Robust Coordination of Small UAVs for Vision-based Target Tracking using Output-Feedback MPC with MHE

Steven A. P. Quintero[1], David A. Copp, João P. Hespanha [2]

June 17, 2016

[1]S. A. P. Quintero performed work while at the Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, CA 93106, U.S.A., steven@stevenquintero.net

[2]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., 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.

Small unmanned aerial vehicles (UAVs) are relatively inexpensive mobile sensing platforms that are finding increasingly widespread use in both commercial and military sectors due to their ability to autonomously perform tasks that would be either too demanding, dangerous, or mundane for a human operator. Such tasks include agricultural monitoring, exploration and mapping, search and rescue, and surveillance and tracking, to name a few. As the proliferation of small UAVs has allowed them to be manufactured at much lower costs than their larger counterparts, these robotic agents are now being deployed as multi-agent teams to perform missions in a cooperative fashion. The increased numbers enable the fleet to overcome the limitations of lower quality sensors and reduced computational resources that typically accompany small, inexpensive mobile robots, as they are able to distribute sensing and computation across the fleet. The plural nature of the fleet also enables redundancy and synoptic coverage of an object or area of interest.

One commonly underlying theme for UAV applications is that these robotic agents are faced with gathering measurements and making decisions without the aid of a human operator. Furthermore, the decision making must not only be made in realtime as new information becomes available, but it must also be done robustly to avoid mission failure. Additionally, as a measure of quality can typically be associated with the information gathered, the motion of the individual vehicles can be optimized so as to gather the best *joint* measurements of a given quantity, object, or area of interest. For example, a team of UAVs might be tasked with obtaining a diverse set of pressure measurements in a hurricane, and in the military reconnaissance application of target tracking, a UAV team might be tasked with tracking a moving ground vehicle whose motion is unpredictable. In both cases, the UAVs must coordinate their present actions in light of future measurement gathering performance, and their joint actions must be robust to real world conditions, such as environmental disturbances and unmodeled dynamics. If fixed-wing aircraft are employed, as they often are for their speed and endurance, the motion constraints of the vehicles, namely airspeed limits and a minimum turning radius, pose a significant challenge for any feedback control or motion planning solution.

One final challenge for a team of UAVs is that typically only noisy, partial information is available for decision making. Hence a state estimation algorithm is typically employed to estimate unmeasured system states and improve the accuracy of measured states. Thus, the corresponding decision-making algorithm, which is usually independent of the state estimator, must be inherently robust to noise in the estimated system states to avert mission failure, and the overall control solution is one of *output feedback*. Additionally, parametric uncertainty may also be present in the system dynamics, and thus estimation of the unknown parameters may be done online in realtime to overcome the uncertainty. While such uncertainty may take the form of time constants in the vehicle dynamics, it may also be present in the form of an environmental disturbance such as wind, which can be detrimental to motion planning algorithms if neglected. This is especially true for small or micro air vehicles where wind constitutes a significant portion of the vehicle's ground speed.

Thus, in this chapter, we firstly aim to highlight solutions for the robust, coordinated control of multi-agent systems that represent the state of the art, where a special emphasis is placed on output-feedback approaches. Secondly, we present a novel, output-feedback approach that both enables the robust, optimal control of smaller multi-agent systems with nonlinear dynamics and is suitable for real-world implementation, as corroborated by realtime, high fidelity simulations. We do this in the context of the particularly challenging problem of *vision-based target tracking* with small, fixed-wing UAVs. In this application, one employs multiple, camera-equipped UAVs traveling at a fixed airspeed and constant altitude to jointly measure and track the position of a moving ground vehicle whose motion is unpredictable. The main benefit of focusing on this problem is that it has nearly all of the aforementioned real-world conditions to which agents in a coordinated, multi-agent system must be robust. Namely, the motion planning for the team must be robust to target vehicle motion that may be either random, evasive, or otherwise, and hence the evolution of some of the overall system states is unpredictable and may be adverse to the given objective. Additionally, planning must be done in spite of noisy, partial information for all vehicle states and unmeasured, non-zero wind conditions, and it must be done with significant foresight, as the motion constraints of the UAVs may cause greedy, myopic approaches to either have poor long-term tracking performance or lose the possibly evasive target altogether. And lastly, one must utilize a motion model for the aircraft dynamics with a sufficient degree of realism, else ignored dynamics will have a detrimental effect on tracking performance.

0.1 Vision-based Target Tracking

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 measurement error's covariance grows, and the associated confidence ellipse becomes significantly elongated in the viewing direction. When a UAV is directly above the target, the measurement error is smallest, as the corresponding confidence 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 confidence ellipses corresponding to the individual measurement errors 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]. These works utilize observers, adaptive control, and extended Kalman filtering, respectively, 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. In [8], Sinha et al. demonstrate a decentralized tracking approach wherein each member of a UAV team maximizes its information regarding the target state while accounting for its teammates' track states and survival and detection probabilities as well. Miller et al. utilize the framework of partially observable Markov decision processes (POMDPs) in [9] 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 [10]. Stachura et al. [11] employ online receding horizon control to enable two variable-airspeed 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 [12], 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 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 UAVs to gather the best joint vision-based measurements of a moving target over considerably longer planning horizons of at least 30 seconds, where no restrictions are placed on the vehicles other than kinematics. Firstly, in [13], Quintero et al. consider two Dubins vehicles tracking a constant-velocity target in multiple scenarios differentiated by target speeds. The results in all scenarios 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. Secondly, in [14], Quintero et al. advance the work in [13] by considering a more realistic problem formulation that instead makes use of stochastic fourth-order motion models for both the UAV team and ground target. Using a novel regression-based dynamic programming approach, the

authors solve the formidable 9-dimensional stochastic optimal control problem offline and show that distance coordination can still be achieved under more realistic settings that include UAV roll dynamics, environmental disturbances, and dynamical uncertainty.

There is other notable work in the area of target tracking that is outside both of the aforementioned categories of standoff tracking and optimization-based control. For example, in [15], Triplett et al. focus on the challenges of communication, control, and state estimation as multiple UAVs track a moving ground vehicle. In this work, a decentralized Extended Kalman Filter is used to estimate the target state while a turn-rate based steering controller is utilized to have the pursuit group's centroid track the target, achieve a desired inter-vehicle spacing, and avoid obstacles. The results show that increased communication frequency and sensor reliability improve the estimation performance of the system. Additionally, a few works, namely [16] and [17], have proposed using sinusoidal turn-rate control inputs that approximate the optimal behavior of [13] at higher speeds.

While all of the preceding works do not constitute an exhaustive literature review on target tracking approaches, they are representative of the primary techniques taken to address this challenging problem. Furthermore, in all of the aforementioned works, at least one or more assumptions are made that create severe hindrances for real-world implementation. More specifically, the works mentioned thus far assume at least one of the following, either explicitly or implicitly:

1. Coordinated circular trajectories are optimal, namely those trajectories resulting from coordinated standoff tracking.

2. The UAVs are controlled by commanding turn rate, so roll dynamics are ignored.

3. The UAV airspeed can be changed quickly and reliably over a significant range.

4. Target motion is predictable.

5. Greedy planning horizons are adequate for optimal tracking.

6. All systems states are known exactly.

Additionally, the approaches not relying on optimization do not always place explicit constraints on the control effort. We also highlight the fact that the target's motion may not only be unpredictable but also evasive, which was not considered in any of the previous works and requires a robust planning solution with adequate foresight to mitigate the effects of the inherent dynamical limitations of the fixed-wing UAVs. The present work avoids the said practices and assumptions to promote a more practical solution that yields robust coordination under the following realistic conditions: unknown constant wind,

non-negligible roll dynamics with roll-angle setpoint limits, evasive target motion, and the availability of only noisy, partial information of the overall system's states.

An output-feedback control approach that can be used to achieve the desired robustness was recently introduced by Copp and Hespanha in [18] and combines robust model predictive control (MPC) with moving horizon estimation (MHE). As described in [19] and [20], 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 [21]. While the two optimizations have traditionally been computed separately, in the framework of [18], 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. Through realtime high-fidelity flight simulations, this combined robust MPC/MHE approach is shown to be a viable, practical solution for the present particularly challenging nonlinear problem of autonomous vehicle coordination.

