

Robust UAV Coordination for Target Tracking using Output-Feedback Model Predictive Control with Moving Horizon Estimation

Steven A. P. Quintero, David A. Copp, and João P. Hespanha

Abstract—We consider the control of two UAVs tracking an evasive moving ground vehicle. The UAVs are small fixed-wing aircraft equipped with gimballed cameras and must coordinate their control actions so that at least one UAV is always close to the target. The control actions of the UAVs are computed based on noisy measurements of the UAVs' current state and vision-based measurements of the target's position corrupted by state-dependent noise. We propose a novel approach for solving this problem in which the state estimate and optimal control are computed simultaneously online using model predictive control with moving horizon estimation. The efficacy of this approach is demonstrated in simulation results using realistic vision-based measurements of the target's position. These results show that while using noisy, partial information about the target state, the UAVs are able to coordinate their distances to the target in the ideal case of constant target velocity as well as perform robustly when the target becomes evasive.

I. INTRODUCTION

Small unmanned aerial vehicles (UAVs) are comparatively inexpensive mobile sensing platforms that are becoming ubiquitous due to their ability to autonomously perform tasks that would be either too demanding, dangerous, or mundane/repetitive for a human operator. Such tasks include agricultural monitoring, exploration and mapping, search and rescue, and surveillance and tracking, to name a few. One particularly challenging problem of interest is that of performing *vision-based target tracking* with a small fixed-wing UAV traveling at a constant airspeed and fixed altitude. In this problem, a camera-equipped UAV is responsible for measuring and tracking the position of a vehicle moving unpredictably in the ground plane.

In vision-based target tracking, image processing software is responsible for determining the centroid pixel coordinates of the ground target moving in the image frame. Using these pixel coordinates, along with the intrinsic and extrinsic camera parameters and terrain data, one can estimate the three-dimensional location of the target in inertial coordinates and compute the associated error covariance [1]. This vision-based measurement of the target's position is also referred to as the *geolocation* estimate. The error associated with the geolocation estimate is highly sensitive to the UAV's position relative to that of the target. As the UAV's planar distance from the target increases, the associated error covariance grows and becomes significantly elongated in the viewing

direction. When a UAV is directly above the target, the measurement error is smallest, as the corresponding error ellipse is circular. Thus, a UAV would ideally hover directly above the target, but the relative dynamics between a fixed-wing UAV and a moving ground target typically preclude this viewing position from being maintained over a period of time. To mitigate a single UAV's inability to maintain close proximity to the target, one can employ multiple UAVs to gather measurements, which are then fused to obtain an improved geolocation estimate. This is referred to as cooperative (or coordinated) target tracking.

Considerable work has been done in the general area of coordinated target tracking, with coordinated standoff tracking comprising the greatest body of work in this area. In standoff tracking, two UAVs orbit the target at a nominal standoff distance while maintaining orthogonal viewing angles. This practice minimizes the joint/fused geolocation error covariance at the fixed nominal standoff distance, as the individual measurement error ellipses are orthogonal [2]. Standoff tracking with perfect knowledge of the target state has been studied in [3] and [4] where the most prevalent control strategies involve the use of vector fields and nonlinear feedback. Approaches with only partial information of the target state are presented in [5], [6], and [7]. The authors of these works utilize observers, adaptive control, and extended Kalman filtering to estimate the full target state. Note that [7] utilizes nonlinear model predictive control to achieve the desired standoff configuration for a target that accelerates but is not necessarily evasive.

The preceding works have designed UAV coordination policies that attempt to improve the estimate of the target state without directly solving a dynamic optimization that minimizes some metric of the estimation error. However, a number of works have employed optimal control to achieve this objective. Miller et al. utilize the framework of partially observable Markov decision processes (POMDPs) in [8] to enable two UAVs to track a moving ground target and present a new approximate solution, as nontrivial POMDP problems are typically intractable to solve exactly [9]. Stachura et al. [10] employ online receding horizon control to enable two variable-air-speed UAVs to track a stochastic ground target using bearing-only sensors in the presence of packet losses when communicating with the base station where target state estimation takes place. In [11], Ding et al. study the problem of optimally controlling two Dubins vehicles and their pan-tilt-zoom cameras to maximize the geolocation information of a stochastic ground target and show that maintaining orthogonal viewing angles is essential in the case of terrestrial

S. A. P. Quintero, D. A. Copp, and J. P. Hespanha are with the Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, CA 93106, U.S.A., quintero@ece.ucsb.edu, dacopp@engr.ucsb.edu, hespanha@ece.ucsb.edu. This material is based upon work supported by the National Science Foundation under Grant No. CNS-1329650.

pursuit vehicles and less pronounced for airborne vehicles. While the preceding optimization-based methods consider short planning horizons, e.g., 2 – 3 seconds, Quintero et al. consider the optimal coordination of two Dubins vehicles to gather the best joint vision-based measurements of a constant-velocity target over long planning horizons of at least one minute [12], where no restrictions are placed on the vehicles other than kinematics. The results show that coordination of the distances to target is more effective for achieving the said goal than the traditional practice of solely coordinating viewing angles.

