2019 American Control Conference (ACC)
Philadelphia, PA, USA, July 10-12, 2019

LQG Reference Tracking with Safety and Reachability Guarantees under False Data Injection Attacks

Luyao Niu, Zhouchi Li, and Andrew Clark

Abstract— **Control systems are increasingly targeted by malicious adversaries, who may inject spurious sensor measurements in order to bias the controller behavior and cause suboptimal performance or safety violations. This paper investigates the problem of tracking a reference trajectory while satisfying safety and reachability constraints in the presence of such false data injection attacks. We consider a linear, time-invariant system with additive Gaussian noise in which a subset of sensors can be compromised by an attacker, while the remaining sensors are regarded as secure. We propose a control policy in which two estimates of the system state are maintained, one based on all sensors and one based on only the secure sensors. The optimal control action based on the secure sensors alone is then computed at each time step, and the chosen control action is constrained to lie within a given distance of this value. We show that this policy can be implemented by solving a quadratically-constrained quadratic program at each time step. We develop a barrier function approach to choosing the parameters of our scheme in order to provide provable guarantees on safety and reachability, and derive bounds on the probability that our control policies deviate from the optimal policy when no attacker is present. Our framework is validated through numerical study.**

I. Introduction

Control systems increasingly rely on real-time measurements from distributed sensors in order to make autonomous decisions. Malicious adversaries may attempt to degrade system performance by introducing false sensor measurements that induce suboptimal control decisions. These false measurements may be introduced by spoofing the sensed physical invariant, as in reported attacks on GPS and other navigation sensors [1], physically compromising the sensor [2], or hijacking the communication channel between the sensor and controller [3]. Spoofing navigation sensors in particular has been highlighted as a critical threat to safety and performance of autonomous vehicles [4].

There is an extensive literature on modeling and mitigating false data injection attacks on control systems [5]–[7]. The worst-case impact of false data injection attacks was considered in [8]. Methodologies for detecting false data [9] and resilient state estimation in the presence of attacks [10], [11] have also been proposed.

At present, less attention has been given to closed-loop control in the presence of possible false data attacks. The problem of Linear Quadratic Gaussian (LQG) control under false data injection attacks was considered in [12]. In [12], the system first chooses a control policy mapping the observed measurements to a sequence of control actions. Based on this policy, the adversary chooses a strategy for introducing false measurements based on the state of the vehicle. The goal of the system is to minimize the worst-case value of a quadratic cost function while the adversary attempts to maximize the cost function.

While an adversary can degrade the system performance by increasing the cost function over a given time horizon, a potentially more devastating attack is to wait for an opportune moment and then introduce false measurements to drive the system to an unsafe state. Such attacks have forced UAVs to land via GPS spoofing [13] and caused autonomous vehicles to steer off the road and crash [14]. While simply ignoring measurements from easily-spoofed sensors could limit the impact of such attacks, it also reduces the performance and safety of systems under normal (non-adversarial) operating conditions. A new approach is therefore needed to ensure safety and reachability in the presence of false data injection without compromising performance.

In this paper, we consider the problem of reference tracking for linear systems in the presence of additive noise and false data injection attacks with safety and reachability constraints. We formulate the problem of selecting a control policy that minimizes the deviation of the system state from the reference trajectory while also ensuring that, when an adversary is present, the system remains outside of an unsafe region and reaches a given region with a desired probability. Our approach is to proactively limit the set of control inputs in order to ensure a safety constraint holds with a desired probability. We make the following specific contributions:

- We propose a set of control laws in which the control input is within a prespecified bound of the optimum control input when the possibly-compromised sensors are removed. We demonstrate that this policy can be executed by solving a quadratically-constrained quadratic program (QCQP) at each time step.
- We present an algorithm for selecting the maximum deviation of the control input in order to ensure a provable bound on the safety property. Our algorithm is based on the barrier method and we show it can be implemented in an offline fashion by solving a sequence of semidefinite programs.
- We analyze the optimality of our approach. In particular, we derive a lower bound on the probability that the utility achieved by our proposed control policy is equal

L. Niu, Z. Li, and A. Clark are with the Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, MA 01609, USA. L. Niu and Z. Li contributed equally to the work. {lniu,zli4,aclark}@wpi.edu

This work was supported by NSF grants CNS-1544173 and CNS-1656981.

