Guidance Laws for Partially-Observable Interception Based on Linear Covariance Analysis

Jasper Arneberg* Ezra Tal* Sertac Karaman

***Abstract*— We consider pursuit-evasion games in which the pursuer is tasked with intercepting the evader using only partial measurements. Motivated by the utilization of visual sensing on board the pursuer, we focus on the case when only bearing measurements are available to the pursuer. The resulting partially-observable interception problem is computationally challenging, and the separation principle does not hold in general. In this paper, we identify a set of maneuvers that improve observability, and we propose an algorithm that utilizes these maneuvers to move the pursuer so that the expected payoff of the differential game is maximized. The algorithm uses in-the-loop uncertainty propagation based on linear covariance analysis to assess the effect of the maneuvers. We evaluate the resulting guidance law in experiments involving a quadcopter in flight representing the pursuer, and a simulated evader.**

I. Introduction

Recent years have witnessed a surge in personal ownership of unmanned aerial vehicle (UAV) systems. Today, off-the-shelf multicopters are widely available, and consumers are able to fly them without any prior training. UAVs flown by amateur pilots may cause dangers, *e.g.*, in the proximity of large crowds and near airports. These concerns are exacerbated in the presence of UAVs that are purposely deployed for surveillance, disturbance, or other malicious activities. In light of these fears, several countries have started to implement more rigorous regulations. However, enforcement of these regulations is not straightforward, especially towards eliminating the risk of adversarial UAVs [1]. Several measures to achieve this goal have been considered, *e.g.*, radio frequency jammers and ground-based net launchers [2].

We are motivated by the development of autonomously pursuing UAVs that can intercept and capture their adversarial evading counterparts. We define interception as the nulling of the relative distance between the pursuer and its adversary. After interception takes place, the adversary can be captured, *e.g.*, through collision or by using a catching mechanism [3]. We model the interception task as a pursuit-evasion differential game, so that the resulting guidance law is a *minmax* optimal control law, meaning that it achieves the optimal performance against the worst-case evasive action [4]. More specifically, we are interested in the case where the pursuer can acquire only measurements of the evader's relative bearing. The scenario represents autonomous pursuit using a low-cost, lightweight UAV that is equipped with only a monocular camera and has no access to extraneous measurements, such as ground radar.

Solving differential games is a challenging task. Most analytical solutions only apply to simple, *e.g.*, linear, problem instances [4]; and most numerical solutions are doomed by the curse of dimensionality [5]. When only partial measurements are available, as in the case we consider in this paper, the problem is further complicated from both an analytical and a computational point of view, and the notion of tracking observability is introduced [6], [7]. In most realistic scenarios where the pursuer can only obtain partial measurements of the state of the evader, the certainty equivalence principle does not hold for nonlinear control [8]; and exact solutions require working with infinitely many belief state variables, as both players need to maintain knowledge of what their opponent knows about what they know about what their opponent knows, *ad infinitum* [9]. Various control modifications that address bearings-only interception can be found in literature, *e.g.*, line-of-sight oscillations [10], filter covariance eigenvalue minimization [11], and nonlinear sliding mode control with variable damping [12].

For our proposed guidance law design, we divide the pursuit into three phases: In the first phase, which we call the *observability maneuvering* phase, the pursuer maneuvers to improve its estimation of the evader's relative position and velocity. In the second phase, which we call the *pre-terminal guidance* phase, the pursuer engages a guidance law based on a continuously-updated estimate of the relative position and velocity of the evader. The main objective of this phase is to get closer to the evader. In the third phase, which we call the *terminal guidance* phase, the pursuer and evader both engage in complex maneuvers.The main objective of this phase is to capture the evader.

In this paper, we focus on the first phase, namely the observability maneuvering phase. The goal of the observability maneuvering is to improve the expected game utility, *i.e.*, to increase the probability of successful interception by improving the quality of the estimate of the evader state. This is complicated by the fact that the system may be locally unobservable. Since there does not exist a global observability condition for nonlinear systems, we utilize the extended output Jacobian (EOJ) to analyze local observability and to propose maneuvers that improve the state estimate. The expected game utility is computed online using closed-loop uncertainty propagation. For this purpose, we utilize linear covariance analysis (LCA), which can be seen as a linearized version of Monte Carlo analysis (MCA). The LCA approximates nonlinear uncertainty propagation in a single

*The first two authors contributed equally.

J. Arneberg, E. Tal and S. Karaman are with the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA. {jaspera,eatal,sertac}@mit.edu

run by linearizing the dynamics, measurement, and control equations around a nominal trajectory [7], [13]. LCA is commonly applied to enable quick iterations in guidance, navigation, and control (GNC) design [13], [14]. We show that its computational efficiency makes it an excellent tool for online uncertainty propagation, as the pursuer assesses the expected value of the game after a certain maneuver.

This paper has three main contributions. First, we derive certain maneuvers, based on an analysis of the extended output Jacobian, that improve the expected value of the game. Second, we propose an uncertainty propagation method, based on the linear covariance analysis, that allows us to estimate the expected value of the game in a computationally efficient manner. Third, we present a practical framework on guidance for partially-observable interception, and we validate the guidance law design using experiments involving a quadrotor in flight pursuing a simulated evader.

The paper is organized as follows. In Section II, we formalize the differential game and analyze its local observability. An overview of the guidance law including observability maneuvers and LCA is given in Section III. Section IV contains computational and real-world experimental results. Finally, conclusions are presented in Section V.