In this work, there are a number of real-world conditions to which we wish to be explicitly robust. First, we assume that the only information available for feedback is noisy measurements of each full UAV state and noisy vision-based measurements of the target’s position. Thus, the target’s velocity is unmeasured and must be estimated. Second, unlike any of the previous approaches on cooperative target tracking, we consider evasive target motion. Since the UAVs are fixed-wing aircraft, and are therefore subject to a minimum turning radius, they must carefully consider the impact of current control actions on future tracking performance in light of state estimation error and evasive target maneuvers. Moreover, robustness to both measurement noise and evasive target motion is crucial to the success of vision-based target tracking.

An output-feedback control approach that can be used to achieve the desired robustness was recently introduced by Copp and Hespanha in [13] and combines robust model predictive control (MPC) with moving horizon estimation (MHE). As described in [14] and [15], robust MPC involves an online dynamic optimization aimed at minimizing a cost function over a finite planning horizon in light of worst-case disturbances on a dynamical system. MHE also involves an online optimization problem but for the purpose of state estimation of nonlinear systems, and it has been shown to have advantages over state-of-the-art alternatives [16]. While the two optimizations have traditionally been done separately, in the framework of [13], the two are combined into a single min-max optimization. More specifically, a desired cost function is maximized with respect to disturbance and measurement noise variables and minimized with respect to control input variables. The min-max optimization provides state estimates over a fixed, finite window into the past and an optimal control input sequence into the future that is simultaneously robust to worst-case estimates of the state as well as worst-case disturbances to the plant. This combined robust MPC/MHE approach is demonstrated here as a viable, practical solution for the present particularly challenging nonlinear problem of autonomous vehicle coordination.

The remainder of the paper is organized as follows. Section II describes the dynamics and measurement model that compose the problem of vision-based target tracking. Section III discusses the cost function and the robust output-feedback MPC/MHE solution. Section IV presents and discusses simulation results for multiple scenarios. Finally, Section V provides conclusions and plans for future work.

II. PROBLEM FORMULATION

Consider two camera-equipped UAVs tasked with estimating the state of a target vehicle moving evasively in the ground plane. The UAVs fly at a fixed airspeed and constant altitude and are subject to a minimum turning radius. The target vehicle moves in the ground plane and is subject to a maximum acceleration and maximum speed that is less than the UAVs’ groundspeed, which is the same as its airspeed in the ideal case of no wind. Each UAV makes measurements of the target’s position using a gimbaled video camera, and we assume that the target is detected at all times and kept in the center of the camera’s field of view by onboard software. We first discuss the dynamical models for each type of vehicle and then proceed to describe their measurement models.

A. UAV Dynamics

The Dubins vehicle is a planar vehicle that moves forward at a fixed speed and has a bounded turning radius. It is commonly used to provide a simple model for UAVs flying at a fixed altitude. We assume that UAV j , where $j \in \{1, 2\}$, flies at a constant speed s_j and at a fixed altitude h_j , and it has a bounded turning rate u_j with maximum absolute upper bound $\bar{u} \in \mathbb{R}_{>0}$, which we take to be the same for both UAVs. Accordingly, $u \in \mathcal{U} := [-\bar{u}, \bar{u}] \times [-\bar{u}, \bar{u}]$. We denote by $\xi^{(j)} = [\xi_1^{(j)} \ \xi_2^{(j)} \ \xi_3^{(j)}]^\top \in \mathbb{R}^3$ the state of UAV j , which comprises its planar position $\mathbf{p}_j := (\xi_1^{(j)}, \xi_2^{(j)})$ and its heading $\psi_j := \xi_3^{(j)}$, all of which are measured in a local East-North-Up coordinate frame. The kinematics of UAV j are given by

$$\frac{d\xi^{(j)}(t)}{dt} = F(\xi^{(j)}, u_j) := \begin{pmatrix} s_j \cos \xi_3^{(j)} \\ s_j \sin \xi_3^{(j)} \\ u_j \end{pmatrix}. \quad (1)$$

While the majority of work on target tracking treats the problem in continuous time, this work addresses the problem in discrete time since measurements of the target’s position are available at discrete time instances $t = kT_s$ seconds, where $k \in \mathbb{Z}_{\geq 0}$ and $T_s > 0$ is the measurement sampling period. Accordingly, we assume a zero-order hold (ZOH) of T_s seconds on each UAV’s control input. The discrete-time equations of motion provided in [12] are given by

$$\xi^{(j)+} = f_a(\xi^{(j)}, u_j), \quad (2)$$

where the subscript “a” refers to the fact that the discrete-time dynamics are those of an air vehicle.