The work in this chapter is based on work that appears in the conference paper [22] in a simpler form. In this chapter, significant extensions to [22] have been made. Specifically, whereas the traditional third-order Dubins vehicle model was utilized in [22], here we consider a more realistic fourth-order model that incorporates roll dynamics. Furthermore, the problem formulation in the earlier conference paper considered the ideal condition of no wind. As wind can have a considerable effect on small unmanned aircraft, the present tracking solution incorporates adaptive estimates of the wind into the online MPC/MHE optimization, thereby explicitly addressing the effects of wind and adding robustness to the motion planning of the UAVs. Lastly, the Dubins vehicle model was used to validate the approach taken in [22]. In constrast, the results section in this chapter illustrate the effectiveness of the approach in realistic environments by using high fidelity simulations that make use of an aircraft model with six degrees of freedom and target logs from actual tracking experiments at Camp Roberts, California, U. S. A.

The remainder of the chapter is organized as follows. Section 0.2 describes the dynamics and measurement model that compose the problem of vision-based target tracking. Section 0.3 discusses the cost function and the robust output-feedback MPC/MHE solution. Section 0.4 presents and discusses simulation results for multiple scenarios. Finally, Section 0.5 provides conclusions and plans for future work.

0.2 Problem Formulation

Consider two camera-equipped UAVs tasked with estimating the state of a target vehicle moving evasively in the ground plane. The UAVs are fixed-wing aircraft that fly at a constant altitude and have an autopilot that regulates roll angle, airspeed, and altitude to the desired setpoints via internal feedback loops. Furthermore, these underactuated aircraft are assumed to fly at a constant airspeed since the range of permissible airspeeds for such small aircraft may be very limited, as noted in [23] and §5.1 of [4]. The roll angle setpoint is hence the sole control input that affects the horizontal plant dynamics. The target vehicle moves in the ground plane and is subject to a maximum acceleration and maximum speed that is considerably less than the slowest UAV's airspeed, which means that the groundspeed of each UAV will always exceed that of the target in light wind conditions. 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.

0.2.1 UAV Dynamics

The basis for our UAV motion model is the kinematic guidance model given by Equation 9.20 of [24], where the altitude, altitude rate, and airspeed states have been omitted since the said quantities are treated as being constant in this work. Note that similar models have also been used in [14], [25], and [26]. To present the model, we assume that UAV j, where $j \in \{1, 2\}$, flies at a constant airspeed s_j and at a fixed altitude h_j. Its roll-angle setpoint is denoted by u_j and has a maximum absolute limit of $\bar{u} \in \mathbb{R}_{>0}$, which we take to be the same for both UAVs without loss of generality. Accordingly, $u \in \mathcal{U} := [-\bar{u}, \bar{u}] \times [-\bar{u}, \bar{u}]$. The state of UAV j comprises it planar position $\mathfrak{p}_j \in \mathbb{R}^2$, its heading angle $\psi_j \in \mathbb{R}$, and its roll angle $\phi_j \in \mathbb{R}$. Furthermore, we denote the state of UAV j by $\xi^{(j)} = [\xi_1^{(j)} \ \xi_2^{(j)} \ \xi_3^{(j)} \ \xi_4^{(j)}]^\top \in \mathbb{R}^4$, where $(\xi_1^{(j)}, \xi_2^{(j)}) := \mathfrak{p}_j$, $\xi_3^{(j)} := \psi_j$, and $\xi_4^{(j)} := \phi_j$. The position is given in an East-North-Up coordinate frame while the roll and heading angles are used in a 3-2-1 (yaw-pitch-roll) Euler angle sequence that describes the transformation from the vehicle-fixed inertial frame to the aircraft body frame, where the pitch angle is assumed to be zero in this work. We also assume that the UAV flies in a non-zero constant wind $w \in \mathbb{R}^2$, which is an unknown quantity that must be estimated. Furthermore, we assume that the wind speed is no more than half the minimum UAV airspeed, i.e., $w \in \mathcal{W}$, where

$$\mathcal{W} := \{v \in \mathbb{R}^2 | \|v\| \leq 0.5 \min\{s_1, s_2\}\}. \tag{1}$$

Thus, the kinematic model for UAV j used in this work is given by

$$\frac{d\xi^{(j)}(t)}{dt} = F(\xi^{(j)}, u_j) := \begin{pmatrix} s_j \cos\xi_3^{(j)} + w_1 \\ s_j \sin\xi_3^{(j)} + w_2 \\ (\alpha_g/s_j)\phi_j \\ -\alpha_{\phi,j}(\phi_j - u_j), \end{pmatrix}. \tag{2}$$

where α_g is the vertical acceleration due to gravity and $1/\alpha_{\phi,j}$ is the time constant for UAV j's roll dynamics that are governed by the implementation of the autopilot's control loops and state estimator. Note that we have used the small angle approximation for the roll angle in the heading rate dynamics, i.e., $\dot{\psi}_j = (\alpha_g/s_j)\tan\phi_j \approx (\alpha_g/s_j)\phi_j$. The reasons for this are discussed next.

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 synchronized zero-order holds of T_s seconds on the UAVs' control inputs. The corresponding discrete-time equations of motion are denoted by

$$\xi^{(j)^+} = f_a(\xi^{(j)}, u_j) = \begin{bmatrix} f_{a,p}(\xi^{(j)}, u_j) \\ f_{a,o}(\xi^{(j)}, u_j) \end{bmatrix},$$

where the subscript "a" refers to the fact that the discrete-time dynamics are those of an air vehicle. Furthermore, we have partitioned the discrete-time dynamics into the position dynamics $f_{a,p}(\xi^{(j)}, u_j) : \mathbb{R}^4 \to \mathbb{R}^2$ corresponding to $\mathfrak{p}_j$ and the orientation dynamics $f_{a,o}(\xi^{(j)}, u_j) : \mathbb{R}^4 \to \mathbb{R}^2$ corresponding to (ψ_j, ϕ_j). As the system of differential equations given by (2) does not readily emit an analytic solution for the position update equations $f_{a,\mathfrak{p}}(\xi^{(j)}, u_j)$, we utilize a second-order Lie series approximation that is well suited for numerical optimization. Namely, we approximate the discrete-time position dynamics as follows:

$$f_{a,p}(\xi^{(j)}, u_j) \approx \begin{bmatrix} I_2 & \mathbf{0}_{2\times 2} \\ \mathbf{0}_{2\times 2} & \mathbf{0}_{2\times 2} \end{bmatrix} \left(\xi^{(j)} + T_s F(\xi^{(j)}, u_j) + \frac{T_s^2}{2} \frac{\partial F}{\partial \xi^{(j)}} F(\xi^{(j)}, u_j) \right),$$

where I_n is an $n \times n$ identity matrix, $\mathbf{0}_{n\times n}$ is an $n \times n$ matrix of zeros, and $\partial F/\partial \xi^{(j)}$ is the Jacobian of (2) and yields the following:

$$\frac{\partial F}{\partial \xi^{(j)}} F(\xi^{(j)}, u_j) = \begin{pmatrix} \alpha_g \xi_4^{(j)} \sin\xi_3^{(j)} \\ -\alpha_g \xi_4^{(j)} \cos\xi_3^{(j)} \\ (\alpha_g \alpha_{\phi,j}/s_j)(\xi_4^{(j)} - u_j) \\ \alpha^2(\xi_4^{(j)} - u_j) \end{pmatrix}. \tag{3}$$

To determine the orientation dynamics, we note that

$$\frac{d}{dt}\begin{bmatrix} \xi_3^{(j)}(t) \\ \xi_4^{(j)}(t) \end{bmatrix} = \begin{bmatrix} 0 & \alpha_g/s_j \\ 0 & -\alpha_{\phi,j} \end{bmatrix} \begin{bmatrix} \xi_3^{(j)} \\ \xi_4^{(j)} \end{bmatrix} + \begin{bmatrix} 0 \\ \alpha_{\phi,j} \end{bmatrix} u_j = A_o \begin{bmatrix} \xi_3^{(j)} \\ \xi_4^{(j)} \end{bmatrix} + b_o u_j.$$

