Multi-Agent Cooperative Pursuit-Evasion Strategies Under Uncertainty

Kunal Shah[1] and Mac Schwager[2]

[1] Department of Mechanical Engineering, Stanford University, Stanford, CA
k2shah@stanford.edu
[2] Department of Aeronautics and Astronautics, Stanford University, Stanford, CA
schwager@stanford.edu

Abstract. We present a method for a collaborative team of pursuing robots to contain and capture a single evading robot. The main challenge is that the pursuers do not know the position of the evader exactly nor do they know the policy of the evader. Instead, the pursuers maintain an estimate of the evader's position over time from noisy online measurements. We propose a policy by which the pursuers move to maximally reduce the area of space reachable by the evader given the uncertainty in the evader's position estimate. The policy is distributed in the sense that each pursuer only needs to know the positions of its closest neighbors. The policy guarantees that the evader's reachable area is non-increasing between measurement updates regardless of the evader's policy. Furthermore, we show in simulations that the pursuers capture the evader despite the position uncertainty provided that the pursuer's measurement noise decreases with the distance to the evader.

Keywords: Multi-Agent Pursuit-Evasion, Game Theoretic Control, Reachability Methods

1 Introduction

The advent of consumer multi-rotors has created a potential hazard for airports and other sensitive airspace, as pilots can easily fly UAVs into restricted areas (either intentionally or accidentally).[3] In recent years, numerous airports have been shut down due to rogue drones.[4] Some countries are even attempting to train eagles to disable unauthorized UAVs flying in sensitive areas.[5] In this paper we propose an online cooperative pursuit algorithm for a team of drones to pursue and ultimately disable a rogue evader drone. We focus on the realistic case in which the pursuers do not know the exact location of the evader but have to estimate the evader position with noisy on-board sensors, such as cameras and LIDAR. This paper builds upon existing works [1–3] which provide cooperative

[3] cnn.com/2015/01/26/technology/security/drone-white-house/
[4] bbc.com/news/uk-40476264
[5] cbsnews.com/news/dutch-police-use-eagles-to-take-down-illegal-drones/

pursuit strategies in bounded environments based on Voronoi cells but assume that the pursuers know the position of the evader exactly.

To conservatively manage the uncertainty in the evader's position we focus on an uncertainty ellipsoid that contains the evader. We extend the notion of safe-reachability to this ellipsoid to characterize the set of points that the evader can reach before any pursuer regardless of the evader's initial position. We derive a controller for the pursuers that maximally reduces the area of this conservative safe-reachable set. We prove that, using our policy, the area of the safe-reachable set is non-increasing between measurement updates of the evader's position. The policy is distributed in the sense that the pursuers only require the relative positions of their closest neighboring pursuers, they do not need to communicate their intended control actions or explicitly coordinate in the computation of their control actions. Instead, coordination is ensured by the geometric construction of the algorithm. We show in simulations that if the noise in the pursuers' sensors decreases sufficiently strongly as they get closer to the evader (as is typically the case for sensors such as cameras and LIDAR) then the pursuers capture the evader despite the position uncertainty.

This paper is organized as follows. The rest of this section discusses related work. Section 2 formulates the pursuit-evasion problem and gives the necessary mathematical background on greedy safe-reachable set minimization. Section 3 formalizes the evader's safe-reachable set given an elliptical uncertainty region and derives our policy for the pursuers to maximally decrease this safe-reachable set. Section 4 discusses a set-based filtering method [4] we use to update the pursuer's estimate of the evader's position with noisy measurements. Simulations are presented in Section 5 to showcase the policy's performance.

1.1 Related Work

Pursuit techniques based on decreasing the size of an evader's safe reachable set, especially in cases of multiple agents, have shown promise due to their scalability. Instead of minimizing the distance between the pursuers and evader directly, these methods seek to decrease the set of points that the evader can safely reach before being intercepted by a pursuer. This eventually reduces the safe options available to the evader and ultimately leads to capture. For arbitrary dynamic agents this set can be found algorithmically by discretizing the space [3], finding intersections of time-optimal reachability boundaries (isochrones) [5], or with respect to a designed cost function (e.g. fuel) [6]. In the case of single integrator agents in an obstacle-free environment, the safe-reachable set of an evader is simply its Voronoi cell. The area-minimization pursuit policy in this case is for each pursuer to move towards the centroid of the Voronoi boundary that it shares with the evader [1, 2, 7] or, in the case of one pursuer and evader, by discretizing the state and action spaces [3]. We present a different method for maximally reducing the safe-reachable set when the evader is known to lie in an elliptical uncertainty region.

A comprehensive review of pursuit-evasion with and without uncertainty is given in [8] where the pursuing agents utilize search algorithms to decrease the

uncertainty in the evader's position. Some proposed solutions [9] first discretize the space and use multiple agents to preform a coordinated grid search while updating a map of the environment. Other works [10, 11] discuss centralized methods to search and pursuit problems of pursuers with limited fields of view. Time-optimal centralized pursuit-evasion techniques have also be studied as a differential game [12] by solving the corresponding Hamilton Jacobi Isaacs (HJI) equations [13, 14]. Other similar methods involve graph based approaches or solving a set of ordinary differential equations using approximations to extrema functions [15]. However, these methods require finding a terminal point from the initial conditions and tend to scale poorly with the number of agents due to the curse of dimensionality. Our work considers the pursuit of a mobile evader in continuous space where the uncertainty, represented by an ellipsoid, is updated from on-board noisy measurements of the evader's position and each pursuer calculates its own control action independently.