While the equations of motion for the Dubins vehicle in discrete time are readily derived in [12], they involve a $1/u_j$ term that becomes problematic for numerical optimizations and is hence avoided. Instead, we approximate the Dubins vehicle model using a second order Lie series:

$$f_a(\xi^{(j)}, u_j) \approx \xi^{(j)} + T_s F(\xi^{(j)}, u_j) + \dots + \frac{T_s^2}{2} \frac{\partial F}{\partial \xi^{(j)}} F(\xi^{(j)}, u_j), \quad (3)$$

where $\partial F / \partial \xi^{(j)}$ denotes the Jacobian of (1) that does not involve division by the control input. The Lie series is a

good approximation to the Dubins vehicle dynamics since the heading dynamics are exact and, with the simulation parameters considered in this work, the position error from (3) corresponding to $u_j = \bar{u}$ is less than 4% of the total distance traveled. In what follows, the approximation of (3) is used only in the numerical optimization of Section III while the simulation results of Section IV utilize (2) to propagate the discrete-time dynamics forward in time.

B. Target Dynamics and Overall State Space

We place no nonholonomic constraints on the ground vehicle and simply model the target as a double integrator moving in the ground plane. The state of the target is denoted by $\eta = [\eta_1 \ \eta_2 \ \eta_3 \ \eta_4]^\top \in \mathbb{R}^4$, where $\mathfrak{p}_g := (\eta_1, \eta_2)$ refers to the target's planar position in the same local East-North-Up coordinate frame as the UAVs. The corresponding target velocity is given by $\mathfrak{v} := (\eta_3, \eta_4)$, and the acceleration inputs of the target are denoted by $d \in \mathbb{R}^2$. We assume a T_s -second ZOH on the target's acceleration input synchronized with that of both UAVs, yielding the straightforward discrete-time linear dynamics

$$\eta^+ = f_g(\eta, d) = A\eta + Bd, \quad (4)$$

where

$$A = \begin{bmatrix} I_2 & T_s I_2 \\ 0 \cdot I_2 & I_2 \end{bmatrix} \text{ and } B = \begin{bmatrix} (T_s^2/2)I_2 \\ T_s I_2 \end{bmatrix}.$$

Here, I_n is an $n \times n$ identity matrix. To keep the problem realistic and well-posed, we take the target's acceleration input d to belong to

$$\mathcal{D} := \{d \in \mathbb{R}^2 \mid \|\mathfrak{v} + dT_s\|_2 \leq \bar{v}, \|d\|_\infty \leq \bar{d}\},$$

where \bar{v} is the maximum allowable target speed and \bar{d} is the maximum absolute acceleration along either the East or North directions. Typically, we take \bar{v} to be less than the smaller of the two UAV airspeeds so that the problem is well-posed.

Now that we have presented all vehicle models, we define the overall state as $x := (\xi^{(1)}, \xi^{(2)}, \eta) \in \mathbb{R}^{10}$. The overall dynamics are thus given by

$$x^+ = f(x, u, d) := \begin{pmatrix} f_a(\xi^{(1)}, u_1) \\ f_a(\xi^{(2)}, u_2) \\ f_g(\eta, d) \end{pmatrix}.$$

C. Measurement Error Models

We turn our attention to the overall measurement model in vision-based target tracking. The measurement vector associated with the state of UAV j is denoted by $y_a^{(j)} \in \mathbb{R}^3$ and is given by

$$y_a^{(j)} = \xi^{(j)} + w_a^{(j)}, \quad w_a^{(j)} = [w_{a,1}^{(j)} \ w_{a,2}^{(j)} \ w_{a,3}^{(j)}]^\top, \quad (5)$$

where $(w_{a,1}^{(j)}, w_{a,2}^{(j)}) \sim \mathcal{N}(0, \sigma_p^2 I_2)$, $w_{a,3}^{(j)} \sim \mathcal{N}(0, \sigma_\psi^2)$, σ_p^2 is the variance of the uncorrelated noise on the UAV's North and East position coordinates, and σ_ψ^2 is the variance on the UAV's heading angle.