978-1-5386-7926-5/$31.00 ©2019 AACC

to the utility achieved by an optimal LQG policy that ignores the adversary.

- Our results are validated through numerical study, in which our approach closely tracks a reference signal while avoiding an unsafe region and providing a cost that is comparable to the optimal LQG control in a benign environment.

The paper is organized as follows. Section II presents the related work. Section III presents the system and adversary models. Section IV presents the problem formulation and solution approach. Section V analyzes the optimality guarantees of this approach. Section VI contains simulation results. Section VII concludes the paper.

II. Related Work

Reference tracking has been extensively studied in the existing literature, including characterization of the optimal control and estimation policies in the LQG case [15], as well as extensions to nonlinear [16] and constrained [17] systems. These existing works, however, focus on systems without intelligent adversaries.

Resilient control in adversarial environments has received relatively recent research attention. Most of the work on false data injection attacks has focused on resilient state estimation and fault detection. Optimal attack strategies were studied in [8]. Fundamental limits on resilient state estimation were presented in [9]. Tractable estimators using Luenberger observers were proposed in [10]. Extensions to noisy systems are considered in [11]. Closing the control loop in the presence of attacks has been less studied. In [12], linear quadratic Gaussian control under false data injection attacks was studied, although reference tracking and safety and reachability constraints were not considered, leading to a fundamentally different approach from this work.

Our approach leverages the barrier function method, which has been developed for deterministic [18] and stochastic [19] systems. To the best of our knowledge, the barrier method has not previously been used for analysis of systems with malicious adversaries present.

III. System and Adversary Model

We consider a linear system with state $\mathbf{x}(t) \in \mathbb{R}^n$, input $\mathbf{u}(t) \in \mathbb{R}^m$, and observations $\mathbf{y}(t) \in \mathbb{R}^p$. The system dynamics are described by the equations

$$\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t) + \mathbf{w}(t) \qquad (1)$$

$$\mathbf{y}(t) = C\mathbf{x}(t) + \mathbf{v}(t) + \mathbf{a}(t) \qquad (2)$$

In (1), $\mathbf{w}(t)$ is a Gaussian process with mean identically zero and autocorrelation function $R_w(\tau) = \Sigma_w\delta(\tau)$. In (2), $\mathbf{v}(t)$ is a Gaussian process with mean identically zero and autocorrelation function $R_v(\tau) = \Sigma_v\delta(\tau)$. The processes $\mathbf{v}(t)$ and $\mathbf{w}(t)$ are independent. The initial state $\mathbf{x}(0)$ is equal to $\mathbf{x}_0$. The control strategy of the system is defined as a function μ that takes as input

$$\{\mathbf{u}(t') : t' < t\} \cup \{\mathbf{y}(t') : t' < t\},$$

and outputs a control signal $\mathbf{u}(t)$. The vector $\mathbf{a}(t)$ describes the impact of the attack as follows.

We let S denote the set of sensors that could be compromised by the adversary. If the sensors are compromised, then the vector $\mathbf{a}(t)$ satisfies support$(\mathbf{a}(t)) \subseteq S$ for all t. The nonzero entries of $\mathbf{a}(t)$ can be chosen arbitrarily by the adversary. Otherwise, if no compromise has taken place, $\mathbf{a}(t) \equiv 0$. At each time t, the adversary is assumed to have knowledge of the control policy, the system states $\mathbf{x}(t')$ for $t' \leq t$, and the values of $\mathbf{u}(t')$ and $\mathbf{y}(t')$ for $t' \leq t$. We define variable $\alpha \in \{0,1\}$ to be 1 if the adversary has compromised the sensors S and 0 otherwise. The adversary's strategy τ is defined as a mapping from

$$\{\mathbf{x}(t'), \mathbf{u}(t'), \mathbf{y}(t') : t' \leq t\}$$

to a vector $\mathbf{a}(t)$. The assumption that the adversary has access to the true system state enables us to model attacks in which the adversary has direct observation of the targeted system.