II. Interception Dynamics and Observability

We consider the interception scenario shown in Fig. 1. The pursuing quadrotor aims to intercept the adversarial aircraft, *i.e.*, the evader, using relative angle measurements. We assume that the quadrotor has accurate knowledge of its own state. Hence, we focus on the estimation of the relative position and velocity of the evader.

A. Pursuit-Evasion Dynamics

The evader is modeled using Dubins vehicle kinematics [15]. It performs smooth turns at constant altitude. Since the pursuer velocity is assumed known, it is considered an input to the interception dynamics. The resulting system of relative dynamics is given by

$$\begin{aligned} \dot{x}_1 &= u_1 - V_e \cos\chi, \\ \dot{x}_2 &= u_2 + V_e \sin\chi, \\ \dot{x}_3 &= u_3, \\ \dot{\chi} &= v_\chi, \end{aligned} \tag{1}$$

where the state $\mathbf{x} = [x_1\ x_2\ x_3\ \chi]^T \in \mathbb{R}^3 \times \mathbb{T}$ consists of the evader position with regard to the pursuer in the inertial reference frame, and the evader heading angle χ. The symbol $\mathbb{T}$ denotes the circle group. The known pursuer velocity vector is indicated by $\mathbf{u} = [u_1\ u_2\ u_3]^T \in \mathbb{R}^3$. The evader control input $v_\chi \in \mathbb{R}$ corresponds to its heading rate, and $V_e \in \mathbb{R}$ is the constant evader speed. For convenience of notation, we define the following vectors:

$$\mathbf{R} = [x_1\ x_2\ x_3]^T, \tag{2}$$

$$\mathbf{R_h} = [x_1\ x_2\ 0]^T, \tag{3}$$

$$\mathbf{V_e} = [-V_e \cos\chi\ V_e \sin\chi\ 0]^T. \tag{4}$$

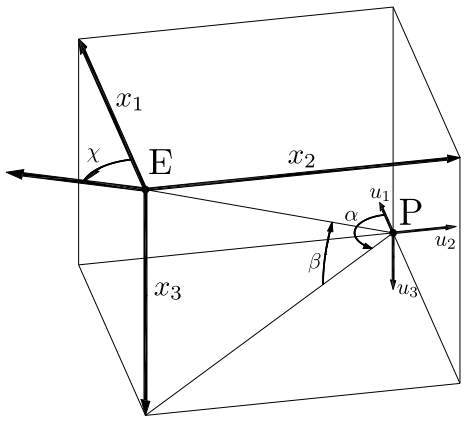

Fig. 1: Pursuit-evasion scenario overview.

An unscented Kalman filter (UKF) based on (1) is used for state estimation during the observability maneuvering and pre-terminal guidance phases. During these phases, it is assumed that no strong evasive maneuvering is performed by the evader. This justifies the constant-altitude Dubins vehicle model. The evader input is modeled as white noise. For an overview of the UKF equations, the reader is referred to [16].

B. Nonlinear Observability

In this section, we consider the observability of the nonlinear dynamics system. We derive necessary and sufficient conditions for local weak observability.

The pursuer obtains the measurement vector $\mathbf{y} = [\alpha\ \beta]^T \in \mathbb{T} \times [-\frac{\pi}{2}, \frac{\pi}{2}]$ from its omnidirectional sensor. The azimuth angle α and attitude angle β are defined as

$$\begin{aligned} \alpha &= h_1(\mathbf{x}) = \arctan2(-x_2, x_1), \\ \beta &= h_2(\mathbf{x}) = \arcsin\left(\frac{x_3}{R}\right) \end{aligned} \tag{5}$$

with $\arctan2$ the multi-valued inverse tangent function, and $R = \|\mathbf{R}\|$ the relative Euclidean distance between pursuer and evader. By combining these two measurements, the pursuer is able to determine the relative direction of the evader. However, the evader orientation χ and distance R are not directly observable.

In contrast to the case where linear measurements and linear differential equations are considered, global observability cannot be established for the scenario considered here. Instead, we use the notion of *local weak observability* through the extended output Jacobian, as introduced in [17]. Intuitively, a state $\mathbf{x_0}$ is said to be locally weakly observable if it can be distinguished from its neighbors instantaneously [17].

The EOJ provides a sufficient observability rank condition for local observability at a state $\mathbf{x_0}$. The nonlinear system consisting of the kinematic model (1), and the measurement equations (5) is locally weakly observable at $\mathbf{x_0}$, if the EOJ matrix with rows

$$\mathbf{J} = \{ \left.\frac{\partial h_i^{j-1}}{\partial \mathbf{x}}\right|_{\mathbf{x}=\mathbf{x_0}} | i = 1, 2; j = 1, \ldots, 4\} \tag{6}$$

is full rank [18]. The superscript j denotes the j-th temporal derivative of the measurement function. Full rankness of at least one 4×4 submatrix of the 8×4 EOJ is sufficient to guarantee local observability [18]. Expressions for the higher-order derivatives of the measurement functions are mathematically quite involved and not suitable for the formulation of a usable observability condition. Moreover, in practice when measurements are noisy, these higher-order derivatives can rarely be obtained with sufficient accuracy, therefore we consider the following submatrix:

$$\mathbf{J'} = \begin{bmatrix} \frac{\partial h_1}{\partial x_1} & \frac{\partial h_1}{\partial x_2} & \frac{\partial h_1}{\partial x_3} & \frac{\partial h_1}{\partial \chi} \\ \frac{\partial \dot{h}_1}{\partial x_1} & \frac{\partial \dot{h}_1}{\partial x_2} & \frac{\partial \dot{h}_1}{\partial x_3} & \frac{\partial \dot{h}_1}{\partial \chi} \\ \frac{\partial h_2}{\partial x_1} & \frac{\partial h_2}{\partial x_2} & \frac{\partial h_2}{\partial x_3} & \frac{\partial h_2}{\partial \chi} \\ \frac{\partial \dot{h}_2}{\partial x_1} & \frac{\partial \dot{h}_2}{\partial x_2} & \frac{\partial \dot{h}_2}{\partial x_3} & \frac{\partial \dot{h}_2}{\partial \chi} \end{bmatrix} \tag{7}$$

with

$$\det(\mathbf{J'}) = V_e \frac{(x_1 \cos\chi - x_2 \sin\chi) u_3}{\sqrt{x_1^2 + x_2^2} R^2} + V_e \frac{x_3 (V_e - u_1 \cos\chi + u_2 \sin\chi)}{\sqrt{x_1^2 + x_2^2} R^2}. \tag{8}$$

We note that the sufficient rank observability condition may require that the inputs are known. However, in this case the evader input does not appear in the measurement equations or in their temporal derivatives, and thus does not affect observability. Interpretation of (8) allows us to formulate practical conditions that render the system unobservable. We consider four events corresponding to the factors of the two terms of (8):

(e-i) $(x_1 \cos\chi - x_2 \sin\chi) = 0 \iff \mathbf{V_e} \perp \mathbf{R_h}$,
(e-ii) $u_3 = 0 \iff \dot{x}_3 = 0$,
(e-iii) $V_e - u_1 \cos\chi + u_2 \sin\chi = 0 \iff \dot{\mathbf{R}}_\mathbf{h} \perp \mathbf{V_e}$,
(e-iv) $x_3 = 0$

The intersection of events between terms corresponds to singularity of $\mathbf{J'}$, *e.g.*, (e-i)$\wedge$(e-iii) $\implies \det(\mathbf{J'}) = 0$. In this manner, the following singularity conditions at which observability is not guaranteed can be constructed: (e-i)$\wedge$(e-iii), (e-i)$\wedge$(e-iv), (e-ii)$\wedge$(e-iii), and (e-ii)$\wedge$(e-iv). Additionally, there is a fifth event that directly causes singularity in itself:

(e-v) $V_e(x_1 \cos\chi - x_2 \sin\chi) u_3 + V_e x_3 (V_e - u_1 \cos\chi + u_2 \sin\chi) = 0$,

which can be rewritten as

$$V_e \frac{x_1 \cos\chi - x_2 \sin\chi}{x_3} = V_e \frac{-V_e + u_1 \cos\chi - u_2 \sin\chi}{u_3} \tag{9}$$

or

$$\frac{-\mathbf{V_e}^T \mathbf{R_h}}{x_3} = \frac{-\mathbf{V_e}^T \dot{\mathbf{R}}_\mathbf{h}}{u_3}. \tag{10}$$

The singularity events present a clear geometric interpretation of when observability is lost, and are visually depicted in Fig. 2. By (8), there exists a non-empty subset of the state-space for which the EOJ is singular, consequently the rank observability condition is sufficient and necessary [18]. Therefore the system is *not* weakly locally observable when the singularity conditions are met.

III. Vision-Based Guidance Using Linear Covariance Analysis

In this section, we detail the various aspects of the vision-based guidance law design. The interception guidance law consists of three phases.

The first, or *observability maneuvering*, phase is the main focus of our analysis in this paper. During this phase, the distance between pursuer and evader is still relatively large. The pursuer performs maneuvers to increase the quality of its relative state estimate.

The second, or *pre-terminal guidance*, phase starts when the pursuer deems its relative state estimate of sufficient quality to start interception. A simple guidance law, such as deviated pure pursuit, may be used during this phase, because the relative distance is still large and therefore no evasive action is expected.

The third and final phase consists of *terminal guidance*. In this phase a differential games-based guidance law is employed to achieve a guaranteed miss distance despite possible aggressive evasive action. This controller is based on state feedback and thus requires a sufficiently accurate state estimate to achieve satisfactory performance.

A. Linear Covariance Analysis

The goal of the observability maneuvering phase is to improve the relative state estimate. The phase is concluded when the state estimate variance is sufficiently low, such that the pursuer can successfully execute the pre-terminal and terminal guidance laws based on the estimated state. In this context, the probability of success is given by the probability of capture, *i.e.*, the probability that the condition $R \leq r_c$, with r_c the capture radius, is satisfied during the terminal guidance phase.

The transition to the pre-terminal guidance phase can be formulated as an uncertainty propagation problem, where the transition depends on the probability of interception as a function of the current state estimate and its covariance. The uncertainty propagation needs to take into account the system dynamics, measurement equations, and guidance law. If dynamics and measurement equations are nonlinear, Monte Carlo analysis is typically used to address the closed-loop uncertainty propagation. However, MCA may require several hundreds or thousands of samples, making online computation infeasible [19].

To enable real-time uncertainty propagation, we resort to linear covariance analysis. This method can be considered a linearized version of Monte Carlo analysis, where nonlinear dynamics, measurement, and control models are linearized around a mean reference trajectory. Granted the validity of the linearization, LCA is able to obtain the same statistics as MCA at a fraction of the computational cost.