Each UAV's camera makes image-plane measurements of the target. The dominant source of geolocation error arises

from the error in the sensor attitude matrix $T_S^T(\theta_j)$ that relates the coordinates of the line-of-sight vector u_j^S from the UAV to the target in the North-East-Down sensor frame (centered at UAV j 's position) to the coordinates of the same vector in the local East-North-Up topographic coordinate frame. This transformation is a nonlinear function of the 3-2-1 Euler-angle sequence of yaw, pitch, and roll denoted by $\theta_j \in \mathbb{R}^3$. Image tracking software controls the camera's gimbal platform to keep the target in the center of the camera's field of view and reports the Euler angles of the camera sensor as well as the line-of-sight vector u_j^S . Here, a superscript "S" denotes a quantity in the sensor coordinate frame while the absence thereof indicates a quantity in the topographic coordinate frame.

The 3-dimensional target position measured by UAV j with 3D position $\mathfrak{s}_j = [\mathfrak{p}_j^\top, h_j]^\top$ is denoted by \mathfrak{o}_j . Its estimate is given by

$$\hat{\mathfrak{o}}_j = \hat{\mathfrak{s}}_j + \hat{r}_j T_S^T(\hat{\theta}_j) u_j^S = \hat{\mathfrak{s}}_j + \hat{r}_j \hat{u}_j,$$

where

$$\begin{aligned} \hat{\mathfrak{s}}_j &= \mathfrak{s}_j + \tilde{\mathfrak{s}}, \quad \tilde{\mathfrak{s}} \sim \mathcal{N}(0, \text{diag}(\sigma_p^2 I_2, \sigma_a^2)), \\ \hat{\theta}_j &= \theta_j + \tilde{\theta}, \quad \tilde{\theta} \sim \mathcal{N}(0, \sigma_\psi^2 I_3), \end{aligned} \quad (6)$$

and σ_a^2 denotes the variance of the measurement noise on the UAVs' altitude h_j . Also, the 3D distance from UAV j to the target is denoted by $r_j = \|\mathfrak{o}_j - \mathfrak{s}_j\|_2$, and its estimate \hat{r}_j is provided by the flat-Earth approximation $\hat{r}_j = (h_0 - \hat{\mathfrak{s}}_{j,3})/\hat{u}_{j,3}$, where h_0 is the height of the ground plane in the topographic coordinate frame and is taken to be zero in this work without loss of generality. Since all camera angles are measured with respect to the UAV attitude, we take the noise on the estimate of the camera's attitude to be the same as that on the estimate of the UAV's attitude. Thus, the noise on the estimate of UAV j 's heading angle is the same as that on the estimate of its camera's yaw angle, i.e., the first element of θ_j .

From the preceding measurement equation, one can show that the covariance $P_{\mathfrak{o},j} \in \mathbb{R}^{3 \times 3}$ associated with the error $\tilde{\mathfrak{o}}_j := \hat{\mathfrak{o}}_j - \mathfrak{o}_j$ in the three dimensional position of the target is proportional to the product of r_j^2 and the covariance of the Euler-angle sequence estimate $\hat{\theta}_j$ given in (6). The exact analytic expression for $P_{\mathfrak{o},j}$ is derived in [12] and is omitted here for brevity. Since we are tracking in the ground plane, only the upper left 2×2 submatrix of $P_{\mathfrak{o},j}$ is relevant and is denoted by P_j .

Since the UAVs collect independent measurements of the target, the fused measurement y_g of the target's true position \mathfrak{p}_g can be computed using the best linear unbiased estimate, which is as follows:

$$\begin{aligned} y_g &= \mathcal{P}(P_1^{-1} \hat{\mathfrak{p}}_g^{(1)} + P_2^{-1} \hat{\mathfrak{p}}_g^{(2)}) \\ &= [I_2 \ \mathbf{0}_{2 \times 2}] \eta + w_g = \mathfrak{p}_g + w_g, \end{aligned} \quad (7)$$

where $\mathcal{P} = (P_1^{-1} + P_2^{-1})^{-1}$, $w_g \sim \mathcal{N}(0, \mathcal{P})$, $\hat{\mathfrak{p}}_g^{(j)} = [I_2 \ \mathbf{0}_{2 \times 1}] \hat{\mathfrak{o}}_j$, and $\mathbf{0}_{m \times n}$ denotes the $m \times n$ zeros matrix. The confidence ellipse corresponding to the fused geolocation

error covariance (GEC) \mathcal{P} has the property that it is small when at least one UAV is close to the target and only slightly less when both aircraft are directly above the target. Therefore, it is advantageous for at least one UAV to be near the target at any given time.