We let $G \subseteq \mathbb{R}^n$ denote a set of goal states, and let $U \subseteq \mathbb{R}^n$ denote a set of unsafe states. We let $T > 0$ denote the final time of the system. The goal of the system is to satisfy $\mathbf{x}(T) \in G$ and $\mathbf{x}(t) \notin U$ for all $t \in [0,T]$. We assume that U and R are defined by

$$U = \{\mathbf{x} \in \mathbb{R}^n : g_U(\mathbf{x}) \geq 0\}, \ G = \{\mathbf{x} \in \mathbb{R}^n : g_G(\mathbf{x}) \geq 0\}.$$

In order to satisfy this goal, a trajectory $\{\mathbf{r}(t) : t \in [0,T]\}$ is chosen that satisfies $\mathbf{r}(t) \notin U$ and $\mathbf{r}(T) \in G$. The system then attempts to track the reference trajectory by choosing a control strategy that minimizes

$$\mathbf{E}\left[\int_0^T ((\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) + \mathbf{u}(t)^T R\mathbf{u}(t)) \ dt\right]$$

where Q and R are given positive-definite matrices.

IV. Problem Formulation and Solution Approach

In this section, we first present the problem formulation and then our proposed solution approach.

A. Formulation

The problem studied in this work is

$$\begin{array}{ll} \min_\mu & \mathbf{E}\left[\int_0^T ((\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) \right. \\ & \left. +\mathbf{u}(t)^T R\mathbf{u}(t)) \ dt|\mu,\overline{\alpha}\right] \\ \text{s.t.} & \max_\tau \left\{Pr\left(\bigcup_{t\in[0,T]} \{\mathbf{x}(t) \in U\}|\mu,\tau,\alpha\right)\right\} \leq \epsilon \\ & \min_\tau \{Pr(\mathbf{x}(T) \in G|\mu,\tau,\alpha)\} \geq 1-\delta \end{array} \qquad (3)$$

Here, $\mathbf{E}(\cdot|\mu,\overline{\alpha})$ denotes the expectation conditioned on the control policy μ when no adversary is present, while $Pr(\cdot|\mu,\tau,\alpha)$ denotes the probability conditioned on control policy μ, adversary policy τ, and $\alpha = 1$. The objective function captures the system's goal of minimizing the expected tracking error when there is no attack. The first constraint implies that the worst-case probability (over all adversary policies τ) of violating the safety constraint is bounded above by $\epsilon > 0$. The second constraint implies that, for any adversary policy, the probability of satisfying the reachability constraint is at least $(1-\delta)$. Hence, when no adversary is present ($\alpha = 0$), the control policy should provide minimal tracking error. When there is an adversary present ($\alpha = 1$), the tracking requirement is relaxed, but the system must still satisfy the safety and reachability constraints with a desired probability regardless of the adversary's strategy. The parameters δ and ϵ are chosen based on the safety requirements of the system, similar to other works on safety verification of stochastic systems [19]. The time horizon T is similarly chosen based on the alloted time to reach the desired state, although infinite-horizon and free-endpoint generalizations of the problem can also be considered.

B. Solution Approach

Our proposed solution approach is based on the following intuition. By injecting false data to modify the observations, the adversary aims to cause the system to choose a suboptimal control action that violates the safety and reachability constraints. Hence, the impact of the attack is determined by how far the control deviates from its desired value. By proactively limiting the set of feasible control inputs at each time step, the system can therefore bound the damage caused by the attack and derive approaches that are guaranteed to satisfy the desired safety and reachability properties.

In order to limit the set of control inputs, we consider the state estimate obtained from the set $\{1,\ldots,p\} \setminus S$ of secure sensor measurements. We then require that the selected control action must lie within a particular neighborhood of the optimal control action based on the computed estimate.

Formally, we let C_{α} denote the matrix obtained by selecting the rows from C indexed in $\{1,\ldots,p\} \setminus S$, i.e., the matrix of observations that are not affected by the adversary. We let $\mathbf{y}_{\alpha}(t) \in \mathbb{R}^{p-|S|}$ denote the vector of observations that are not compromised in the attack and $\mathbf{v}_{\alpha}(t) = (v_i(t) : i \in \{1,\ldots,p\} \setminus S)$, so that $\mathbf{y}_{\alpha}(t) = C_{\alpha}\mathbf{x}(t) + \mathbf{v}_{\alpha}(t)$. Let $\Sigma_{v_{\alpha}}$ denote the covariance matrix of $\mathbf{v}_{\alpha}$.