The LCA algorithm consists of two stages: First, a single noiseless nonlinear simulation run is performed to obtain the nominal trajectory. Next, the Gaussian state variance is propagated along this nominal trajectory, using the linearized dynamics, measurement, and control models. The reader is referred to [13] for further details on linear covariance

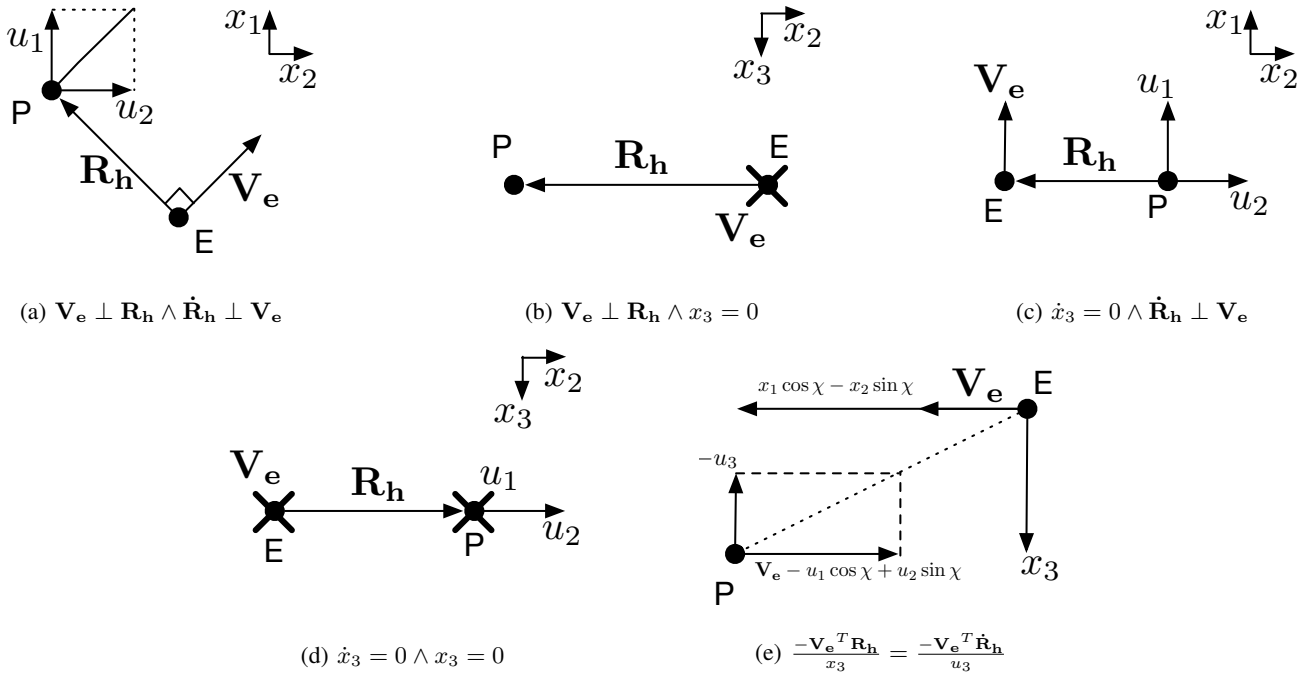

(a) $\mathbf{V_e} \perp \mathbf{R_h} \wedge \dot{\mathbf{R}}_\mathbf{h} \perp \mathbf{V_e}$ (b) $\mathbf{V_e} \perp \mathbf{R_h} \wedge x_3 = 0$ (c) $\dot{x}_3 = 0 \wedge \dot{\mathbf{R}}_\mathbf{h} \perp \mathbf{V_e}$

(d) $\dot{x}_3 = 0 \wedge x_3 = 0$ (e) $\frac{-\mathbf{V_e}^T \mathbf{R_h}}{x_3} = \frac{-\mathbf{V_e}^T \dot{\mathbf{R}}_\mathbf{h}}{u_3}$

Fig. 2: Geometric representation of singularity conditions of $\mathbf{J'}$. Using the geometric interpretation, we can formulate maneuvers to avoid the singularity conditions. Singularities are caused by cancellation of measurements, as follows: (a) $\dot{\mathbf{R}}_\mathbf{h} \parallel \mathbf{R_h}$, so $\mathbf{R_h}$ does not rotate. Hence, $\dot{\alpha} = 0$. (b) $x_3 = 0$, so $\beta = 0$; $\dot{\alpha} \neq 0$, but does not provide sufficient information to estimate both R_h and χ, if $\beta = 0$. (c) Inspection of $\mathbf{J'}$ shows that under this condition the measurement derivatives only have non-zero partials with regard to χ, leaving three states to be observed with the two angular measurements. (d) $x_3 = \dot{x}_3 = 0$, so that $\beta = \dot{\beta} = 0$. (e) The pursuer is climbing straight towards the evader, leading to $\dot{\beta} = 0$.

analysis. For brevity purposes, we will only state the current implementation in this section.

The LCA truth model is expressed as

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) + \mathbf{B}\mathbf{w} \tag{11}$$

with f given by (1), $\mathbf{B}$ the process noise gain matrix designed to represent model uncertainty due to unknown evader control input, and $\mathbf{w}$ continuous Gaussian process noise with spectral density $\mathbf{S_w}$. The control input is computed according to some control law

$$\mathbf{u} = g(\hat{\mathbf{x}}), \tag{12}$$

and the system measurement vector is obtained as

$$\tilde{\mathbf{y}} = h(\mathbf{x}) + \boldsymbol{\nu} \tag{13}$$

with h given by (5), and $\boldsymbol{\nu}$ a Gaussian measurement noise vector with covariance $\mathbf{R}_{\boldsymbol{\nu}}$. LCA is based on a closed-loop system with extended Kalman filter (EKF). Incorporation of the UKF that is used for state estimation is not straightforward, as it is based on nonlinear propagation equations.