Finally, the measurement model corresponding to the overall state x is given by combining (5) and (7) as follows:

$$y := (y_a^{(1)}, y_a^{(2)}, y_g) = Cx + w,$$

where $w := (w_a^{(1)}, w_a^{(2)}, w_g)$ and $C := [I_8 \ \mathbf{0}_{8 \times 2}]$. Since the target velocity is not measured directly, the control law used in this framework will be based on output feedback.

III. ROBUST OUTPUT-FEEDBACK MPC/MHE

When considering only one UAV, the target tracking problem can be regarded as a two-player zero-sum game in which the UAV tries to minimize its 3D distance to the target, r , and the target tries to maximize r . In the two-UAV case, the UAVs ideally coordinate their movements in order to ensure that at least one UAV is close to the target to keep the fused GEC comparatively low. Additionally, the UAVs should keep their individual distances to the target sufficiently small to maintain adequate resolution of the target in the camera's image plane for effective visual detection. This motivates us to choose the following criterion

$$g(x) := \beta_1 \frac{r_1^2 r_2^2}{r_1^2 + r_2^2} + \beta_2 (r_1^2 + r_2^2), \quad (8)$$

where β_1 and β_2 are positive weighting constants, that the UAVs (Player 1) would like to minimize and the target (Player 2) would like to maximize. The first term in (8) is motivated by noting that the size of the confidence ellipse associated with P_j is proportional to r_j^2 and that the fused GEC has the form $\mathcal{P} = (P_1^{-1} + P_2^{-1})^{-1}$. Moreover, the previous matrix expression is simplified to one that is scalar and more compatible with numerical optimization by replacing the individual covariances with the respective 3D distances. This term enforces distance coordination so that one UAV is always close to the target to improve measurement quality, just as in [12]. The second term in (8) penalizes the individual UAV distances to the target to ensure that the size of the target in each UAV's image plane is sufficiently large for reliable detection by image processing software. While other optimality criterion may be considered, we aim to utilize a simpler expression than those found in [11] and [12] that achieves similar behavior and lends itself to efficient numerical computation. We shall see that the distance coordination of [12] is indeed induced by choosing the criterion (8).

For our control approach, we use the output-feedback MPC with MHE approach described in [13]. This requires us to solve a finite-horizon online optimization problem at each time k . Solving this online optimization problem uses output measurements from the last L steps into the past in order to give us an estimate of the current state at time k (the MHE problem), and from that, give us policies for both

Players 1 and 2 to use for the next K steps into the future (the MPC problem).

Specifically, the control objective is to select the control signal $u_k \in \mathcal{U}$, $\forall k \in \mathbb{Z}_{\geq 0}$ so as to minimize a criterion of the form

$$\sum_{\ell=k}^{k+K} g(x_\ell) - \sum_{\ell=k-L}^k \lambda_n \|w_\ell\|_2^2, \quad (9)$$

where $g(x)$ is given by (8) and λ_n is a positive scalar, for worst-case values of $d_k \in \mathcal{D}$ and measurement noise $w_k \in \mathbb{R}^{n_w}$, $\forall k \in \mathbb{Z}_{\geq 0}$.

Given a discrete-time signal $z : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}^n$ and two times $k_0, k \in \mathbb{Z}_{\geq 0}$ with $k_0 \leq k$, we denote by $z_{k_0:k}$ the sequence $\{z_{k_0}, z_{k_0+1}, \dots, z_k\}$. This notation allows us to re-write (9) and define our cost function as

$$J_k(x_{k-L}, u_{k-L:k+K-1}, d_{k-L:k+K-1}, y_{k-L:k}) := \sum_{\ell=k}^{k+K} g(x_\ell) - \sum_{\ell=k-L}^k \lambda_n \|y_\ell - Cx_\ell\|_2^2 \quad (10)$$

which emphasizes the dependence of (9) on the unknown initial state x_{k-L} , the unknown input sequence for the target $d_{k-L:k+K-1}$, the measured output sequence $y_{k-L:k}$, and the control input sequence $u_{k-L:k+K-1}$. This control input sequence $u_{k-L:k+K-1}$ comprises two distinct sequences: the (known) past inputs $u_{k-L:k-1}$ that have already been applied and the future inputs $u_{k:k+K-1}$ that still need to be selected.