By defining the matrix $\Phi \in \mathbb{R}^{2\times 2}$ and vector $\Gamma \in \mathbb{R}^2$ as $\Phi := \exp(A_o T_s)$ and $\Gamma := \int_0^{T_s} \exp(A_o \tau) b_o d\tau$, respectively, we have

$$f_{a,o}(\xi^{(j)}, u_j) = \Phi \begin{bmatrix} \xi_3^{(j)} \\ \xi_4^{(j)} \end{bmatrix} + \Gamma u_j.$$

Before presenting the target's dynamics, we note that the use of $\tan\phi$ for the heading rate dynamics in (2) leads to a more complicated expression for (3), which includes multiplication by $\tan\phi$ and division by $\cos^2\phi$. Its use would also necessitate an approximation to the heading rate dynamics as well. Furthermore, if one wished to employ a higher order Lie series for the solution of (2), the use of $\tan\phi$ adds even more complexity to the results. As such complexity can pose a great deal of difficulty for numerical optimization, we have avoided the use of $\tan\phi$ in (2). Nonetheless, we shall see that the small angle and Lie series approximations lend themselves to tractable nonlinear optimizations that can be solved in realtime to produce good tracking performance.

0.2.2 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 target's planar position is denoted by $\mathfrak{p}_g \in \mathbb{R}^2$ and is measured in the same local East-North-Up coordinate frame as the UAVs. The corresponding velocity vector is $\mathfrak{v} \in \mathbb{R}^2$, and the acceleration input vector is $d \in \mathbb{R}^2$. The overall target state is $\eta = [\eta_1\ \eta_2\ \eta_3\ \eta_4]^\top \in \mathbb{R}^4$, where $(\eta_1, \eta_2) := \mathfrak{p}_g$ and $(\eta_3, \eta_4) := \mathfrak{v}$. 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

$$\eta^+ = f_g(\eta, d) = A\eta + Bd, \tag{4}$$

where

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

To keep the problem realistic and well-posed, we take the target's acceleration input d to belong to

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

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 $x := (\xi^{(1)}, \xi^{(2)}, \eta) \in \mathbb{R}^{12}$. The overall dynamics are thus given by

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

0.2.3 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 $y_a^{(j)} \in \mathbb{R}^4$ and is given by

$$y_a^{(j)} = \xi^{(j)} + n_a^{(j)}, \qquad n_a^{(j)} = [n_{a,1}^{(j)}\ n_{a,2}^{(j)}\ n_{a,3}^{(j)}\ n_{a,4}^{(j)}]^\top, \tag{5}$$

where $(n_{a,1}^{(j)}, n_{a,2}^{(j)}) \sim \mathcal{N}(0, \sigma_{\mathfrak{p}}^2 I_2)$, $(n_{a,3}^{(j)}, n_{a,4}^{(j)}) \sim \mathcal{N}(0, \sigma_\theta^2 I_2)$, $\sigma_{\mathfrak{p}}^2$ is the variance of the uncorrelated noise on the UAV's North and East position coordinates, and σ_θ^2 is the variance on the UAV's Euler angles, namely roll and yaw (heading).

To describe the measurement error on the target's position, we primarily follow the work of [1] and summarize key ideas here. Since we use similar notation for vectors and matrices, we refer the reader to the said work for illustrations of these quantities. We begin by noting that 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 $T_S^T(\boldsymbol{\theta}_j)$ that relates the coordinates of the line-of-sight vector $\mathfrak{u}_j^S$, which is measured 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 $\boldsymbol{\theta}_j \in \mathbb{R}^3$. 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 $\mathfrak{u}_j^S$. 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

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

where

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

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 $\hat{r}_j = (h_0 - \hat{\mathfrak{s}}_{j,3})/\hat{\mathfrak{u}}_{j,3}$, 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, for example, 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 $\boldsymbol{\theta}_j$.

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 $\hat{\boldsymbol{\theta}}_j$ given in (6). The exact analytic expression for $P_{\mathfrak{o},j}$ is derived in both [1] and [13] 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 $P_j \in \mathbb{R}^{2\times 2}$.

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:

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

where $\mathcal{P} = (P_1^{-1} + P_2^{-1})^{-1}$, $\hat{\mathfrak{p}}_g^{(j)} = [I_2 \ \mathbf{0}_{2\times 1}]\hat{\mathfrak{o}}_j$, $\mathbf{0}_{m\times n}$ denotes the $m \times n$ zeros matrix, and n_g is noise that is adequately approximated by a zero-mean Gaussian distribution with covariance $\mathcal{P}$. 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. Note that the expressions for the true covariances P_j are based on the true Euler angles and true vehicle altitudes, which are unknown quantities, and hence estimates of the covariances $\hat{P}_j$ based on the raw Euler angle estimates $\hat{\boldsymbol{\theta}}_j$ and raw altitude estimates $\hat{\mathfrak{s}}_{j,3}$ are used in (7) in practice.

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

$$y := (y_a^{(1)}, y_a^{(2)}, y_g) = Cx + n, \tag{8}$$

where $n := (n_a^{(1)}, n_a^{(2)}, n_g)$ and $C := [I_{10} \ \mathbf{0}_{10\times 2}]$. Since the target velocity is not measured directly, the control law used in this framework will be based on output feedback.

0.3 Robust Output-Feedback MPC/MHE

If only one UAV is considered in the target tracking problem, the problem can be posed as a two-player zero-sum game in which the UAV tries to minimize its squared 3D distance to the target while the target tries to maximize this same quantity. Moreover, by seeking to minimize a scalar quantity proportional to its individual covariance, the UAV's aim becomes that of gathering the best individual vision-based measurements of the target. 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. This practice translates into the UAVs gathering the best joint vision-based measurements of

the target. 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 that the UAVs (Player 1) would like to minimize and the target (Player 2) would like to maximize:

$$g(x) := \beta_1 \frac{r_1^2 r_2^2}{r_1^2 + r_2^2} + \beta_2 (r_1^2 + r_2^2) \tag{9}$$

In this criterion, β_1 and β_2 are positive weighting constants. The first term in (9) 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, which improves measurement quality just as in [13]. The second term in (9) 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 criteria may be considered, we aim to utilize a simpler expression than those found in [12] and [13] that achieves similar behavior and lends itself to efficient numerical computation. We shall see that the distance coordination of [13] is indeed induced by choosing the criterion (9).

For our control approach, we use the output-feedback MPC with MHE approach described in [18]. This requires us to solve a finite-horizon online optimization problem at each step k. Solving this online optimization problem uses the last $L \in \mathbb{Z}_{\geq 1}$ output measurements in order to give us an estimate of the current state at step k (the MHE problem) while simultaneously computing policies for both Players 1 and 2 to use for the next $K \in \mathbb{Z}_{\geq 1}$ steps into the future (the MPC problem).

As previously discussed in §0.1, the problem we address in this chapter is similar to the problem of coordinating UAVs for target tracking described in [22], but here we incorporate the more realistic UAV model given by (2) that includes unknown wind disturbances. To solve this problem, while simultaneously computing the state estimates and future control sequence, we can estimate the unknown wind w by including the estimate $\hat{w} \in \mathcal{W}$ as an optimization variable, where $\mathcal{W}$ is given by (1). More details about using this combined MPC/MHE approach for adaptive control and learning can be found in [27].