2 Problem Formulation

Consider a group of N pursuing agents and a single evader in an unbounded, obstacle free environment. The position of pursuer i for $i \in \{1, \ldots, N\}$ at time t is $p_i(t)$ and the evader's position is $e(t)$. Let all agents have single integrator dynamics and the same fixed maximum speed,

$$\begin{aligned} \dot{p_i} = u_i \quad \dot{e} = v, \\ ||u_i|| \leq 1 \quad ||v|| \leq 1, \end{aligned}$$

where u_i, v denotes the control actions for pursuer i and the evader, respectively. **Capture** occurs when $||e - p_i|| \leq r_{cap}$ for some $i \in \{1, \ldots, N\}$ for some chosen $r_{cap} > 0$. The pursuers measure the evader's position with noisy on-board sensors to generate an estimate of e. The sensor noise naturally causes uncertainty in the estimate, which we assume is described by an ellipsoid as is common with the Kalman filter and its many variants. Our goal is to find a policy u_i for the pursuers to capture the evader by only using their estimate of e and without any knowledge of the evader's control policy v. We make no assumptions about the policy of the evader, or the information available to the evader (i.e. we seek a pursuit policy that will lead to capture for any evader policy). Before we present our policy, we review an existing pursuit policy that guarantees capture when the evader position is known exactly by the pursuers. We later build upon this policy for the case of uncertain evader position.

2.1 Background

Instead of directly minimizing the distance between the evader and a pursuer, which is the classic approach to pursuit-evasion, the works [1–3] introduce a greedy area-minimization strategy which is designed to decrease the set of points the evader can safely move to without being captured. This set is known as the

evader's **safe-reachable set** and is the region reachable by the evader before any other agents. As discussed by [3], decreasing the area of the evader's safe-reachable set will eventually lead to capture for any arbitrarily small capture radius. For the case of single integrator dynamics, the set of points reachable by the evader before any pursuer is equivalent to the set of points closer in the Euclidean sense to the evader. Thus, the evader's safe-reachable set V_e can be represented as a cell in a Voronoi tessellation generated by the positions of all the agents $(p_1, \ldots, p_N, e)$ [2],

$$V_e = \{q \mid ||q - e|| \leq ||q - p_i|| \ \forall i\}.$$

This set is a time varying quantity as each of the agent's positions are changing in time. Let $A_e(t) = \int_{V_e(t)} dq$ be the area of the evader's safe-reachable set at time t. We assume that the evader's safe-reachable set is initially bounded, which necessitates a minimum number of pursuers. If this set were initially unbounded A_e would be infinite and the evader could trivially escape. The greedy area-minimization problem can be thought of as a classical game where the payoff, $\dot{A}_e$, is greedily minimized by the pursuer and maximized by the evader. The corresponding optimization problem is

$$\begin{aligned} \underset{u_i}{\text{minimize}} \ \underset{v}{\text{maximize}} \quad & \dot{A}_e = \sum_i^N \frac{\partial A_e}{\partial p_i}^T \dot{p_i} + \frac{\partial A_e}{\partial e}^T \dot{e} \\ \text{subject to} \quad & \dot{p_i} = u_i, \ ||u_i|| \leq 1 \ \forall i \\ & \dot{e} = v, \ ||v|| \leq 1. \end{aligned} \tag{1}$$

We refer to the terms $\frac{\partial A_e}{\partial e}$ and $\frac{\partial A_p}{\partial p_i}$ in (1) as the area flux quantities as they relate the "flow" of area into the evader's safe-reachable set to the movement of the agents. While it is generally difficult to find saddle-point equilibrium to min-max games, the additive structure of (1) lets each agent calculate its optimal control independently. Thus the game can be solved by finding the solution to $N+1$ separate optimization problems, which naturally results in a decentralized policy for the pursuers. Once the flux quantity $\frac{\partial A_e}{\partial p_i}$ is known each pursuer solves the following optimization,

$$\begin{aligned} \underset{u_i}{\text{minimize}} \quad & \frac{\partial A_e}{\partial p_i}^T \dot{p_i} \\ \text{subject to} \quad & \dot{p_i} = u_i, \ ||u_i|| \leq 1, \end{aligned}$$

with solution

$$u_i^\star = -\frac{\frac{\partial A_e}{\partial p_i}}{\| \frac{\partial A_e}{\partial p_i} \|}. \tag{2}$$

When the safe-reachable set of the evader is a cell of a Voronoi tessellation, the optimal action with respect to the above optimization is for the pursuer to move towards the centroid of the shared Voronoi boundary at maximum speed. In 3D these boundaries are sections of planes and in 2D they are lines segments,

hence the safe-reachable set is a polyhedron in 3D or polygon in 2D. Figure 1 shows the safe-reachable set of an evader in 3D. For this case, the area flux terms can be calculated by noting that each boundary segment is a section of a plane that is always orthogonal to the vector between the evader and the corresponding pursuer [16]. The flux terms can be found by perturbing each agent and finding the resulting variational change in the safe-reachable area [1, 3]. While a worst-case area-maximizing evader policy can be found by using the evader's flux term $\frac{\partial A}{\partial e}$, each pursuer does not need this information to calculate its area-minimizing policy.