At a given time $k \in \mathbb{Z}_{\geq 0}$, we do not know the values x_{k-L} and $d_{k-L:k+K-1}$ on which the criterion (10) depends, so we optimize this criterion under worst-case assumptions on these variables, leading to the following finite-dimensional min-max optimization

$$\min_{\hat{u}_{k:k+K-1|k} \in \mathcal{U}} \max_{\hat{x}_{k-L|k} \in \mathcal{X}, \hat{d}_{k-L:k+K-1|k} \in \mathcal{D}} J_k(\hat{x}_{k-L|k}, u_{k-L:k-1}, \hat{u}_{k:k+K-1|k}, \hat{d}_{k-L:k+K-1|k}, y_{k-L:k}), \quad (11)$$

where the arguments $u_{k-L:k-1}, \hat{u}_{k:k+K-1|k}$ correspond to the sequence $u_{k-L:k+K-1}$ in the definition of $J_k(\cdot)$ in (10). The subscript $|\cdot|_k$ in the (dummy) optimization variables in (11) emphasizes that this optimization is repeated at each time step $k \in \mathbb{Z}_{\geq 0}$. At different time steps, these optimizations typically lead to different solutions, which generally do not coincide with the real control input, target input, and noise. We can view the optimization variables $\hat{x}_{k-L|k}$ and $\hat{d}_{k-L:k+K-1|k}$ as (worst-case) estimates of the initial state and target input, respectively, based on the past inputs $u_{k-L:k-1}$ and outputs $y_{k-L:k}$ available at time k .

Inspired by model predictive control, at each time k , we use as the control input the first element of the sequence

$$\hat{u}_{k:k+K-1|k}^* = \{\hat{u}_{k|k}^*, \hat{u}_{k+1|k}^*, \hat{u}_{k+2|k}^*, \dots, \hat{u}_{k+K-1|k}^*\} \in \mathcal{U}$$

that minimizes (11), leading to the following control law:

$$u_k = \hat{u}_{k|k}^*, \quad \forall k \geq 0.$$

TABLE I
SIMULATION PARAMETERS

Parameter	Description	Value	Units
\bar{u}	Max UAV turn rate	0.46	rad./s
(s_1, s_2)	UAV speeds	(15, 15)	m/s
\bar{d}	Max target accel.	$3/\sqrt{2}$	m/s^2
\bar{v}	Max target speed	10	m/s
σ_p^2	N/E position variance	2.5^2	m^2
σ_h^2	Altitude variance	4^2	m^2
σ_ψ^2	Euler angle variance	3^2	deg.^2
(h_1, h_2)	UAV altitudes	(40, 45)	m
β_1	Coord. coefficient	$4 \cdot 10^{-2}$	-
β_2	Dist. coefficient	$4 \cdot 10^{-3}$	-
λ_n	Noise coefficient	10	-
T_s	Sampling period	1	s
L	Backward horizon	7	-
K	Forward horizon	7	-

IV. SIMULATION RESULTS

We now demonstrate the effectiveness of the proposed robust MPC/MHE control approach to the problem of vision-based target tracking with two UAVs. In particular, we simulate two scenarios with realistic levels of noise in the models presented in Section II. Firstly, we quantify the performance of the proposed approach in an ideal scenario wherein the target is traveling at a constant velocity and secondly in a scenario in which the target follows its worst-case strategy determined by (11).

The parameters pertaining to both simulation scenarios are provided in Table I. Given the maximum UAV turn rate, the total time it takes a UAV to make a full loop is $2\pi/\bar{u} \approx 13.66$ seconds, and hence the future planning horizon of $KT_s = 7$ seconds allows the UAVs to consider the impact of beginning to loop around the target. Additionally, in the cost function (8), the coefficients β_1 and β_2 were chosen to place a greater emphasis on distance coordination than on keeping individual distances small.

To solve the min-max optimization problem (11), we use a primal-dual-like interior point method that is designed specifically for this MPC/MHE approach and is described in [17]. Typical execution times for solving the optimization problem at each time step using the C programming language on a laptop with a 2.3 GHz Intel[®] Core[™] i7 processor are near 5 milliseconds with a maximum execution time not exceeding 20 milliseconds. Since these execution times are much less than the $T_s = 1$ second sampling period, the approach presented here is suitable for online real-world implementation.

Because the simulations incorporate stochastic measurement noise, we quantify performance based on $M = 1,000$ Monte Carlo simulations to determine the steady-state tracking performance of each scenario. More specifically, 3 minutes of steady-state behavior are considered, where the effects of initial conditions have been removed by discarding 30 seconds of initial data. We initialize the problem with the UAVs circling a stationary target until the past measurement buffer is filled. Figure 1 depicts an instance of the output of




Fig. 1. The output of the robust MPC/MHE optimization. Red and blue markers depict quantities related to the UAVs while black markers indicate quantities related to the target. The “+” markers illustrate the vehicles’ noisy position measurements while the “*” markers depict ground truth. The “o” markers indicate state estimates from the min-max optimization while “x” markers illustrate future positions corresponding to the optimal control sequence determined by the same optimization. A “·” marker indicates the estimate of each vehicle’s current position. The target trajectory begins near the center of the plot and ends heading Northwest while the leftmost UAV markers indicate their starting positions.

the optimization (11) at time k for the parameters given in Table I. This plot illustrates the past positions, measurements, and estimates of the UAVs and target as well as the planned future trajectory for each agent generated at the current time k . In this instance, we see that the target just executed a sharp turn in order to take advantage of the UAVs’ minimum turning radius of $r \approx 32.6$ [m].