Given a discrete-time signal $z : \mathbb{Z}_{\geq 0} \to \mathbb{R}^n$ and two steps $k_0, k \in \mathbb{Z}_{\geq 0}$ with $k_0 \leq k$, we denote by $z_{k_0:k}$ the sequence $\{z_{k_0}, z_{k_0+1}, ..., z_k\}$. Then, formally, 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

$$J_k(x_{k-L}, u_{k-L:k+K-1}, d_{k-L:k+K-1}, y_{k-L:k}) :=$$

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

where $g(x)$ is given in (9), the second term in the right hand side penalizes the difference between the roll command pair u_ℓ and the roll angle pair $\varphi_\ell := (\phi_{1,\ell}, \phi_{2,\ell}) \in \mathbb{R}^2$ at time ℓ, and λ_u, λ_d, and λ_n are positive scalars. Here $\phi_{j,\ell}$ refers to UAV j's roll angle at time ℓ. The scalar λ_u is a tuning weight that determines the penalty for the UAVs applying roll commands that are different from their current roll angles. Therefore, λ_u may be chosen large enough to ensure that the rate of change of the roll angle does not exceed an actuator's limitations. The scalar λ_d is a tuning weight that determines how much of a penalty the target incurs when using large accelerations. Finally, the scalar λ_n is a tuning weight that can be thought of as determining how much of a penalty is incurred for choosing less likely values for the measurement noise (in this case, larger values).

Defining the criterion as in (10) emphasizes the dependence 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 input sequence for the UAVs $u_{k-L:k+K-1}$. The input sequence $u_{k-L:k+K-1}$ is comprised of 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 step $k \in \mathbb{Z}_{\geq 0}$, 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

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

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 $\hat{x}_{k-L|k}$, $\hat{d}_{k-L:k+K-1|k}$, and $\hat{w}$ as (worst-case) estimates of the initial state, target input, and wind speed, respectively, based on the past inputs $u_{k-L:k-1}$ and outputs $y_{k-L:k}$ available at time k. The past measurement noises $n_{k-L:k}$ are not included as explicit optimization variables because they are uniquely defined by the output equation (8) and the choices of the optimization variables $\hat{x}_{k-L|k}$ and $\hat{d}_{k-L:k-1|k}$.

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

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

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

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

0.4 Simulation Results

Now we demonstrate the effectiveness of the MPC/MHE control approach to the problem of two fixed-wing UAVs performing vision-based target tracking of a moving ground vehicle. To underscore the robustness of the approach to real-world conditions such as constant winds and unmodeled aircraft dynamics, we present the results of several simulations that were performed in realtime using the flight simulator Aviones [28]. This simulator utilizes an aircraft model with six degrees of freedom and allows the user to simulate constant winds. Furthermore, we simulate three different scenarios. The first scenario involves tracking a constant-velocity target, which represents the "easiest" target to be tracked. The second scenario involves tracking an evasive target, where the worst-case target acceleration d^*_k, computed by solving the optimization (11), is applied as the target's actual input. This results in the most "difficult" target to be tracked. Finally, the third scenario uses an experimental target log that was generated by a human physically driving a vehicle casually yet unpredictably. This scenario represents the most realistic target motion. In each of these scenarios, realistic levels of noise as well as unmeasured wind disturbances are included in the models presented in Section 0.2. For each of the three target scenarios, three different wind speeds are also considered, namely 0 m/s, 3 m/s, and 6 m/s.

The parameters pertaining to all three simulation scenarios are provided in Table 1. Given the maximum UAV turn rate, the total time it takes a UAV to make a full loop is $2\pi s_1/(\alpha_g \tan \bar{u}) \approx 13.72$ seconds, where we have used the relationship of bank angle to turn rate given by Equation 9.20 of [24]. Hence the future planning horizon of $KT_s = 10$ seconds allows the UAVs to consider the impact of beginning to loop around the target. Additionally, in the cost function (9), the coefficients β_1 and β_2 are chosen to place a greater emphasis on distance coordination than on keeping individual distances to the target small. The maximum target speed of 10 m/s is chosen smaller than the fixed airspeed of the UAVs (15 m/s) in order to make the problem well-posed for wind speeds less than 5 m/s. The UAVs fly at fixed altitudes of 40 and 45 m, respectively. Other parameters included in Table 1 include the scalar cost function weights λ_u, λ_d, and λ_n, the forward and backward horizon lengths K and L, respectively, as well as the variances for the UAVs' positions, altitudes, and Euler angles.

To solve the nonlinear, non-convex, min-max optimization problem (11) at each time step, we use a primal-dual-like interior point method that is designed specifically for this MPC/MHE approach and is described in [29]. Table 2 shows the computation times using the C programming language on a laptop with a 2.3 GHz Intel® Core™ i7 processor for each scenario that we investigated.

Table 1: Simulation Parameters

Parameter	Description	Value	Units
$\bar{u}$	Max UAV roll angle	35	deg.
(s_1, s_2)	UAV airspeeds	$(15, 15)$	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_{p}^2	Altitude variance	4^2	m^2
σ_{θ}^2	Euler angle variance	3^2	deg.^{2}
(h_1, h_2)	UAV altitudes	$(40, 45)$	m
β_1	Coord. coefficient	0.04	-
β_2	Dist. coefficient	0.002	-
λ_u	UAV roll-angle coefficient	5	-
λ_d	Target acceleration coefficient	10	-
λ_n	Noise coefficient	40	-
T_s	Sampling period	1	s
L	Backward horizon	8	-
K	Forward horizon	10	-

If the solver does not converge to a solution within a reasonable number of iterations (a maximum of 140 in this case), the optimization is reattempted up to a maximum of 4 additional times with a new random start condition each time. If the solver fails to converge after 5 attempts, the optimal control from the previous time step is used, i.e., $u_k = \hat{u}^*_{k|k-1}$. The last column in the table shows the percentage of time steps for which the solver successfully converges in one of those 5 allowed attempts and the total computation time is less than the sampling time of $T_s = 1$ second. Even for this difficult nonlinear, non-convex problem, the solver converges over 90% of the time for every scenario. We note that the percent of convergence is highest for the evasive target scenario because the true trajectories in this scenario best match the solution of the min-max optimization which computes trajectories based on worst-case evasive target motion.

The other columns in Table 2 show the minimum and mean total computation times, as well as their standard deviations for all of the time steps in which the solver converged. Regardless of the scenario, the minimum computation time is about 30 milliseconds. The mean computation times vary, but all of them are between 37 milliseconds and 90 milliseconds. These computation times clearly show that this problem can be solved in real-time using the primal-dual-like interior point method described in [29].

To determine the steady-state tracking performance of each scenario, 5 minutes of steady-state behavior is considered, where the effects of initial conditions have been removed by discarding 30 seconds of initial data. During the initial 30 seconds, the UAVs circle a stationary target, and the past measurement buffer is filled.

Table 2: Computation Time and Percent of Convergence

Target Scenario	min. [ms]	mean [ms]	std. dev. [ms]	% converge
Const.-Vel., 0m/s wind	30	37	19	93%
Const.-Vel., 3m/s wind	30	90	61	94%
Const.-Vel., 6m/s wind	29	82	56	95%
Evasive, 0m/s wind	31	49	26	97%
Evasive, 3m/s wind	31	45	26	97%
Evasive, 6m/s wind	30	45	30	99%
Log, 0m/s wind	30	42	24	92%
Log, 3m/s wind	30	51	37	91%
Log, 6m/s wind	30	54	29	95%

0.4.1 Constant-Velocity Target

We first consider a constant-velocity target, i.e., $d_k \equiv 0$ in equation (4), where the target travels at 7.5 m/s (one half of the UAVs' fixed speed) in the positive x-direction. The UAVs' inputs are computed by solving the min-max optimization (11) and hence are prepared for worst-case target motion as well as worst-case noise and wind.

The results for this scenario with the UAVs subjected to 6 m/s wind to the East, i.e., $w = [6 \;\; 0]^\top$, where w is presented in (2), are provided in Figures 1, 2, and 3. In Figure 1, one can see how the UAVs make loops so that their average (ground) speed matches that of the target. Figure 2 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 solid black curve depicting $\min\{r_1, r_2\}$. This can also be seen by the alternating blue (dashed) and red (dotted) curves that depict the individual UAV distances to the target. In this case, the wind is blowing at 6 m/s in the positive x-direction, so it is helping the UAVs catch up to the target. The estimate of the wind is initialized at $[0 \;\; 0]^\top$ but quickly converges close to the true value of the wind and stays close to the true value for the rest of the simulation, as shown in Figure 3.