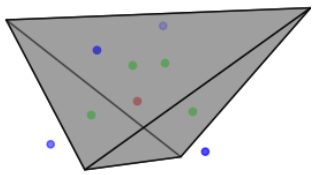

Fig. 1. A group of pursuers (blue) in 3D with a single evader (red). The pursuers' control action when the evader position is known exactly is to move towards the centroid (green) of their shared boundary (gray) with the evader, as proposed in [1, 3].

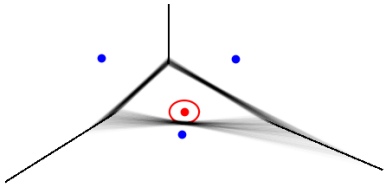

Fig. 2. The effect on the evader's safe-reachable set due to imperfect measurements is shown here. Overlaid instances of possible safe-reachable sets (black) when the pursuers jointly sample the evader's position from a Gaussian distribution. The 99% confidence ellipsoid is in red.

3 Pursuit-Evasion with Uncertain Evader Position

Suppose that the evader's position is not known exactly, but is known to lie within an uncertainty ellipse. Under these circumstances, we will show that the safe-reachable set is not polygonal. Furthermore, implementing the above area-minimization method without consideration of the evader's uncertainty may not always achieve capture and could lead to the evader escaping. Figure 2 showcases the variability in the safe-reachable set when the pursuers are uncertain about the evader's position. This necessitates the formulation of a safe-reachable set that takes into account the uncertainty region of the evader as well as a method to calculate the corresponding area flux $\frac{\partial A_e}{\partial p_i}$ to formulate a pursuit policy.

Figure 4 shows a similar configuration of agents but takes into account the uncertainty in the evader's position to construct the safe-reachable set. In the next section we solve for the corresponding safe-reachable and present a principled way of calculating the area flux terms for two general classes of boundary representations: explicit and implicit. We first show that this method recovers the known policy (i.e. move to the centroid of the shared Voronoi boundary) for the noiseless case found in [2]. We then use our method to find the area flux for

the case with elliptical uncertainty, resulting in our uncertainty-aware pursuit policy.

Depending on the measurement model used, the uncertainty quantifies a spatial region (normally an ellipsoid) within which the evader is expected to lie. For typical Gaussian noise models a confidence region has a probabilistic interpretation, such as the 99% confidence ellipsoids. In the case of bounded noise models, we can find a confidence region that contains the evader with certainty. For simplicity, in the remainder of this work we assume the bounded noise model (so the evader is known to lie in the ellipse with certainty), however the extension to an arbitrary Gaussian confidence ellipse is straightforward. In order to incorporate uncertainty ellipse into the safe-reachable set we consider generalized Voronoi tessellations which allow sets as generators [17–19]. Let each point in V_E be closer to the set $E \subset \mathbb{R}^n$ than any point p_i,

$$V_E = \{q \mid ||E - q|| \leq ||p_i - q|| \ \forall i\},$$

where $||E - q|| = \inf_x\{||x - q||_2 \mid x \in E\}$ is the standard distance-to-set metric. We also assume that the pursuers use a distributed filtering algorithm [20–22] that results in a common ellipsoidal uncertainty region E shared by all pursuers. Since $e \in E$ it follows that $V_e \subseteq V_E$ and $A_e \leq A_E$. Much like how $\dot{A}_e$ was minimized in the previous section we aim to minimize $\dot{A}_E$ by letting the pursuers use V_E as an upper bound to V_e.

3.1 Uncertainty Aware Safe-Reachable Set

We can form an expression for the boundary between E, parametrized by (μ, Σ), and p_i, where $E(\mu, \Sigma) = \{x \mid (x - \mu)^T \Sigma^{-1}(x - \mu) \leq 1\}$. Since E is convex, any point q along the boundary of the safe-reachable set can be decomposed as $q = a + rn_a$ where a is the closest point to q on the boundary of E and n_a is the outward normal of E at a. Figure 3 shows the construction of the boundary. For each q on the boundary we can find $r > 0$ such that $||E - q|| = ||p_i - q||$. Thus, in solving for r we form a map between the boundary of the ellipsoid and the boundary of the uncertainty aware safe-reachable set. After some simplification,

$$r = \frac{||p_i - a||^2}{2(p_i - a)^T n_a}.$$

Since E is an ellipsoid, there are parameterizations of the boundary $a \in \partial E$ and corresponding surface normal n_a of the ellipsoid. Let $F^T F = \Sigma$ and $d(\theta)$ be a parametrization of an n-sphere with $\theta \in \mathbb{R}^{n-1}$ with appropriate bounds. For a given θ

$$a(\theta) = \mu + Fd(\theta) \ \text{ and } \ n_a(\theta) = \Sigma^{-1} Fd(\theta)$$

fully determine a unique point and corresponding normal on the boundary of the ellipsoidal set E. It follows that the uncertainty aware boundary is

$$q_i(\theta) = \mu + Fd(\theta) + \frac{||p_i - \mu - Fd(\theta)||^2}{2[(p_i - \mu - Fd(\theta))^T \Sigma^{-1} Fd(\theta)]} \Sigma^{-1} Fd(\theta). \tag{3}$$