We first consider a constant-velocity target, i.e., $d_k \equiv 0$ in equation (4), and the target travels at just over one third of the UAVs’ fixed speed. The UAVs are still performing the min-max optimization of (11) and hence are planning for an evasive target. The results for one of the Monte Carlo simulations are provided in Figures 2 and 3. In Figure 2, one can see how the UAVs make loops and “S” turns so that their average speed matches that of the target. Figure 3 indicates that this is done in a coordinated fashion so that at least one UAV is never very far from the target, as indicated by the dashed cyan curve depicting $\min\{r_1, r_2\}$. We shall soon verify that coordination is consistently achieved across all runs for a constant-velocity target using the robust MPC/MHE approach with the given stage cost.

We now consider an evasive target, i.e., it is using the optimal worst-case d_k^* computed from the min-max optimization (11). Results for one of the Monte Carlo simulations are provided in Figures 4 and 5. By observing the vehicle trajectories in Figure 4, one can see that the optimal trajectory for the target is quite erratic. Indeed, the target takes advantage of the UAVs’ kinematic constraints by making sharp turns and forcing the UAVs to make loops at their maximum allowable turn rate. This is corroborated by the




Fig. 2. Actual trajectories of two UAVs tracking a constant-velocity target over a 3-minute window. The starting positions of all vehicles are denoted by an “o” while the ending positions are indicated by an “x”. In the legend, \mathcal{T} corresponds to the target while \mathcal{A}_1 and \mathcal{A}_2 refer to the UAVs.




Fig. 3. 3D distances r_j and stage cost $g(x)$ for two UAVs tracking a constant velocity target.

fact that over the 1,000 runs, the modes of the sample distributions (not shown) of the absolute control input values for all vehicles are concentrated at their maximum allowable values. From Figure 5, one can see that distance coordination has diminished, and the stage cost peaks over 300 [m²] twice, whereas it never exceeded 200 [m²] in the case of a constant-velocity target.

A quantitative summary of the $M = 1,000$ Monte Carlo runs is provided in Table II, where sample statistics are computed over both time and samples. From the data, one should notice the detrimental effect that adversarial target motion has on the coordination effort of the UAVs. This is indicated by the sample Pearson correlation coefficient ϱ for the UAV distance pairs r_1 and r_2 , where ϱ is in general a measure of the linear correlation between two random variables and belongs to the interval $[-1, 1]$. A more negative value for ϱ indicates stronger anti-correlation, which in the present setting implies that when one UAV is relatively far from the target, the other is likely to be rather close. Thus, one can see that the constant velocity scenario has substantial distance coordination across all simulation runs; however, the evasive target is able to disrupt the UAV coordination and induce positive correlation. Nonetheless, the average cost sees only a moderate increase of nearly 18% when the target becomes evasive while the maximum value of the individual UAV distances from the target only experiences an increase of about 10%.

Concerning estimation of the target’s position, one should




Fig. 4. Actual trajectories of two UAVs tracking an evasive target over a 3-minute window. The same notation as Figure 2 is employed here.




Fig. 5. 3D distances r_j and stage cost $g(x)$ for two UAVs tracking an evasive target.

TABLE II
STATISTICS FOR 1,000 3-MINUTE MONTE CARLO SIMULATIONS

Statistic	Const. Vel.	Evasive	Units
avg $g(x)$	118.4	139.3	m ²
var $g(x)$	1,163.8	3,439.4	m ⁴
max $\{r_1, r_2\}$	154.2	169.7	m
avg $\ \hat{p}_g - \hat{p}_g\ $	4.92	4.66	m
avg $\ w_g\ $	4.72	5.24	m
ϱ	-0.307	0.206	N/A

first observe that average measurement error (w_g) on the target’s position increases when the target is evasive, as the fused error covariance increases with the UAVs’ inability to keep at least one aircraft close at all times. However, the estimation error becomes smaller, which shows that the target’s deviation from its optimal worst-case evasion leads to a smaller cost for (10). This means that in the evasive case, the first term of (10) that quantifies tracking performance has increased and the magnitude of the second (negative) term regarding state estimation has decreased, which can be observed in the first and fourth rows of Table II.