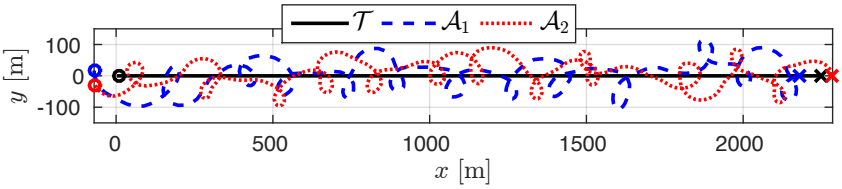

Figure 1: Trajectories of two UAVs, subject to 6 m/s wind, tracking a constant-velocity target over a 5-minute window. The starting positions of all vehicles are denoted by an "$\circ$" while the ending positions are indicated by an "$\times$". In the legend, $\mathcal{T}$ corresponds to the target while $\mathcal{A}_1$ and $\mathcal{A}_2$ refer to the UAVs.

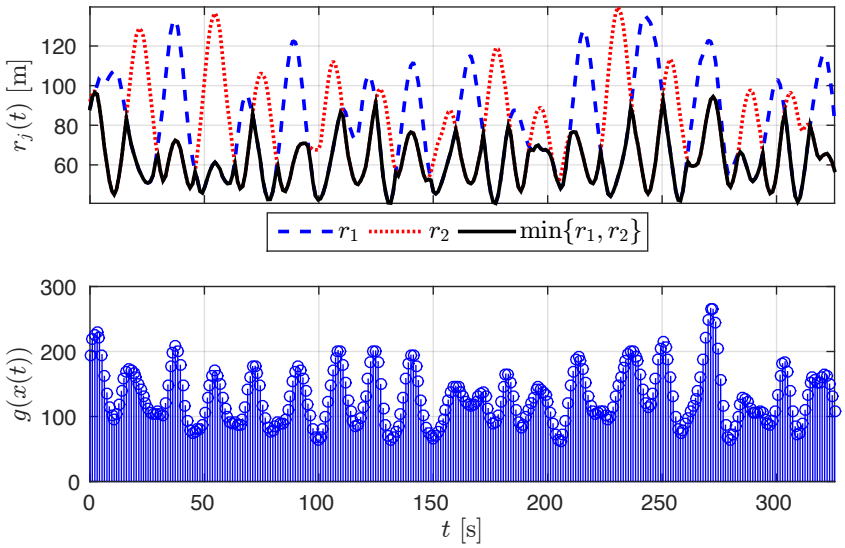

Figure 2: 3D distances r_j and stage cost $g(x)$ for two UAVs, subject to 6 m/s wind, tracking a constant-velocity target. The solid black line shows $\min\{r_1, r_2\}$.

A summary of the performance for this scenario and all three different wind speeds considered is given in Table 3. Note that for the various cases the wind's magnitude changed while its direction remained the same, which holds true for all target motion scenarios considered. From the table, one can see statistics associated with the cost $g(x)$, the individual 3D distances r_1 and r_2, the estimation error $\|\mathfrak{p}_g - \hat{\mathfrak{p}}_g\|$, and the actual measurement noise n_g on the ground target's position. Additionally, the degree of coordination between the UAVs 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. Overall, while the average cost is quite close for the various wind speeds, one should notice that the cost variance $\operatorname{var} g(x)$ is quite different for each case and decreases as the wind speed increases. Additionally, for the wind speed of 6 m/s, the Pearson correlation coefficient was most negative, indicating that the greatest degree of coordination was achieved by the UAVs at this speed. Moreover, the wind speed of 6 m/s in the direction of travel of the target proved to be beneficial to the UAVs' coordination efforts, as it allowed them to generally have at least one agent relatively close to the target and thereby reduce the variability in the cost.

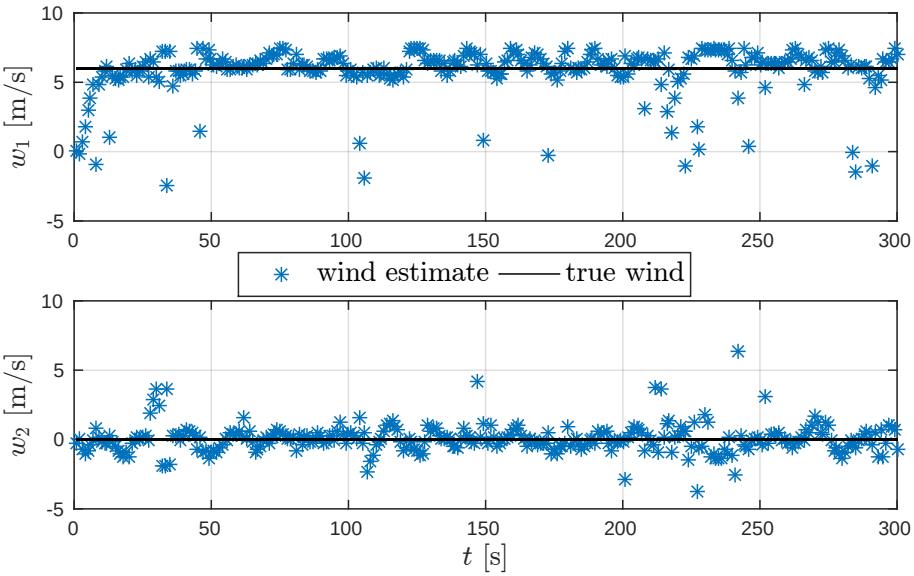

Figure 3: Estimate of the wind speed in the x-direction (top subplot) and y-direction (bottom subplot) for the constant-velocity target scenario.

Table 3: Constant Velocity Target

Statistic	**Wind = 0m/s**	**Wind = 3m/s**	**Wind = 6m/s**	**Units**
avg $g(x)$	132.8	126.0	130.7	m^2
var $g(x)$	3226	2608	1768	m^4
avg$(\min\{r_1, r_2\})$	62.01	62.10	61.22	m
$\max\{r_1, r_2\}$	169.3	138.0	139.7	m
avg $\|\mathfrak{p}_g - \hat{\mathfrak{p}}_g\|$	5.65	5.29	5.31	m
avg $\|n_g\|$	5.44	5.17	5.48	m
ϱ	-0.182	-0.079	-0.430	N/A

0.4.2 Evasive Target

We now consider an evasive target, where the optimal worst-case d_k^* computed from the min-max optimization (11) is applied as the target's actual input. Results for the UAVs flying in 3 m/s winds to the East, i.e., $w = [3 \;\; 0]^\top$ with w given in (2), are provided in Figures 4, 5, and 6.

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. The target also makes use of the wind in this case as the target generally moves in the negative x-direction, thereby forcing the UAVs to fly against the wind.

From Figure 5, the minimum UAV distance from the target frequently peaks above 100 m, which means that both UAVs were often far from the target simultaneously. This is also reflected in the fact that the stage cost peaks at values much higher than in the case of a constant-velocity target. Furthermore, these peaks correspond to the periods of time just after the sharp movements in the target's trajectory shown in Figure 4.

A summary of the performance for this scenario and all three different wind speeds considered is given in Table 4. From the third column of the table, one can see that the evasive target can especially take advantage of the wind in the case of 6 m/s wind speed because the fixed airspeed of the UAVs is 15 m/s, and the maximum speed of the ground target is 10 m/s. Therefore, in this case, the evasive target can move at its maximum speed directly into the wind so that its speed of 10 m/s is faster than the effective ground speed of the UAVs flying directly into the wind (9 m/s), which renders the UAVs unable to catch the target. While the increase in all quantities considered is significant for this case, of particular interest is the fact that the measurement noise magnitude $\|n_g\|$ went up considerably, as it is related to $\min\{r_1, r_2\}$, which also saw a drastic increase. Of course, the increased measurement noise in turn caused a rise in the estimation error magnitude $\|\mathfrak{p}_g - \hat{\mathfrak{p}}_g\|$, which shows the overall importance of coordinating UAV distances to the target when possible.

