_i can acquire in $\mathcal{G}_{CPG}\left(\tilde{f}, x_j(t_i), \phi_j(t_i), \tilde{n}_j\right)_{j=1,2}$. Therefore, in scheduling her turns, C_i considers not only the density she acquires prior to turning, but also the density gathered in the equilibria associated with the resultant game. To solve (18) using dynamic programming, we must first solve the CPG for the relevant base case scenarios.

Having solved the 1T-CPG in Section VI, the remaining base cases are those in which one cow, say C_i, may turn $n_i \geq 2$ times, and the other cow, C_{-i}, has used up all of her turns. Any other finite turn game will degenerate to one of the aforementioned cases following a sufficient number of turns, as depicted in Figure 3. For $n_i \geq 2$, C_i's best strategy is to immediately orient herself so that she is traveling toward C_{-i}. Establishing this alignment takes at most one turn. Then, C_i proceeds to travel toward C_{-i} before turning one last time at the instant just before running into C_{-i}. This last turn ensures C_i captures any unclaimed density still remaining on $\mathcal{R}_1$.

We remark that although the multi-turn games degenerate to simpler games as the cows turn, the dynamic programming approach suggested by Figure 3 requires these reduced games be solved for a variety of initial cow headings and positions. Therefore, employing dynamic programming-based to study multi-turn CPGs may exact a rather steep computational price. Nevertheless, it is reassuring to know there exists a well-developed methodology to address, at least at a theoretical level, multi-turn CPGs.

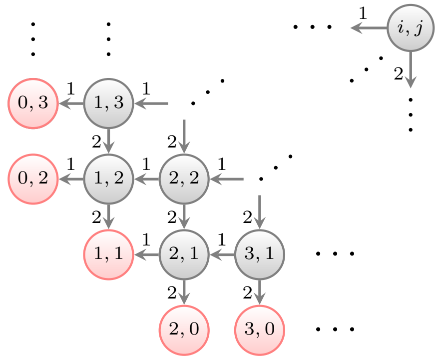

Fig. 3. Transitional diagram relating inter-dependence of families of multi-turn CPGs. The node labelled with the pair (i, j) denotes the family of games in which C_1 and C_2 may turn up to i and j times, respectively. Base case games for which equilibria strategies may be found directly are shaded in red. All other families of games are shaded in gray. The arrows indicate how one family of game degenerates to another when a cow turns. For example, the $(2, 2)$-CPG becomes an instance of the $(1, 2)$-CPG when C_1 turns, and an instance of the $(2, 1)$-CPG when C_2 turns.

VIII. Conclusions and Future Directions

We introduced the Cow-Path Game to study strategic decision-making in systems where multiple, self-interested, mobile agents compete to find a target. It was assumed the agents have minimal sensing capabilities and limited prior knowledge of the target's location. In the scenarios of interest, each agent must devise a search strategy for exploring the environment that is efficient in light of inter-agent competition.

As a prelude, we considered the problem of a single hungry cow that searches for a target on the unit ring given a density function on the target's location. A second cow was subsequently added to the ring, making the search strategic in a game-theoretic sense: each cow now structuring her search based not only on her position and the prior information she possesses, but also the presence of a rival cow. We established the existence of equilibria strategies in a variant of the game for which each cow can turn at most once and discussed how these ideas may be extended to the case of multi-turn games.

Looking forward, it remains to consider versions of the CPG in which three or more cows are involved in searching for the target. Additionally, although our efforts have focused on the ring, there is no conceptual barrier to proposing CPG-like games in other environments, e.g., star graphs, networks, or regions of $\mathbb{R}^2$. These alterations fundamentally change the nature of the game and it is unclear in what capacity the results reported herein can be used to jumpstart the analysis of such games.

Also of interest are persistent scenarios, in which optimal strategies are those that prove viable in a long-run or steady-state setting. In this context, we envision targets arriving in $\mathcal{Q}$ on an on-going basis, according to an appropriate spatio-temporal process, and either expiring with time or being consumed by search agents upon detection. Such a setup, especially in a planar environment, could prove a useful model for animals foraging in the wilderness, or the diffusion of bacteria over a nutrient-laden plate. Shifting focus to these persistent scenarios requires redefining agent utility function s to reflect the ongoing nature of the game, which, in turn, is likely to require the adoption of new strategic search protocols.

Game theorists have delineated various definition of equilibria to articulate precise notions of efficiency and social equality among self-interested agents. To date, we have focused exclusively on ε-Nash equilibria. To this end, it may prove advantageous to investigate the existence of other, perhaps more relevant, notions of equilibria. Along similar lines, it could prove telling to contrast the expected time until $\mathcal{T}$ is discovered in the CPG, with the associated time in settings where the cows cooperate to find $\mathcal{T}$. Any result in this direction would shed valuable insight into the cost of anarchy in competitive search games.

As a final new direction, we wish to consider settings in which select targets broadcast their position to search agents at the time they enter the environment. More in line with the traditional dynamic vehicle routing framework, one can imagine search agents capturing a certain fraction of their targets by exploiting this new source of information, with the remaining fraction obtained, as in the CPG, through local, short-range searching. Such a model could serve as a viable abstraction of taxi systems that permits customers to either flag down a taxi in realtime, or call a centralized dispatcher and request pick up at a designated location.

References

[1] L. Liu, C. Andris, and C. Ratti, "Uncovering cabdriver's behavior patterns from their digital traces," *Computers, Engineers and Urban Systems*, vol. 34, 2010.
[2] T. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics: A survey," *Autonomous Robots*, 2011.
[3] L. D. Stone, *Theory of Optimal Search*. Academic Press, 1975.
[4] R. Bellman, "An optimal search cream," *SIAM review*, 1963.
[5] A. Beck, "On the linear search problem," *Israel Journal of Mathematics*, 1964.
[6] P. Seiler, A. Pant, and K. Hedrick, "On the linear search problem," *Israel Journal of Mathematics*, vol. 2, pp. 221–228, 1964.
[7] A. Beck, "More on the linear search problem," *Israel Journal of Mathematics*, 1965.
[8] A. Beck and D. Newman, "Yet more on the linear search problem," *Israel Journal of Mathematics*, vol. 8, pp. 419–429, 1970.
[9] R. Baeza-Yates, J. Culberson, and J. Rawlins, "Searching the plane," *Information and computation*, vol. 106, pp. 234–252, 1993.
[10] M. Kao, J. Rief, and S. Tate, "An optimal randomized algorithm for the cow-path problem," *Information and computation*, vol. 131, pp. 63–79, 1996.
[11] E. Demaine, S. Kekete, and S. Gal, "Online search with turn cost," *Theoretical computer science*, 2006.
[12] D. Song, C. Kim, and J. Yi, "Stochastic modeling of the expected time to search for an intermittent signal source under a limited sensing range," *Proc. of Robotics: Science and Systems*, 2010.
[13] V. Huynh, J. Enright, and E. Frazzoli, "Persistent patrol in stochastic environments with limited sensors," *Proc. AIAA Conference on Guidance, Navigation, and Control*, 2010.
[14] S. Gal, *Search Games*. Academic Press, 1980.
[15] S. Alpern and S. Gal, *The Theory of Search Games and Rendezvous*. Springer, 2010.
[16] T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. SIAM, 1999.
[17] M. Zhu and E. Frazzoli, "On competitive search games for multiple vehicles," *Proc. of Conference on Decision and Control*, 2012.
[18] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.