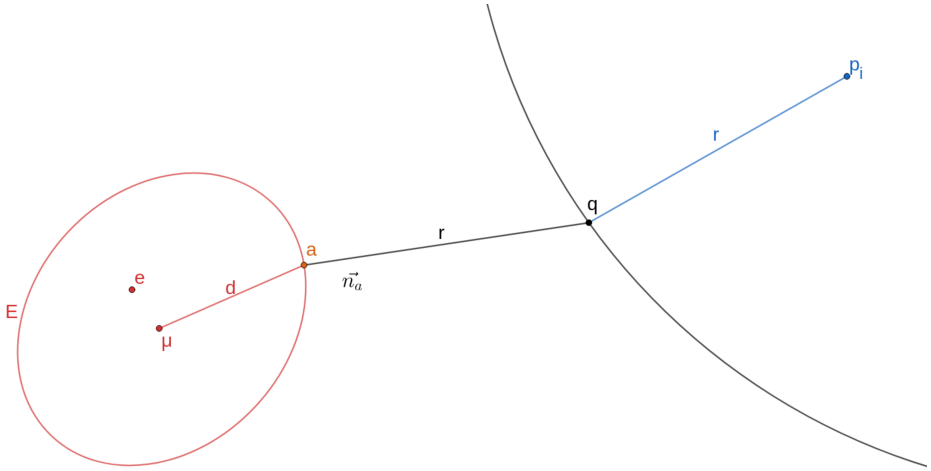

Fig. 3. Boundary Construction of an Uncertainty Aware Safe-Reachable Set

Figure 4 shows the uncertainty aware safe-reachable set of the evader with three pursuers. Each segment of the boundary is only shared with a single pursuer, which is key to the decentralized policy as each pursuer can only affect its corresponding boundary.

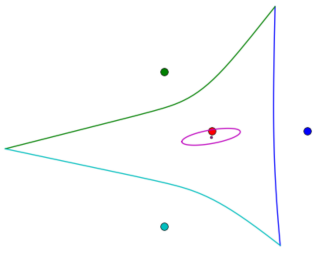

Fig. 4. Uncertainty aware safe-reachable set of an evader (red), with an uncertainty ellipse (magenta). The boundary is decomposed into three segments, one for each pursuer (blue, green, cyan).

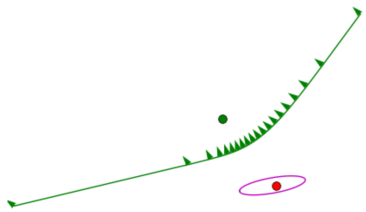

Fig. 5. A single boundary segment with various normals (triangles) along the boundary which are used to calculate the pursuer's policy

3.2 Area Flux

In order to discuss how the pursuers alter the area of the evader's safe-reachable set we invoke the Leibniz Integral Rule [2], which allows us to differentiate a

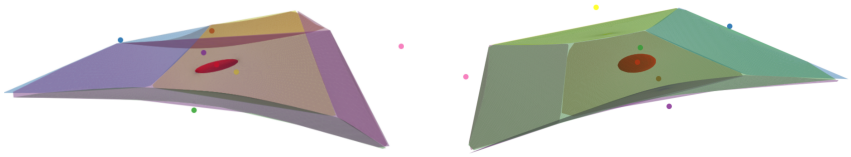

Fig. 6. Two views of an uncertainty aware safe-reachable set in 3D with six pursuers and one evader and uncertainty ellipsoid (red)

quantity defined by an integral, such as area. Using this Rule, we can equate the time derivative of a quantity in a control region to the flux of that quantity across the region's boundary and the change in size of the region. This allows us to relate the effect a pursuer has on the evader's safe-reachable area to changes in the shared boundary between them. We can decompose the boundary of the reachable set $\partial V_e = \bigcup_i^N \partial V_{ei}$ into N segments and formulate the change in area due to each pursuer i

$$\frac{\partial A_e}{\partial p_i}^T \dot{p_i} = \int_{\partial V_{ei}} n_i^T \frac{\partial q}{\partial p_i} \dot{p_i} \, ds. \tag{4}$$

Specifically, (4) relates the area flux of pursuer i to the unit outward normal vector of the boundary n_i and the boundary Jacobian $\frac{\partial q}{\partial p_i}$ at a point q on the shared boundary of the evader and pursuer i. Figure 5 shows a single boundary segment with various normals along the boundary. Geometrically, $\frac{\partial q}{\partial p_i}\dot{p_i}$ can be thought of as the velocity of a boundary point q induced by the movement of pursuer i. The infinitesimal amount of area flux due to the change in the pursuer's position can be decomposed as $n_i^T \frac{\partial q}{\partial p_i}\dot{p_i} \, ds$ where ds is the infinitesimal length/area of the boundary surface. The flux $\frac{\partial A_e}{\partial e}$ can be found similarly by summing the area change over each segment to find the evader's area maximization policy. We provide two methods for finding the flux $\frac{\partial A_e}{\partial p_i}$ based on representations of the boundary ∂V_{ei}.