Overall, the target's evasive maneuvers hinder UAV coordination efforts and thereby increase measurement noise. This is especially evident by the fact that the Pearson correlation coefficient is positive for all evasive target cases, which means that the UAVs were often relatively far from the target at the same time. Additionally, the cost variance was also much higher in all cases in comparison with the constant-velocity target scenario. Nonetheless, for an evasive target, 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 is also only slightly larger than in the case of constant target velocity, provided that the UAVs' groundspeed is able to match or exceed the target's speed at all times.

Table 4: Evasive Target

Statistic	**Wind = 0m/s**	**Wind = 3m/s**	**Wind = 6m/s**	**Units**
$\operatorname{avg} g(x)$	157.3	159.9	343.6	m^2
$\operatorname{var} g(x)$	6466	6956	9673	m^4
$\operatorname{avg}(\min\{r_1, r_2\})$	71.06	69.19	91.47	m
$\max\{r_1, r_2\}$	159.3	173.7	235.5	m
$\operatorname{avg} \|\mathfrak{p}_g - \hat{\mathfrak{p}}_g\|$	5.12	4.88	6.83	m
$\operatorname{avg} \|n_g\|$	5.86	6.14	10.4	m
ϱ	0.582	0.272	0.622	N/A

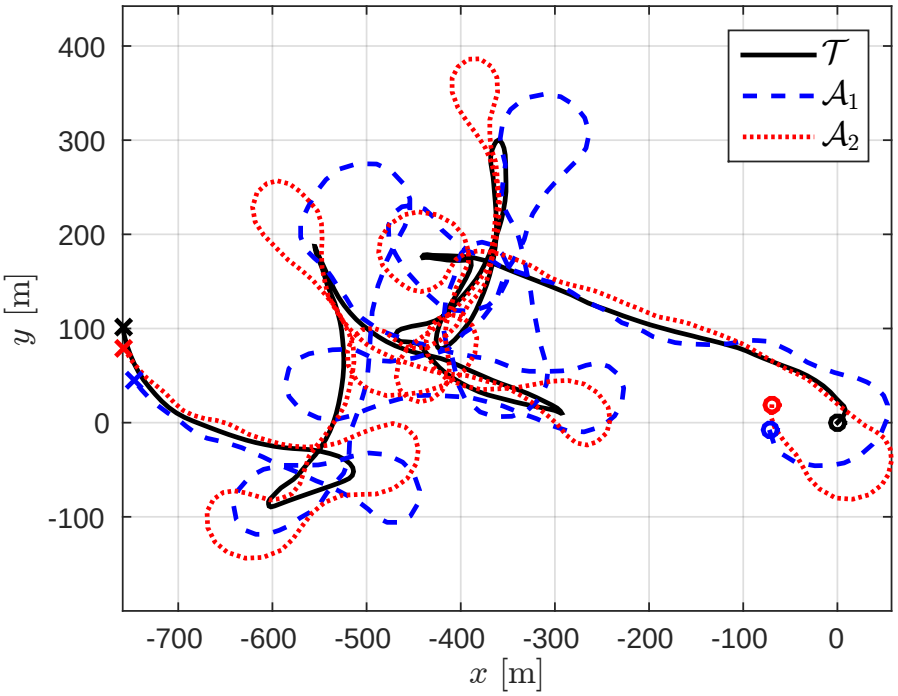

Figure 4: Trajectories of two UAVs, subject to 3 m/s wind, tracking an evasive target over a 5-minute window. The starting positions of all vehicles are denoted by an "$\circ$" while the ending positions are indicated by an "$\times$". In the legend, $\mathcal{T}$ corresponds to the target while $\mathcal{A}_1$ and $\mathcal{A}_2$ refer to the UAVs.

0.4.3 Experimental Target Log

Finally, we now consider target motion taken from the log of a live tracking experiment at Camp Roberts, California, U.S.A., in which a person drove a car casually yet unpredictably. Results with the UAVs subjected to 0 m/s wind are provided in Figures 7, 8, and 9. One should observe from these results that the minimum 3D distance to the target is kept below 100 m and the cost is always less than 300 m^2. Thus, even though the target does not maintain a constant velocity, the UAVs are able to coordinate and keep the cost quite low.

A summary of the performance for this scenario and all three different wind speeds considered is given in Table 5, where the wind direction is to the southeast in all instances of nonzero wind, i.e., for $\|w\| \neq 0$, $w/\|w\| = (1/\sqrt{2})[1 \quad -1]^\top$. From the table, one can see that the average cost and average minimum distances are quite comparable to the constant-velocity scenario, albeit the cost variances are somewhat higher in this case. Nonetheless, the cost variances are still considerably lower than in the evasive case. One should also take notice that the correlation coefficient is always negative, meaning that the UAVs are able to coordinate in all cases considered, whereas the coefficients were always positive

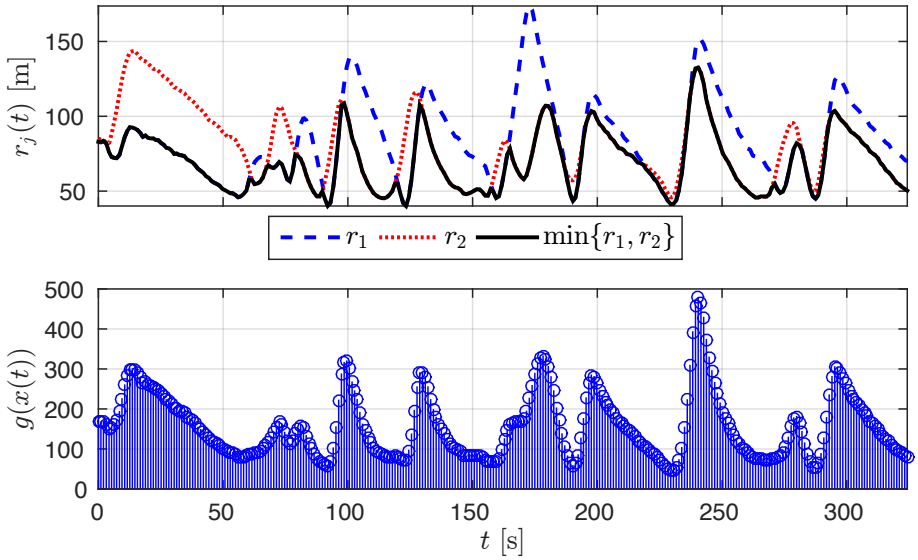

Figure 5: 3D distances r_j and stage cost $g(x)$ for two UAVs, subject to 3 m/s wind, tracking an evasive target. The solid black line shows $\min\{r_1, r_2\}$.

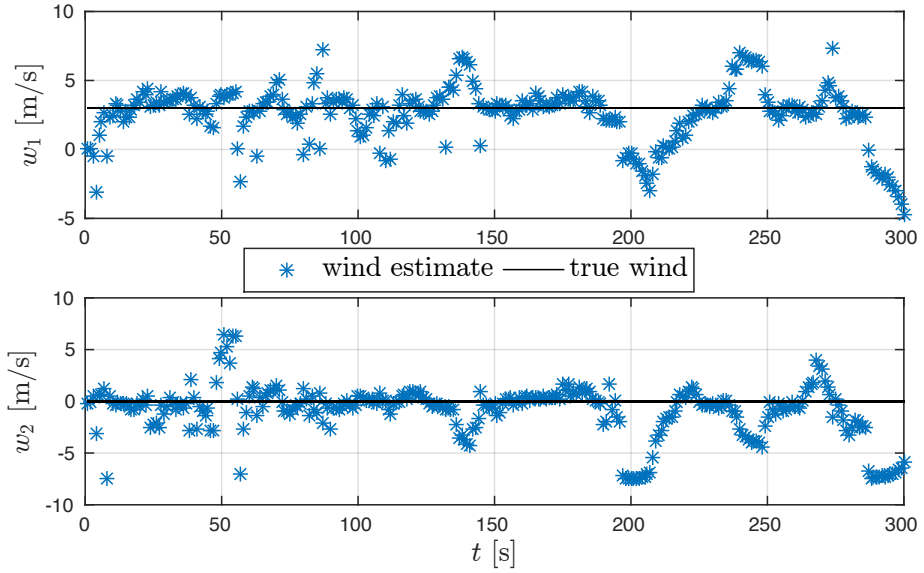

Figure 6: Estimate of the wind speed in the x-direction (top subplot) and y-direction (bottom subplot) for the evasive target scenario.