Overall, the target’s evasive maneuvers hinder UAV coordination efforts and thereby increase measurement noise. Nonetheless, the UAVs are still able to robustly track the target in the sense that their maximum 3D distance from the

target is not only bounded, but also only slightly larger than in the case of constant target velocity.

V. CONCLUSION AND FUTURE WORK

We considered two UAVs performing vision-based target tracking of a moving ground vehicle. We showed that a novel approach based on min-max MPC combined with MHE is viable for solving this high-dimensional, very nonlinear (non convex) problem. Using this approach with the commonly assumed case of having a constant-velocity target, the UAVs coordinate their distances from the target and achieve very small tracking cost, even though only noisy, partial information about the target state is available. In practice, not only is the full state of the target and UAVs not precisely known, but sharp changes in target velocity occur when the target vehicle makes turns and moves erratically or evasively. Though past work on UAV target tracking does not often consider this case, we have shown that our approach is robust not only to evasive target motion, but also to noisy measurements in the output-feedback setting of vision-based target tracking.

Future work involves utilizing a slightly more realistic UAV model that addresses roll dynamics. With more realistic aircraft dynamics in place, we plan to verify the proposed approach in a high fidelity flight simulator for the UAVs with actual target trajectories captured from real-world tracking experiments.

REFERENCES

- [1] M. Mallick, "Geolocation using video sensor measurements," in *IEEE International Conf. on Information Fusion*, Quebec, Canada, July 2007.
- [2] E. W. Frew, "Sensitivity of cooperative geolocation to orbit coordination," in *AIAA Guid., Nav., and Control Conf.*, Hilton Head, SC, aug 2007, pp. 3869–3892.
- [3] L. Ma and N. Hovakimyan, "Cooperative target tracking in balanced circular formation: Multiple UAVs tracking a ground vehicle," in *American Control Conf.* IEEE, 2013, pp. 5386–5391.
- [4] D. B. Kingston, "Decentralized control of multiple UAVs for perimeter and target surveillance," Ph.D. dissertation, Brigham Young University, Dec. 2007.
- [5] C. Peterson and D. A. Paley, "Multivehicle coordination in an estimated time-varying flowfield," *Journal of guidance, control, and dynamics*, vol. 34, no. 1, pp. 177–191, 2011.
- [6] T. H. Summers, M. R. Akella, and M. J. Mears, "Coordinated standoff tracking of moving targets: Control laws and information architectures," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 32, no. 1, pp. 56–69, 2009.
- [7] S. Kim, H. Oh, and A. Tsourdos, "Nonlinear model predictive coordinated standoff tracking of a moving ground vehicle," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 2, pp. 557–566, 2013.
- [8] S. A. Miller, Z. A. Harris, and E. K. P. Chong, "A POMDP framework for coordinated guidance of autonomous UAVs for multitarget tracking," *EURASIP J. Adv. Signal Process.*, pp. 1–17, 2009.
- [9] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, 2005.
- [10] M. Stachura, A. Carfang, and E. W. Frew, "Cooperative target tracking with a communication limited active sensor network," in *International Workshop on Robotic Wireless Sensor Networks*, Marina Del Ray, CA, June 2009.
- [11] C. Ding, A. A. Morye, J. A. Farrell, and A. K. Roy-Chowdhury, "Coordinated sensing and tracking for mobile camera platforms," in *American Control Conf.* IEEE, 2012, pp. 5114–5119.
- [12] S. A. P. Quintero, F. Papi, D. J. Klein, L. Chisci, and J. P. Hespanha, "Optimal UAV coordination for target tracking using dynamic programming," in *IEEE Conf. on Decision and Control*, Atlanta, GA, December 2010.
- [13] D. A. Copp and J. P. Hespanha, "Nonlinear output-feedback model predictive control with moving horizon estimation," in *IEEE Conf. on Decision and Control*, Los Angeles, California, 2014, note: In press.
- [14] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*. Springer, 1999, pp. 207–226.
- [15] H. S. Chen and F. Allgöwer, "A game theoretic approach to nonlinear robust receding horizon control of constrained systems," in *American Control Conference*, 1997.
- [16] J. B. Rawlings and B. R. Bakshi, "Particle filtering and moving horizon estimation," *Computers & Chemical Engineering*, vol. 30, no. 10, pp. 1529–1541, 2006.
- [17] D. A. Copp and J. P. Hespanha, "Nonlinear output-feedback model predictive control with moving horizon estimation," Univ. California, Santa Barbara, Tech. Rep., 2014, <http://www.ece.ucsb.edu/~hespanha/techrep.html>.