Area Flux: Explicit

The flux terms can be directly calculated when the boundary segments of the evader's safe-reachable set are **explicitly** parametrized. Suppose there exists a function $g_i(\theta, e, p_i)$ such that each boundary segment is

$$\partial V_{ei} = \{q \mid q = g_i(\theta, e, p_i) \; \theta \in \Theta_i\} \tag{5}$$

where the set Θ_i is found by locating the intersection with the other boundaries. Given g_i we can find the boundary normal and Jacobian needed in (4) directly. At a particular $\theta_0 \in \Theta_i$ we find the vector n_i that is normal to the boundary at $q = g_i(\theta_0, e, p_i)$ by solving $(\nabla_\theta \; g_i)^T n_i = 0$ (Figure 5). The boundary Jacobian $\frac{\partial q}{\partial p_i}$ can be found by partially differentiating g_i with respect to p_i. In 2D $\Theta_i \subset \mathbb{R}$ is an interval $[\underline{\theta_i}, \bar{\theta_i}]$ and $g_i(\theta, e, p_i)$ is a curve. Thus, in 2D the area flux due to each pursuer i and evader is given by the following proposition.

Proposition 1.

$$\begin{aligned}\frac{\partial A_e}{\partial p_i} &= \int_{g_i(\underline{\theta_i},e,p_i)}^{g_i(\bar{\theta_i},e,p_i)} \frac{\partial g_i}{\partial p_i}^T \hat{n}_i \; ds \\ \frac{\partial A_e}{\partial e} &= \sum_i^N \int_{g_i(\underline{\theta_i},e,p_i)}^{g_i(\bar{\theta_i},e,p_i)} \frac{\partial g_i}{\partial e}^T \hat{n}_i \; ds\end{aligned} \tag{6}$$

In 3D $\Theta_i \subset \mathbb{R}^2$ and $g_i(\theta, e, p_i)$ is a bounded surface or patch. The area flux formulation changes to a double integral over the boundary surface ∂V_{ei}.

Example 1. In order to demonstrate correctness, we show that with the explicit boundary method (6) we recover the "move-to-centroid" pursuit policy found in [1, 2] when the evader's safe-reachable set is a Voronoi cell with straight boundaries. In this case, the pursuers' control action is to move towards the centroid c_i of the shared boundary, and the area flux term is

$$\frac{\partial A}{\partial p_i} = -\frac{l_i}{||p_i - e||}(c_i - p_i).$$

Where l_i and c_i,

$$l_i = \int_{\partial V_{ei}} dq \quad \text{and} \quad c_i = \frac{1}{l_i}\int_{\partial V_{ei}} q \; dq,$$

are the length and centroid of the boundary between pursuer i and the evader, respectively.

Proof. In 2D each boundary is a straight line and can be explicitly parametrized [23] as

$$q_i(\theta) = m_i + s_i\theta, \tag{7}$$

where $m_i = \frac{e+p_i}{2}$ and s_i is the unit tangent to the boundary. The unit tangent is a simple function of the unit normal $n_i = Rs_i$ where R is a $\frac{\pi}{2}$ rotation matrix. The outward unit normal and boundary Jacobian expressions are

$$n_i = \frac{p_i - e}{||p_i - e||} = \frac{d_i}{||d_i||} \quad \text{and} \quad \frac{\partial q}{\partial p_i} = \frac{I}{2} + \frac{R - s_i d_i^T}{||d_i||}\theta,$$

which we can substitute directly into (6),

$$\begin{aligned}\frac{\partial A}{\partial p_i}^T &= \int_{\partial V_{ei}} n_i^T \frac{\partial q}{\partial p_i} \; dq = \int_{\partial V_{ei}} {n_i}^T \left[\frac{I}{2} + \frac{R - s_i d_i^T}{||d_i||}\theta\right] \; dq \\ = \int_{\partial V_{ei}} \frac{{n_i}^T}{2} + \frac{{n_i}^T R}{||d_i||}\theta \; dq &= \int_{\partial V_{ei}} \frac{1}{||d_i||}\left[\frac{d_i^T}{2} - s_i^T\theta - p_i^T + p_i^T\right] \; dq \\ = \frac{1}{||d_i||}\int_{\partial V_{ei}} -m_i^T - {s_i}^T\theta + p_i^T \; dq &= -\frac{1}{||d_i||}\int_{\partial V_{ei}} q^T - p_i^T \; dq \\ = -\frac{1}{||d_i||}(l_i c_i - l_i p_i)^T &= -\frac{l_i}{||p_i - e||}(c_i - p_i)^T.\end{aligned}$$

We can also recover the "move-to-centroid" policy in 3D with the appropriate boundary expression.

Area Flux: Implicit

The flux terms can be directly calculated when the boundary segments of the evader's safe-reachable set are **implicitly** expressed. For situations where a parametric closed form expression for the boundary is unavailable, an implicit boundary representation,

$$\partial V_{ei} = \{q \mid f_i(q, e, p_i) = 0\},$$

can be used. In this context we assume that we can not explicitly solve for q to create an explicit expression. Differentiating f_i with respect to p_i results in

$$\frac{df_i}{dp_i} = \frac{\partial f_i}{\partial q}^T \frac{\partial q}{\partial p_i} + \frac{\partial f_i}{\partial p_i}^T = 0.$$

Noting that $\frac{\partial f_i}{\partial q}$ is normal to the boundary, but not unit length, it follows that

$$n_i^T \frac{\partial q}{\partial p_i} = -\frac{\partial f_i}{\partial p_i}^T \left\| \frac{\partial f_i}{\partial q} \right\|^{-1}.$$

Integrating over the boundary yields the area flux.