The nominal nonlinear simulation run is performed according to (11), (12), and (13). The run is initialized using the UKF state estimate, and $\mathbf{S_w}$ and $\mathbf{R}_{\boldsymbol{\nu}}$ are set to zero. The resulting trajectory is considered the nominal trajectory $\bar{\mathbf{x}}$. Due to its idealized noiseless nature, the nominal trajectory virtually always terminates with perfect capture of the evader, *i.e.*, $R_f = 0$ up to numerical precision with R_f the distance between pursuer and evader at termination of the guidance law. We are interested in uncertainty propagation to determine the expected dispersion of the state with regard to the nominal trajectory. Using the dispersion distribution, we are able to calculate confidence intervals, and thus the probability of capture for some capture radius r_c.

For the main LCA run, the state and covariance propagation are linearized along the nominal trajectory. We will first give the nonlinear equations and then show the linearized system that is used in LCA.

The EKF propagation equations for the estimated state and its covariance are

$$\dot{\hat{\mathbf{x}}} = f(\hat{\mathbf{x}}, \mathbf{u}) \tag{14}$$

and

$$\dot{\mathbf{P}} = \mathbf{F}_{\hat{\mathbf{x}}}\mathbf{P} + \mathbf{P}{\mathbf{F}_{\hat{\mathbf{x}}}}^T + \mathbf{S_w} \tag{15}$$

with $\mathbf{F}_{\hat{\mathbf{x}}} = \left.\frac{\partial f(\mathbf{x},\mathbf{u})}{\partial \mathbf{x}}\right|_{\mathbf{x}=\hat{\mathbf{x}}}$. When a measurement is received the estimated state and covariance are updated according to

$$\hat{\mathbf{x}}^+ = \hat{\mathbf{x}}^- + \mathbf{K}(\tilde{\mathbf{y}} - h(\hat{\mathbf{x}}^-)) \tag{16}$$

with $\mathbf{K}$ the Kalman gain (see, *e.g.*, [16]), and

$$\mathbf{P}^+ = (\mathbf{I} - \mathbf{K}\mathbf{H}_{\hat{\mathbf{x}}}\mathbf{P}^-(\mathbf{I} - \mathbf{K}\mathbf{H}_{\hat{\mathbf{x}}})^T + \mathbf{K}\mathbf{R}_{\boldsymbol{\nu}}\mathbf{K}^T \tag{17}$$

with $\mathbf{H}_{\hat{\mathbf{x}}} = \left.\frac{\partial h(\mathbf{x})}{\partial \mathbf{x}}\right|_{\mathbf{x}=\hat{\mathbf{x}}}$.

We define the augmented state

$$\mathbf{X} = \begin{bmatrix} \delta\mathbf{x} \\ \delta\hat{\mathbf{x}} \end{bmatrix}, \tag{18}$$

consisting of the true dispersion $\delta\mathbf{x}$ of the actual state with regard to the nominal trajectory $\bar{\mathbf{x}}$, and the navigation dispersion $\delta\hat{\mathbf{x}}$ of the filter state with regard to the nominal trajectory. The update equations are given by

$$\dot{\mathbf{X}} = \mathcal{F}\mathbf{X} + \mathcal{W}\mathbf{w}, \tag{19}$$

$$\mathbf{X}^{+} = \mathcal{A}\mathbf{X}^{-} + \mathcal{D}\boldsymbol{\nu} \tag{20}$$

with

$$\mathcal{F} = \begin{bmatrix} \mathbf{F}_{\bar{\mathbf{x}}} & \mathbf{F}_{\mathbf{u}}\mathbf{G}_{\hat{\mathbf{x}}} \\ \mathbf{0} & \mathbf{F}_{\hat{\mathbf{x}}} + \mathbf{F}_{\mathbf{u}}\mathbf{G}_{\hat{\mathbf{x}}} \end{bmatrix}, \tag{21}$$

$$\mathcal{W} = \begin{bmatrix} \mathbf{B} \\ \mathbf{0} \end{bmatrix}, \tag{22}$$

$$\mathcal{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{K}\mathbf{H}_{\bar{\mathbf{x}}} & \mathbf{I} - \mathbf{K}\mathbf{H}_{\hat{\mathbf{x}}} \end{bmatrix}, \tag{23}$$

$$\mathcal{D} = \begin{bmatrix} \mathbf{B} \\ \mathbf{0} \end{bmatrix}, \tag{24}$$

and $\mathbf{G}_{\hat{\mathbf{x}}} = \left.\frac{\partial g(\mathbf{x})}{\partial \mathbf{x}}\right|_{\mathbf{x}=\hat{\mathbf{x}}}$. Since we are interested in the variance $\mathbf{C}_{\mathbf{A}} = \mathbb{E}\left[\mathbf{X}\mathbf{X}^{T}\right]$, we directly calculate it, as follows:

$$\dot{\mathbf{C}}_{\mathbf{A}} = \mathcal{F}\mathbf{C}_{\mathbf{A}} + \mathbf{C}_{\mathbf{A}}\mathcal{F}^{T} + \mathcal{W}\mathbf{S}_{\mathbf{w}}\mathcal{W}^{T}, \tag{25}$$

$$\mathbf{C}_{\mathbf{A}}^{+} = \mathcal{A}\mathbf{C}_{\mathbf{A}}^{-}\mathcal{A}^{T} + \mathcal{D}\mathbf{R}_{\boldsymbol{\nu}}\mathcal{D}^{T}. \tag{26}$$

We initialize $\mathbf{C}_{\mathbf{A}}$ using the current UKF state covariance. By integration of (25) and (26), the state dispersion variance of the closed-loop system can be propagated in a single run. By doing so, we are able estimate the probability of capture $\Pr_{cap} = \Pr(R_f < r_c)$, where R_f is taken at the termination time according to the nominal trajectory.