We define $\hat{\mathbf{x}}_{\alpha}(t)$ to be the least-squares estimate of $\mathbf{x}(t)$ based on the measurements $\{\mathbf{y}_{\alpha}(t') : t' \leq t\}$. This estimate can be obtained as the output of a continuous-time Kalman filter defined by [15]

$$\begin{aligned}
\dot{\hat{\mathbf{x}}}_{\alpha}(t) &= A\hat{\mathbf{x}}_{\alpha}(t) + \Theta(t)(\mathbf{y}_{\alpha}(t) - C_{\alpha}\hat{\mathbf{x}}_{\alpha}(t)) + B\mathbf{u}(t) \\
\Theta(t) &= \Phi(t)C_{\alpha}^{T}\Sigma_{\mathbf{v}_{\alpha}}^{-1} \\
\dot{\Phi}(t) &= A\Phi(t) + \Phi(t)A^{T} + \Sigma_{\mathbf{w}} - \Phi(t)C_{\alpha}^{T}\Sigma_{\mathbf{v}_{\alpha}}^{-1}C_{\alpha}\Phi(t)^{T}
\end{aligned}$$

where $\Phi(0) = 0$ and $\hat{\mathbf{x}}_{\alpha}(0) = \mathbf{x}_0$. The optimal control based on the measurements $\mathbf{y}_{\alpha}(t)$, denoted $\mathbf{u}_{\alpha}(t)$, is then given by

$$\begin{aligned}
\mathbf{u}_{\alpha}(t) &= \frac{1}{2}K(t)\hat{\mathbf{x}}_{\alpha}(t) - \frac{1}{2}R^{-1}B^{T}\mathbf{s}(t) \\
K(t) &= -R^{-1}B^{T}P(t) \\
-\dot{P}(t) &= A^{T}P(t) + P(t)A - \frac{1}{2}P(t)BR^{-1}B^{T}P(t) + 2Q \\
\dot{\mathbf{s}}(t) &= (-A^{T} + \frac{1}{2}P(t)BR^{-1}B^{T})\mathbf{s}(t) + 2Q\mathbf{r}(t)
\end{aligned}$$

where $\mathbf{s}$ and P have boundary conditions $\mathbf{s}(T) = 0$ and $P(T) = 0$. We then define the set of feasible control inputs at time t as

$$\mathcal{U}_{\gamma}(t) \triangleq \{\mathbf{u}(t) : ||\mathbf{u}(t) - \mathbf{u}_{\alpha}(t)||_2 \leq \gamma\}$$

for some $\gamma \geq 0$. Under this constraint, the control problem becomes

$$\begin{array}{ll}
\text{minimize} & \mathbf{E}\left[\int_0^T ((\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) \right. \\
 & \left. +\mathbf{u}(t)^T R\mathbf{u}(t))\ dt\right] \\
\mathbf{u}(t) & \\
\text{s.t.} & ||\mathbf{u}(t) - \mathbf{u}_{\alpha}(t)||_2 \leq \gamma \quad \forall t \in [0,T]
\end{array} \tag{4}$$

The solution to (4) can be computed by solving a stochastic HJB equation

$$\begin{aligned}
0 = \min_{u \in \mathcal{U}_{\gamma}(t)} & \left\{ (\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) \right. \\
& + \mathbf{u}(t)^T R\mathbf{u}(t) + V_t(t,x) \\
& \left. + V_x(t,x)(A\mathbf{x}(t) + B\mathbf{u}(t)) + \frac{1}{2}\mathbf{tr}(V_{xx}(t,x)\Sigma_{\mathbf{w}}) \right\}.
\end{aligned}$$

Since solving a PDE of this form is computationally challenging, we adopt the relaxation of assuming that the value function V is equal to the value function of the unconstrained problem

$$\min_{\mathbf{u}(t)} \mathbf{E}\left\{ \int_0^T ((\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) + \mathbf{u}(t)^T R\mathbf{u}(t))\ dt \right\}.$$

This value function is given by [15]

$$V(x,t) = \frac{1}{2}\mathbf{x}(t)^T P(t)\mathbf{x}(t) + \beta(t) + \mathbf{s}(t)^T\mathbf{x}(t) + s_0(t),$$

where

$$\begin{aligned}
-\dot{\beta}(t) &= \frac{1}{2}\mathbf{tr}(S(t)\Sigma_{\mathbf{w}}) \\
\dot{s}_0(t) &= \frac{1}{4}\mathbf{s}(t)^T BR^{-1}B^T\mathbf{s}(t) - \mathbf{r}(t)^T Q\mathbf{r}(t)
\end{aligned}$$

We present a bound on the probability that this approximation is tight, which occurs when the optimal control action in the non-adversarial case lies in $\mathcal{U}_{\gamma}(t)$ for all t, as Theorem 1 in Section V.