Proposition 2.

$$\begin{aligned} \frac{\partial A_e}{\partial p_i} &= \int_{\partial V_{ei}} -\frac{\partial f_i}{\partial p_i} \left\| \frac{\partial f_i}{\partial q} \right\|^{-1} dq \\ \frac{\partial A_e}{\partial e} &= \sum_i^{n_p} \int_{\partial V_{ei}} -\frac{\partial f_i}{\partial e} \left\| \frac{\partial f_i}{\partial q} \right\|^{-1} dq \end{aligned} \tag{8}$$

Example 2. The implicit boundary method (8) recovers the "move-to-centroid" policy as well for the case where the evader's safe-reachable set is a Voronoi cell with straight boundaries. Pursuer's i area flux term is

$$\frac{\partial A}{\partial p_i} = -\frac{l_i}{||p_i - e||}(c_i - p_i)^T.$$

Proof. The boundary segment is implicitly given by

$$f_i(q, e, p_i) = ||q - e|| - ||q - p_i|| = 0. \tag{9}$$

Applying (8) to (9) and noting that $||q - p_i|| = ||q - e|| \ \forall q \in \partial V_{ei}$ we obtain

$$\frac{\partial A_e}{\partial p_i} = \int_{\partial V_{ei}} -\frac{(q - p_i)^T}{||q - p_i||} \left\| \frac{q - e}{||q - e||} - \frac{q - p_i}{||q - p_i||} \right\|^{-1} ds = \int_{\partial V_{ei}} -\frac{(q - p_i)^T}{||p_i - e||} \ ds$$

which yields the same result as the explicit boundary case.

Area Flux: Uncertainty Aware Boundary

We have presented two methods for calculating the area flux terms and showed that they recover a known result [1, 2]. Now we directly apply our method for finding the flux terms to the uncertainty aware boundaries. Specifically we apply the explicit boundary method (6) to the parametrized curve (3). An implicit formulation of the uncertainty aware boundary can also be found in the case where the uncertainty is spherical, but for brevity that case is not discussed in detail here. The area of the uncertainty aware safe-reachable set $A(V_E)$ is now a function of the pursuers' positions p_i and the estimate $E(\mu, \Sigma)$. We show that the evader's safe-reachable set is non-increasing. First we show that area flux terms have the following symmetric relationship.

Lemma 1. *For the uncertainty aware boundary, the following relationship holds between the area flux terms:*

$$\frac{\partial A_E}{\partial \mu} = \sum_i^N -\frac{\partial A_E}{\partial p_i}.$$

Proof. First we note that given the form of the boundary in (3)

$$\frac{\partial q_i}{\partial \mu} = I + -\frac{\partial q_i}{\partial p_i}.$$

For each boundary i we can calculate the partials required by (1). Since the boundaries are described by explicit curves like (5) we can directly apply (6),

$$\begin{aligned}\frac{\partial A_E}{\partial \mu} &= \sum_i^N \int_{\partial V_{Ei}} (I + -\frac{\partial q_i}{\partial p_i})^T n_i ds &&= \sum_i^N \int_{\partial V_{Ei}} \hat{n_i} ds - \frac{\partial A_E}{\partial p_i} \\ &= \oint_{\partial V_E} \hat{n_i} ds + \sum_i^N -\frac{\partial A_E}{\partial p_i} &&= \sum_i^N -\frac{\partial A_E}{\partial p_i}.\end{aligned}$$

The contour integral vanishes since its integrand is a conservative (constant) vector field.

In between measurement updates, since the estimate ellipsoid translates no faster than the pursuers, the area of the uncertainty aware safe-reachable set is non-increasing.

Theorem 1. *If the ellipsoidal estimate does not translate faster than the pursuers, then there exists u_i where $\dot{A}_E \leq 0$.*

Proof. Let $z_i = \frac{\partial A_E}{\partial p_i}$, then from Lemma 1 we rewrite (1) as

$$\begin{aligned}\underset{u_i}{\text{minimize}}\ \underset{v}{\text{maximize}} \quad & \dot{A}_e = \sum_i^N z_i^T \dot{p}_i - z_i^T \dot{\mu} \\ \text{subject to} \quad & \dot{p_i} = u_i,\ ||u_i|| \leq 1\ \forall i \\ & \dot{\mu} = v,\ ||v|| \leq 1.\end{aligned}$$

From (2) we substitute the optimal polices therefore

$$\dot{A}_e^* = \sum_i^N -||z_i|| + \left\|\sum_i^N z_i\right\|.$$

From the triangle inequality it follows that $\sum_i^N ||z_i|| \geq \left\|\sum_i^N z_i\right\|$, thus $\dot{A}_e^* \leq 0$.

The "move-to-centroid" interpretation of $\frac{\partial A_E}{\partial p_i}$ no longer applies due to the curved boundaries and there may be configurations where $u_i = v \ \forall i$, resulting in no area change as the configuration of agents simply translates. This results in a "containment" scenario where A_E remains constant. Figure 7 shows how the uncertainty affects the pursuers' control action.

Fig. 7. The pursuers' control action (cyan) is affected by their uncertainty in the evader's potions. In the no uncertainty case (left) the pursuers move towards the evader, however with uncertainty (right) some control actions change significantly.