The distribution of R_f is to be obtained from the upper-left block of $\mathbf{C}_{\mathbf{A}}$. Note that this is the variance of the state vector as defined in (1), which must first be transformed to obtain the distribution of R_f. The transformation can be done either analytically considering the distribution of the Euclidean norm of the Gaussian vector $\mathbf{R}$ [20], or numerically, *e.g.*, using the unscented transform [16].

Based on the probability of capture, it is then decided to switch between observability maneuvers, or – when the probability is sufficiently large – to initiate the pre-terminal guidance phase. Details on the switching strategy are given in Section III-C.

B. Observability Maneuvers

If linear covariance analysis determines that the probability of interception is not sufficient to commence the pre-terminal guidance phase, this may indicate that the observability matrix is singular or may enter a singularity during the LCA trajectory. In this case, observability can be improved by moving the state away from the singularity. This is done using so-called observability maneuvers. Based on the singularity events defined in Section II, four maneuvering directions are defined: (m-i) in the direction of $\mathbf{V_e}$, (m-ii) perpendicular to $\mathbf{V_e}$ towards the evader in the horizontal plane, (m-iii) along x_3 towards the evader, and (m-iv) diagonally in the combined direction of (m-i)–(m-iii). The trajectories are planned with respect to the evader position and velocity as estimated by the UKF. Comparison to Fig. 2 shows how these maneuvers are able to resolve or avoid the singularity conditions:

- The horizontal maneuvers (m-i) and (m-ii) are perpendicular, therefore at least one will cause a rotation of $\mathbf{R_h}$ and will resolve the condition in Fig. 2a.
- The vertical maneuver (m-iii) and — to a lesser extent — the diagonal maneuver (m-iv) cause vertical separation and will thereby resolve the condition in Fig. 2b.
- The vertical maneuver (m-iii) and the diagonal maneuver (m-iv) will enforce a non-zero vertical separation speed. Maneuver (m-i) in the direction of $\mathbf{V_e}$ will avoid the event $\dot{\mathbf{R}}_\mathbf{h} \perp \mathbf{V_e}$. Hence, all of these maneuvers can be used to address the condition in Fig. 2c.
- The vertical and diagonal maneuvers (m-iii) and (m-iv), respectively, both guarantee non-zero vertical separation and speed, and thereby address the condition in Fig. 2d.
- Any maneuver that is not perpendicular to $\mathbf{V_e}$ or aimed directly at or away from the evader is able to resolve the condition in Fig. 2e. Hence, (m-i), (m-iii), and (m-iv) are generally appropriate.

The maneuver selection method is based on LCA and described in the next section.

C. Switching Strategy

The switching strategy forms the heart of the guidance law design. Based on the probability of capture obtained from the LCA, the next observability maneuver or phase transition is commanded. For comparison to a desired value $\Pr_{des}$, which is deemed sufficient to start the pre-terminal guidance phase, the probability of capture is computed using

- a single LCA instance under the pre-terminal guidance law, *e.g.*, deviated pure pursuit.

If the condition $\Pr_{cap} \geq \Pr_{des}$ is satisfied, the pre-terminal guidance phase is immediately started. If this condition is not satisfied, the guidance law first aims to improve the quality of the state estimate. For this purpose, the effects of different sequences of N observability maneuvers are assessed by

- 4^N LCA instances, each under a sequence of N maneuvers followed by the pre-terminal guidance law.

The maneuver sequence corresponding to the largest probability of capture is selected. During execution of the sequence, the switching algorithm continuous evaluation as described above, and may switch to the pre-terminal guidance phase or a different maneuver sequence. The resulting switching strategy is summarized in Algorithm 1.

We propose to evaluate sequences of N maneuvers, to prevent the selection of only locally optimal maneuvers. To reduce the computational load, the 4^N LCA instances can be executed in a tree-like fashion, where maneuver sequences with identical starting subsequences share intermediate LCA results.

D. Interception Guidance

Interception guidance consists of two phases: the pre-terminal guidance phase, and the terminal guidance phase.

Algorithm 1 Switching Strategy

while $Phase = ObservabilityManeuvering$ **do**
  $\Pr_{cap} \leftarrow \text{LCA}(PreTerminalGuidanceLaw)$
  if $\Pr_{cap} \geq \Pr_{des}$ **then**
    $Phase \leftarrow PreTerminalGuidance$
  else
    for $MS \in \{\text{(m-i)}, \ldots, \text{(m-iv)}\}^N$ **do**
      $\Pr_{cap}^{MS} \leftarrow \text{LCA}(MS)$
    end for
    Start execution of $\text{argmax}_{MS} \Pr_{cap}^{MS}$
  end if
end while

During the pre-terminal guidance phase, the distance between pursuer and evader is still relatively large. Hence, no aggressive evasive action is anticipated. In the face of only small evader maneuvers, simple guidance laws may perform adequately. In this case, we employ pure pursuit for pre-terminal guidance [21].

Terminal guidance employs a differential games-based guidance law that takes into account the pursuer's maximum accelerations in terms of the quadrotor attitude and thrust limits [22]. The control law is numerically obtained using approximate dynamic programming based on a compressed representation of the differential game value function. The terminal guidance law is highly nonlinear due to the presence of singular surfaces in the value function [23]. This makes it unsuitable for linear covariance analysis. Simulations have shown that the probability of capture obtained under pure pursuit is a conservative estimate compared to the differential games guidance law. Therefore, we maintain the pre-terminal guidance law up to termination for the purpose of LCA.