for an evasive target. Moreover, through distance coordination, the UAVs are

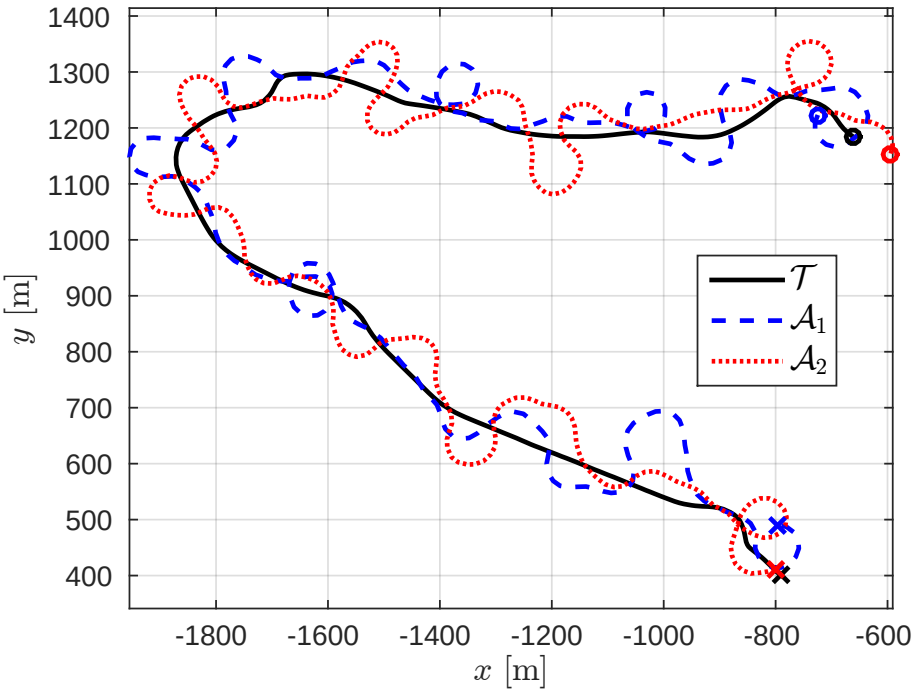

Figure 7: Trajectories of two UAVs, subject to 0 m/s wind, tracking a target moving based on an experimental target log over a 5-minute window. The starting positions of all vehicles are denoted by an "$\circ$" while the ending positions are indicated by an "$\times$". In the legend, $\mathcal{T}$ corresponds to the target while $\mathcal{A}_1$ and $\mathcal{A}_2$ refer to the UAVs.

able to keep the average measurement noise n_g on the target's position smaller than in the evasive scenario, wherein the target was able to thwart the UAVs' coordination efforts and drive the average magnitude of the state-dependent measurement noise higher.

Table 5: Experimental Target Log

Statistic	**Wind = 0m/s**	**Wind = 3m/s**	**Wind = 6m/s**	**Units**
avg $g(x)$	117.8	135.0	132.6	m^2
var $g(x)$	2269	4112	4406	m^4
avg$(\min\{r_1, r_2\})$	57.98	62.71	60.36	m
$\max\{r_1, r_2\}$	167.4	204.8	169.6	m
avg $\|\mathfrak{p}_g - \hat{\mathfrak{p}}_g\|$	5.99	6.45	5.68	m
avg $\|n_g\|$	4.98	5.35	5.63	m
ϱ	-0.219	-0.030	-0.123	N/A

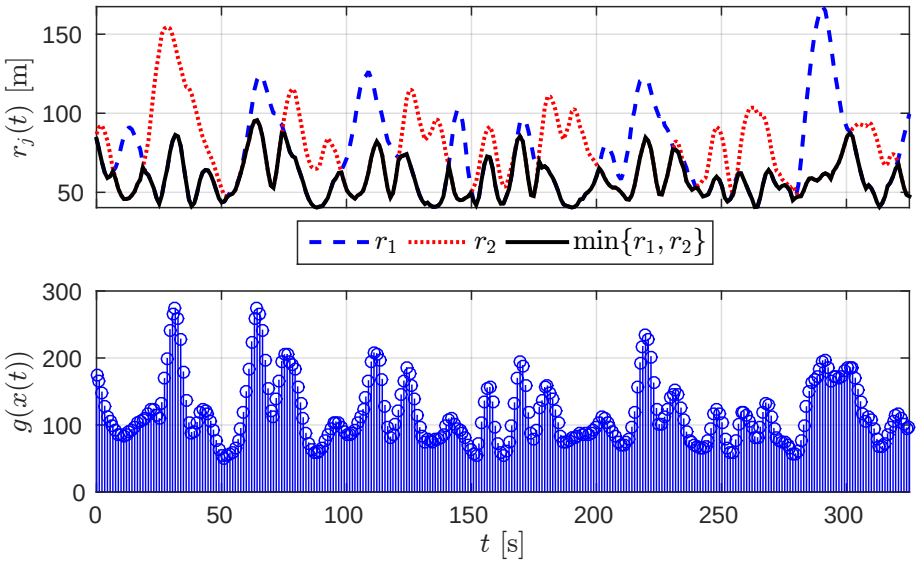

Figure 8: 3D distances r_j and stage cost $g(x)$ for two UAVs tracking a target moving according to an experimental target log. The minimum distance r_j is shown as a solid black line.

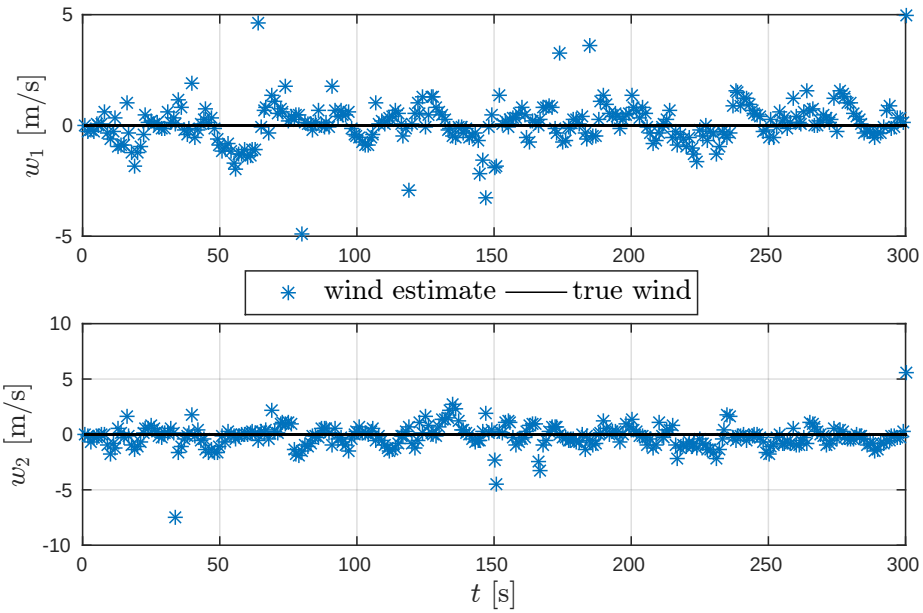

Figure 9: Estimate of the wind speed in the x-direction (top subplot) and y-direction (bottom subplot) for the experimental target log scenario.