4 Set-Membership Filters

As is common with most estimation schemes, we assume our estimate is updated by a measurement received at discrete times $t_j = j\Delta t$ where Δt is some time step. During each estimation update, the pursuers jointly share the measurement $y(t_j)$ at time t_j that is corrupted by the bounded noise $w(t_j)$,

$$y(t_j) = e(t_j) + w(t_j),$$

where $e(t_j)$ is the true position of the evader and $w(t_j) \sim \Omega(0, W(t_j))$ at time t_j. We define $\Omega(0, W(t_j))$ to be an arbitrary probability distribution over the ellipsoidal support $\{x \mid x^T W(t_j)^{-1} x \leq 1\}$, where $W(t_j)$ is analogous to a time dependent measurement noise. Since the evader's velocity is bounded, we can represent its action by a disturbance proportional to the time step, i.e. drawn from the set $\{x \mid x^T x \leq \Delta t^2\}$. We also assume that our measurement prior $(\mu(0), \Sigma(0))$ contains the evader's initial position $e(0)$,

$$e(0) \in E(\mu(0), \Sigma(0)).$$

Using a Set-Membership filter [4], a variant of the Kalman filter designed for bounded disturbance and measurement noise models, we can employ a predict-update like process to update the estimate given the measurement. These filters use a prior ellipsoidal estimate, disturbance bound, and a noise corrupted measurement to generate a posterior ellipsoidal estimate $(\mu(t_{j+1}), \Sigma(t_{j+1}))$ that guarantees the quantity being measured, in this case the evader's position, lies within the new ellipsoid [24–26],

$$e(t_{j+1}) \in E(\mu(t_{j+1}), \Sigma(t_{j+1})).$$

5 Simulations

Combining the pursuer control policy discussed in Section 3 and using the filter described in [26], we tested our method in a 2D planar environment with four pursuers and a single evader. We conducted simulations to validate our method by generating 100 feasible scenarios where positions, initial measurements, and uncertainties were all randomly generated. In order to simulate a realistic situation, the magnitude of the measurement noise was proportional to the smallest distance between the pursuers and evader. Each agent, modeled as a point, had the same maximum speed of 1 unit per time step and r_{cap} was set to 5 units. Figure 8 shows the distance between the evader and its closest pursuer and Figure 9 shows the maximum eigenvalue of Σ at each time step. Both quantities were normalized against the initial conditions and the averages over all runs are plotted in black. The pursuer policy was tested against 3 different evader strategies, which are aggregated in Figures 8 and 9.

- Area maximizing policy with perfect position information (red)
- Move towards the largest gap between pursuers (green)
- Random walk (blue)

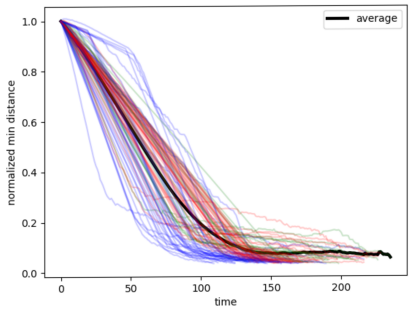

Fig. 8. The smallest distance between the evader and the pursuers vs. time for each simulation, normalized to the initial distance. The average is plotted in black.

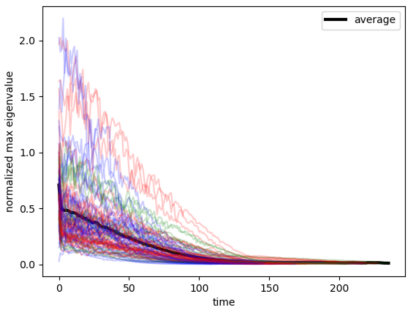

Fig. 9. The largest eigenvalue of Σ for each simulation vs. time for each simulation, normalized to the initial value. The average is plotted in black.

Figure 10 shows one instance of a batch of simulations.[6]

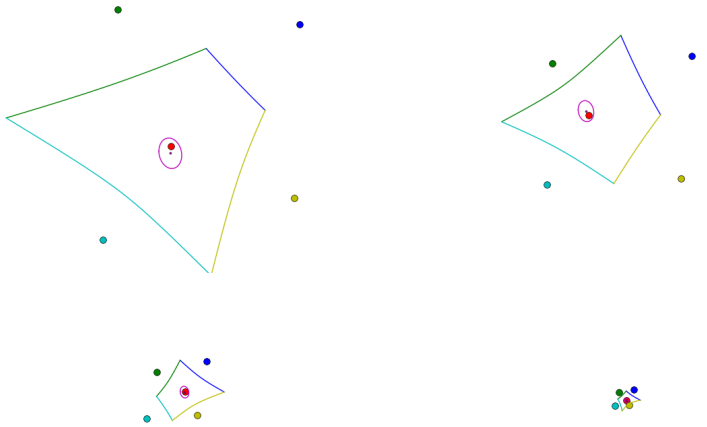

Fig. 10. Snap shots of four time steps of one simulation.

6 Conclusion

This work presents a pursuit policy using an area-minimization approach for pursuit-evasion games in unbounded environments where the pursuers maintain a noisy estimate of the evader's position. This necessitated the need for new methods to calculate the area flux terms as current methods are insufficient. By extending the idea of safe-reachability to account for the uncertainty in the evader's position we are able to determine a decentralized pursuit policy. Simulations were conducted to validate the approach as well as observe how uncertainty can affect the pursuit policy. Future extensions of this work will consider other Set-Membership filters and the effects of each pursuer having an independent estimate. We also plan to study the effects of obstacles have on safe-reachable boundaries.

Acknowledgements This work was supported in part by Ford Motor Company. We are grateful for this support.