IV. Experimental Results

In this section, we present results from computational and real-world experiments. We verify the accuracy of linear covariance analysis in the pursuit-evasion scenario by comparison to Monte Carlo analysis, and present an evaluation of the integrated guidance law using a quadrotor in flight.

A. Computational Experiments

In this section, we compare statistics obtained by LCA to results from MCA. Both analyses are performed using the dynamics and measurement equations presented in Section II with the addition of Gaussian process and measurement noise. Figure 3 shows the state trajectories from a 1000-run MCA, and the 3-σ dispersal bounds obtained from LCA. It can be seen that the 3-σ bounds agree with the maximum state deviation of the MCA runs. This result is confirmed when we compare the statistics for the pursuer-evader distance at termination. In Table I, it can be seen that as the number of MCA runs increases, the 3-σ bound on the terminal miss distance converges to approximately the value obtained using LCA in a single run. The table also shows that the LCA result is obtained in a fraction of the computation time that is required to perform MCA. This reduction in computational burden enables in-the-loop uncertainty propagation using linear covariance analysis.

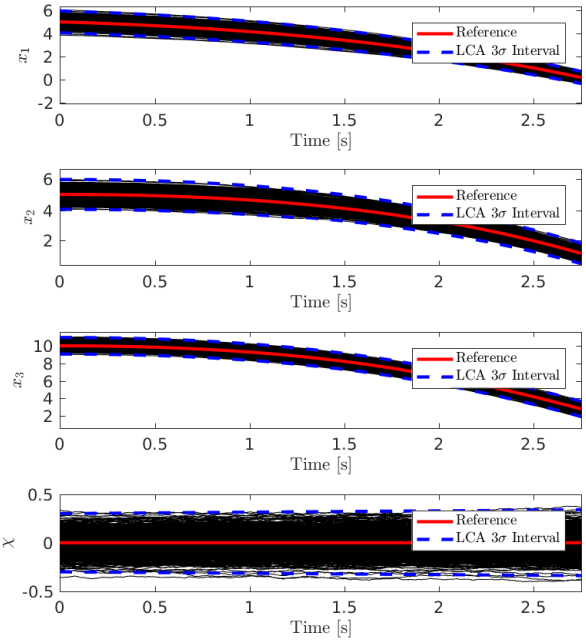

Fig. 3: Nominal trajectory, 1000 Monte Carlo trajectories, and state dispersion 3σ-bounds obtained from LCA.

TABLE I: Comparison of terminal miss distance 3σ-bounds, and computational time for Monte Carlo Analysis and LCA.

Monte Carlo Runs	10	100	1000	10,000	**LCA**
3-σ interval [m]	1.916	2.477	2.598	2.657	2.810
Computation time [s]	1.48	13.34	143.40	1331.8	0.6279

B. Flight Experiments

In this section, we demonstrate the guidance law with observability maneuvers in real-life flight using a quadrotor. While the quadrotor flies in a motion capture space, it is pursuing a simulated evader. An overview of the experimental setup is shown in Fig. 4. The evader dynamics are simulated according to the Dubins vehicle dynamics given in (1) with $V_e = 0.5$ m/s and added process noise with $\mathbf{S_w} = \text{diag}(\begin{bmatrix}10^{-4} & 10^{-4} & 10^{-4} & 10^{-3}\end{bmatrix})$ m^2/s and rad^2/s, respectively. The noisy bearing and attitude angle measurements are simulated using $\mathbf{R_\nu} = \text{diag}(\begin{bmatrix}10^{-4} & 10^{-4}\end{bmatrix})$ rad^2. The pursuer estimates the relative state using a UKF, which is initialized off the first measurements, with large state covariance and a large initial range estimate. The quadrotor is carrying an NVIDIA Jetson TX2 system-on-chip. All computation is performed in real-time using onboard resources. Maneuver sequences are generated with $N = 3$, and tracked using the inner-loop controller described in [24]. Pure pursuit guidance is used in the pre-terminal and terminal guidance phases.

Figure 5 shows the probability of capture as computed using LCA. The observability maneuvering phase is started at 2.6 s with a capture probability of 0.53. Through the observability maneuvers, this value is increased to the desired value of 0.95, so that the pre-terminal guidance phase starts at 4.2 s. The UKF range estimate error is shown as a function of the true range in Fig. 6. The red curve shows that a large

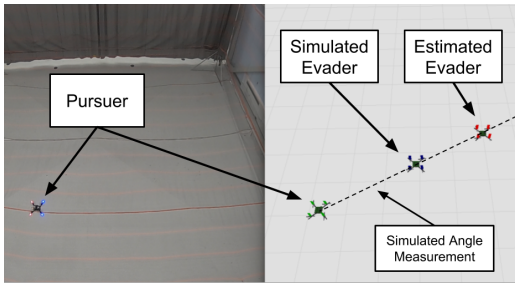

Fig. 4: Flight space (left), and simulated evader and measurement (right).

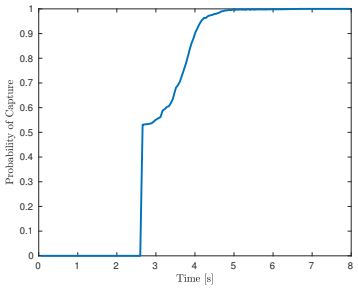

Fig. 5: LCA probability of capture.