0.5 Conclusion and Future Work

We have presented an effective, optimization-based control approach for two fixed-wing UAVs to robustly perform vision-based target tracking of a moving ground vehicle under realistic conditions. Namely, we utilized the novel approach based on min-max MPC combined with MHE that was first presented in [22] and adapted the approach to use fourth-order UAV models with roll dynamics along with the parameter estimation scheme of [27] to estimate constant winds. With simulations performed in realtime using aircraft models having six degrees of freedom for the UAVs and target logs taken from live tracking experiments, we showed that the MPC/MHE approach produces a controller that is robust to unmodeled aircraft dynamics, the unmeasured environmental disturbance of wind, and information about the target state that is both noisy and incomplete.

Regarding the effectiveness of the approach for target tracking, the robust optimal controller enables the UAVs to coordinate their distances to the target when the wind speeds and target motion are not exceptionally adverse, even though the target and wind velocities are unmeasured quantities. This coordination keeps the state-dependent measurement noise at modest levels, thereby allowing the UAVs to effectively track the target. When the target is evasive, the UAVs are generally not able to coordinate their distances to the target. Nonetheless, for wind conditions in which each UAV's groundspeed always exceeds the target's speed, the UAVs are still able to track the evasive target, albeit at a somewhat reduced performance level. Finally, we note that casual, yet unpredictable, target motion from live tracking experiments with a nonholonomic ground vehicle results in performance that is considerably better than that arising from the worst-case evasive motion of a 2D double integrator, which has fewer motion constraints than a typical ground vehicle and is therefore more difficult for the underactuated UAVs to track.

Overall, we have shown that the output-feedback MPC/MHE approach is a viable approach for addressing high-dimensional, very nonlinear (non-convex) problems involving the robust coordination of mobile robots under realistic settings. We foresee this being a powerful tool for solving nonlinear, multi-agent control problems with non-convex objectives that require explicit robustness to both adversarial disturbances and parametric uncertainty in the dynamics, as well as noisy, partial information about the system state.

Future research interests involve addressing the problem of target tracking for a single UAV in three dimensions, as well as the possible use of airspeed control to provide an additional degree of freedom in the motion planning. Additionally, we wish to investigate both the behavior and tracking performance of three UAVs with the present problem formulation in order to evaluate whether or not there is a diminishing return on investment.

Bibliography

[1] Mallick, M. (2007) Geolocation using video sensor measurements, in *IEEE International Conf. on Information Fusion*, Quebec, Canada.

[2] Frew, E.W. (2007) Sensitivity of cooperative geolocation to orbit coordination, in *AIAA Guid., Nav., and Control Conf.*, Hilton Head, SC, pp. 3869–3892.

[3] Ma, L. and Hovakimyan, N. (2013) Cooperative target tracking in balanced circular formation: Multiple UAVs tracking a ground vehicle, in *American Control Conf.*, IEEE, pp. 5386–5391.

[4] Kingston, D.B. (2007) *Decentralized Control of Multiple UAVs for perimeter and target surveillance*, Ph.D. thesis, Brigham Young University.

[5] Peterson, C. and Paley, D.A. (2011) Multivehicle coordination in an estimated time-varying flowfield. *Journal of guidance, control, and dynamics*, **34** (1), 177–191.

[6] Summers, T.H., Akella, M.R., and Mears, M.J. (2009) Coordinated standoff tracking of moving targets: Control laws and information architectures. *AIAA Journal of Guidance, Control, and Dynamics*, **32** (1), 56–69.

[7] Kim, S., Oh, H., and Tsourdos, A. (2013) Nonlinear model predictive coordinated standoff tracking of a moving ground vehicle. *Journal of Guidance, Control, and Dynamics*, **36** (2), 557–566.

[8] Sinha, A., Kirubarajan, T., and Bar-Shalom, Y. (2005) Autonomous ground target tracking by multiple cooperative UAVs, in *Aerospace Conf., 2005 IEEE*, pp. 1–9, doi:10.1109/AERO.2005.1559601.

[9] Miller, S.A., Harris, Z.A., and Chong, E.K.P. (2009) A POMDP framework for coordinated guidance of autonomous UAVs for multitarget tracking. *EURASIP J. Adv. Signal Process*, pp. 1–17.

[10] Thrun, S., Burgard, W., and Fox, D. (2005) *Probabilistic Robotics*, The MIT Press.

[11] Stachura, M., Carfang, A., and Frew, E.W. (2009) Cooperative target tracking with a communication limited active sensor network, in *International Workshop on Robotic Wireless Sensor Networks*, Marina Del Ray, CA.

[12] Ding, C., Morye, A.A., Farrell, J.A., and Roy-Chowdhury, A.K. (2012) Coordinated sensing and tracking for mobile camera platforms, in *American Control Conf.*, IEEE, pp. 5114–5119.

[13] Quintero, S.A.P., Papi, F., Klein, D.J., Chisci, L., and Hespanha, J.P. (2010) Optimal UAV coordination for target tracking using dynamic programming, in *IEEE Conf. on Decision and Control*, Atlanta, GA.

[14] Quintero, S.A.P., Ludkovski, M., and Hespanha, J.P. (2015) Stochastic optimal coordination of small UAVs for target tracking using regression-based dynamic programming. *Journal of Intelligent & Robotic Systems*, pp. 1–28, doi:10.1007/s10846-015-0270-7. URL `http://dx.doi.org/10.1007/s10846-015-0270-7`.

[15] Triplett, B.I., Klein, D.J., and Morgansen, K.A. (2007) Distributed estimation for coordinated target tracking in a cluttered environment, in *Proc. of Robocomm*, Athens, Greece.

[16] Lalish, E., Morgansen, K., and Tsukamaki, T. (2007) Oscillatory control for constant-speed unicycle-type vehicles, in *IEEE Conf. on Decision and Control*.

[17] Regina, N. and Zanzi, M. (2011) UAV guidance law for ground-based target trajectory tracking and loitering, in *Aerospace Conf.*, IEEE, doi:10.1109/AERO.2011.5747522.

[18] Copp, D.A. and Hespanha, J.P. (2014) Nonlinear output-feedback model predictive control with moving horizon estimation, in *IEEE Conf. on Decision and Control*, Los Angeles, California, pp. 3511–3517.

[19] Bemporad, A. and Morari, M. (1999) Robust model predictive control: A survey, in *Robustness in identification and control*, Springer, pp. 207–226.

[20] Chen, H.S. and Allgöwer, F. (1997) A game theoretic approach to nonlinear robust receding horizon control of constrained systems, in *American Control Conf.*

[21] Rawlings, J.B. and Bakshi, B.R. (2006) Particle filtering and moving horizon estimation. *Computers & Chemical Engineering*, **30** (10), 1529–1541.

[22] Quintero, S.A.P., Copp, D.A., and Hespanha, J.P. (2015) Robust UAV coordination for target tracking using output-feedback model predictive control with moving horizon estimation, in *American Control Conf.*, pp. 3758–3764.

[23] Collins, G.E., Stankevitz, C.R., and Liese, J. (2011) Implementation of a sensor guided flight algorithm for target tracking by small UAS, in *Ground/Air Multi-Sensor Interoperability, Integration, and Networking for Persistent ISR II*, vol. 8047, SPIE, vol. 8047.

[24] Beard, R.W. and McLain, T.W. (2012) *Small unmanned aircraft: Theory and practice*, Princeton University Press.

[25] Quintero, S.A.P. and Hespanha, J.P. (2014) Vision-based target tracking with a small UAV: Optimization-based control strategies. *Control Engineering Practice*, **32**, 28 – 42, doi:10.1016/j.conengprac.2014.07.007.

[26] Saunders, J. and Beard, R.W. (2008) Tracking a target in wind using a micro air vehicle with a fixed angle camera, in *American Control Conf.*, Seattle, WA, pp. 3863–3868.

[27] Copp, D.A. and Hespanha, J.P. (2016, in press) Addressing adaptation and learning in the context of MPC and MHE. *Control of Complex Systems: Theory and Applications*.

[28] (2007) Aviones. Avaliable at `http://aviones.sourceforge.net`.

[29] Copp, D.A. and Hespanha, J.P. (2014) Nonlinear output-feedback model predictive control with moving horizon estimation, *Tech. Rep.*, Univ. California, Santa Barbara. `http://www.ece.ucsb.edu/~hespanha/techrep.html`.