[6] Videos of the simulations can be found here, https://youtu.be/3GUlEf0fjyc

References

1. H. Huang, W. Zhang, J. Ding, D. M. Stipanovi, and C. J. Tomlin, "Guaranteed decentralized pursuit-evasion in the plane with multiple pursuers," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 4835–4840, Dec 2011.
2. A. Pierson, Z. Wang, and M. Schwager, "Intercepting rogue robots: An algorithm for capturing multiple evaders with multiple pursuers," *IEEE Robotics and Automation Letters*, vol. 2, pp. 530–537, April 2017.
3. Z. Zhou, W. Zhang, J. Ding, H. Huang, D. M. Stipanović, and C. J. Tomlin, "Cooperative pursuit with voronoi partitions," *Automatica*, vol. 72, pp. 64–72, 10 2016.
4. D. Bertsekas and I. Rhodes, "Recursive state estimation for a set-membership description of uncertainty," *IEEE Transactions on Automatic Control*, vol. 16, pp. 117–128, Apr 1971.
5. D. W. Oyler, *Contributions to Pursuit-Evasion Game Theory*. PhD thesis, University of Michigan, 2016.
6. R. E. Allen, A. A. Clark, J. A. Starek, and M. Pavone, "A machine learning approach for real-time reachability analysis," in *Intelligent Robots and Systems*, pp. 2202–2208, IEEE, 2014.
7. S. Pan, H. Huang, J. Ding, W. Zhang, D. M. S. vi, and C. J. Tomlin, "Pursuit, evasion and defense in the plane," pp. 4167–4173, June 2012.
8. T. H. Chung, G. A. Hollinger, and V. Isler, "Search and pursuit-evasion in mobile robotics," *Autonomous Robots*, vol. 31, pp. 299–316, nov 2011.
9. R. Vidal, O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, "Probabilistic pursuit-evasion games: Theory, implementation, and experimental evaluation," *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 662–669, oct 2002.
10. S. LaValle and J. Hinrichsen, "Visibility-based pursuit-evasion: the case of curved environments," *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 196–202, apr 2001.
11. B. P. Gerkey, S. Thrun, and G. Gordon, "Visibility-based pursuit-evasion with limited field of view," *International Journal of Robotics Research*, vol. 25, no. 4, pp. 299–315, 2004.
12. R. Isaacs, *Differential Games*. New York: Courier Corporation, 1967.
13. T. Basar and G. Olsder, *Dynamic Noncooperative Game Theory*. Classics in Applied Mathematics, Society for Industrial and Applied Mathematics, 1999.
14. H. Huang, J. Ding, W. Zhang, and C. J. Tomlin, "A differential game approach to planning in adversarial scenarios: A case study on capture-the-flag," in *2011 IEEE International Conference on Robotics and Automation*, pp. 1451–1456, May 2011.
15. D. M. Stipanović, A. Melikyan, and N. Hovakimyan, "Guaranteed strategies for nonlinear multi-player pursuit-evasion games," *International Game Theory Review*, vol. 12, no. 01, pp. 1–17, 2010.
16. M. Santos, Y. Diaz-Mercado, and M. Egerstedt, "Coverage control for multirobot teams with heterogeneous sensing capabilities," *IEEE Robotics and Automation Letters*, vol. 3, pp. 919–925, April 2018.
17. F. Anton, D. Mioc, and C. Gold, "The Voronoi Diagram of Circles and Its Application to the Visualization of the Growth of Particles," pp. 20–54, Springer, Berlin, Heidelberg, 2009.
18. A. Okabe, *Spatial tessellations : concepts and applications of Voronoi diagrams*. Wiley, 2000.

19. I. Z. Emiris, E. P. Tsigaridas, and G. M. Tzoumas, "The predicates for the Voronoi diagram of ellipses," in *Proceedings of the twenty-second annual symposium on Computational geometry - SCG '06*, (New York, New York, USA), p. 227, ACM Press, 2006.
20. R. Olfati-Saber, "Distributed kalman filter with embedded consensus filters," in *Proceedings of the 44th IEEE Conference on Decision and Control*, pp. 8179–8184, Dec 2005.
21. T. Ngeli, C. Conte, A. Domahidi, M. Morari, and O. Hilliges, "Environment-independent formation flight for micro aerial vehicles," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1141–1146, Sept 2014.
22. R. Tron, R. Vidal, and A. Terzis, "Distributed pose averaging in camera networks via consensus on se(3)," in *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pp. 1–10, Sept 2008.
23. Y. Kantaros and M. M. Zavlanos, "Communication-aware coverage control for robotic sensor networks," in *53rd IEEE Conference on Decision and Control*, pp. 6863–6868, Dec 2014.
24. C. Durieu, É. Walter, and B. Polyak, "Multi-Input Multi-Output Ellipsoidal State Bounding," *Journal of Optimization Theory and Applications*, vol. 111, pp. 273–303, nov 2001.
25. Y. Becis-Aubry, M. Boutayeb, and M. Darouach, "A stable recursive state estimation filter for models with nonlinear dynamics subject to bounded disturbances," in *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 1321–1326, Dec 2006.
26. Y. Liu, Y. Zhao, and F. Wu, "Ellipsoidal state-bounding-based set-membership estimation for linear system with unknown-but-bounded disturbances," *IET Control Theory & Applications*, vol. 10, pp. 431–442, feb 2016.