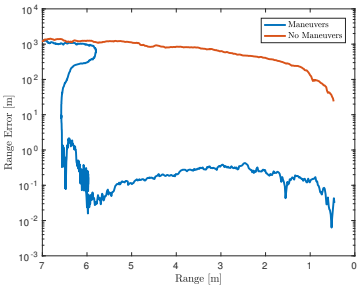

Fig. 6: UKF range estimate error with and without observability maneuvers.

error persists if the pre-terminal guidance phase is entered right away and no observability maneuvers are performed. In contrast, the blue curve first shows a reduction in the error through observability maneuvering, and then a reduction in true range — while maintaining small error — during the pre-terminal guidance phase. Consequently, the terminal guidance phase can be entered with a much better state estimate, which may greatly improve performance against an aggressively maneuvering evader.

V. Conclusions

In this paper, we used the extended output Jacobian to identify a set of maneuvers that improve observability. We proposed an algorithm that utilizes these maneuvers to move the pursuer, so that the expected payoff of the differential game is maximized. We established that linear covariance analysis is a valuable tool for real-time uncertainty propagation and verified its accurate application through computational experiments. Finally, we demonstrated that the observability maneuvers result in an improvement in range estimation accuracy in real-life flight using a pursuing quadcopter and a simulated evader.

ACKNOWLEDGMENT

Arneberg was supported by a Draper Laboratory Fellowship. The authors gratefully acknowledge early discussions with Gian Luca Mariottini of Draper Laboratory, which have contributed to the ideas presented in this paper. This work was financially supported in part by Lockheed Martin.

References

[1] "Small unmanned aircraft systems: FAA should improve its management of safety risks," Tech. Rep. GAO-18-110, US Government Accountability Office, Washington, DC, 2018.

[2] B. Yates, *MITRE Names C-UAS Challenge Winners*. The MITRE Corporation, Washington, DC, 2016.

[3] T. K. Horiuchi, M. Paras, T. Cervi, B. Oursler, S. Sanz, J. Gaus, D. Woodbury, I. Faruque, and H. Xu, "An autonomous, visually-guided, counter-sUAS aerial vehicle with net countermeasure," in *AIAA Atmospheric Flight Mechanics Conference*, pp. 3397–3406, 2016.

[4] T. Başar and G. Olsder, *Dynamic Noncooperative Game Theory*. SIAM, Philadelphia, 1999.

[5] R. Bellman, *Dynamic programming*. Princeton University Press, Princeton, 1957.

[6] A. N. Payne, "Observability conditions for angles-only tracking," in *Signals, Systems and Computers, 1988. Twenty-Second Asilomar Conference on*, pp. 451–457, IEEE, 1988.

[7] D. C. Woffinden and D. K. Geller, "Relative angles-only navigation and pose estimation for autonomous orbital rendezvous," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1455–1469, 2007.

[8] F. Mazenc, L. Praly, and W. Dayawansa, "Global stabilization by output feedback: examples and counterexamples," *Systems & Control Letters*, vol. 23, no. 2, pp. 119–125, 1994.

[9] D. Fudenberg and J. Tirole, *Game theory*. MIT Press, Cambridge, 1991.

[10] T. L. Song and T. Y. Um, "Practical guidance for homing missiles with bearings-only measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 1, pp. 434–443, 1996.

[11] S. Battistini and T. Shima, "Differential games missile guidance with bearings-only measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 2906–2915, 2014.

[12] C.-H. Lee, T.-H. Kim, M.-J. Tahk, and K.-S. Kim, "Design of guidance law for passive homing missile using sliding mode control," in *Control Automation and Systems (ICCAS), 2010 International Conference on*, pp. 2380–2385, IEEE, 2010.

[13] D. K. Geller, "Linear covariance techniques for orbital rendezvous analysis and autonomous onboard mission planning," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 29, no. 6, pp. 1404–1414, 2006.

[14] D. K. Geller and D. Christensen, "Linear covariance analysis for powered lunar descent and landing," *AIAA Journal of Spacecraft and Rockets*, vol. 46, no. 6, pp. 1231–1248, 2009.

[15] S. M. LaValle, *Planning algorithms*. Cambridge University Press, Cambridge, 2006.

[16] E. A. Wan and R. van der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pp. 153–158, IEEE, 2000.

[17] R. Hermann and A. Krener, "Nonlinear controllability and observability," *IEEE Transactions on Automatic Control*, vol. 22, no. 5, pp. 728–740, 1977.

[18] G. L. Mariottini, F. Morbidi, D. Prattichizzo, N. Vander Valk, N. Michael, G. Pappas, and K. Daniilidis, "Vision-based localization for leader–follower formation control," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1431–1438, 2009.

[19] C. P. Robert and G. Casella, *Monte Carlo statistical methods*. Springer, New York, 2004.

[20] P. G. Moschopoulos, "The distribution of the sum of independent gamma random variables," *Annals of the Institute of Statistical Mathematics*, vol. 37, no. 1, pp. 541–544, 1985.

[21] N. A. Shneydor, *Missile guidance and pursuit: kinematics, dynamics and control*. Horwood Pub., Chichester, 1998.

[22] E. Tal, A. Gorodetsky, and S. Karaman, "Continuous tensor train-based dynamic programming for high-dimensional zero-sum differential games," in *American Control Conference (ACC), 2018*, pp. 6086–6093, IEEE, 2018.

[23] J. Lewin, *Differential Games: Theory and Methods for Solving Game Problems with Singular Surfaces*. Springer, London, 1994.

[24] E. Tal and S. Karaman, "Precision tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness," in *Decision and Control (CDC), 2018 IEEE 57th Conference on*, IEEE, 2018.