Under this approximation, the optimal control at each time step is equal to the minimizer of

$$\begin{aligned}
&(\mathbf{x}(t) - \mathbf{r}(t))^T Q(\mathbf{x}(t) - \mathbf{r}(t)) \\
&\quad + \mathbf{u}(t)^T R\mathbf{u}(t) + \mathbf{x}^T P(t)(A\mathbf{x} + B\mathbf{u}) + \mathbf{x}^T\dot{s}(t) + \\
&\quad\quad \frac{1}{2}\mathbf{x}^T\dot{P}(t)\mathbf{x}(t) + \dot{s}_0(t) + \mathbf{s}(t)^T(Ax + Bu)
\end{aligned} \tag{5}$$

over $\mathcal{U}_{\gamma}(t)$. Computing this minimizer is equivalent to solving the quadratically constrained quadratic program (QCQP)

$$\begin{array}{ll}
\text{minimize} & \mathbf{u}^T R\mathbf{u} + \hat{\mathbf{x}}(t)^T P(t)B\mathbf{u} + \mathbf{s}(t)^T B\mathbf{u} \\
\mathbf{u} & \\
\text{s.t.} & (\mathbf{u} - \mathbf{u}_{\alpha})^T(\mathbf{u} - \mathbf{u}_{\alpha}) \leq \gamma^2
\end{array} \tag{6}$$

at each time step, where $\hat{\mathbf{x}}(t)$ is equal to the expected value of $\mathbf{x}(t)$ conditioned on all previous measurements $\mathbf{y}$ under the assumption that no adversary is present. Eq. (6) is obtained by minimizing (5) (neglecting terms that do not contain $\mathbf{u}$) subject to the constraint that $\mathbf{u}(t) \in \mathcal{U}_\gamma(t)$. Problems of this type can be readily solved using standard convex optimization toolboxes. Our approach is shown as Figure 1.

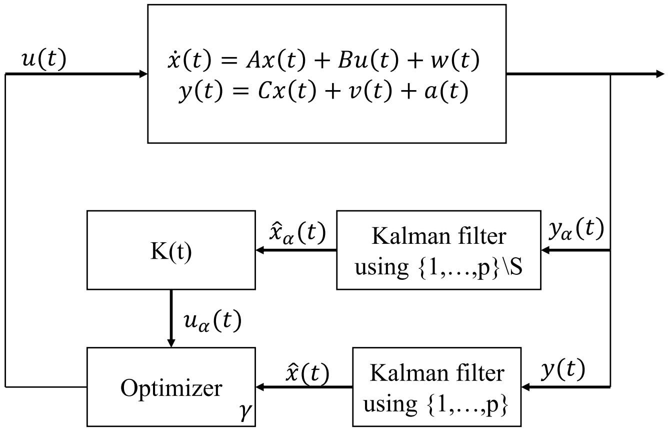
$u(t)$
$\dot{x}(t) = Ax(t) + Bu(t) + w(t)$
$y(t) = Cx(t) + v(t) + a(t)$
K(t)
$\hat{x}_\alpha(t)$
Kalman filter using {1,…,p}\S
$y_\alpha(t)$
$u_\alpha(t)$
Optimizer
γ
$\hat{x}(t)$
Kalman filter using {1,…,p}
$y(t)$

Fig. 1: Schematic illustration of our proposed approach. The optimizer solves the QCQP (6).

C. Selection of Parameter γ

The key design parameter in our approach is γ, which determines the range of admissible inputs at each time t. A high value of γ leads to a wider range of admissible inputs and hence better controller performance when there is no attacker present, however, a high γ value may also mean that the adversary is able to bias the system towards reaching an unsafe state.

Our approach is to perform a search over the possible values of γ in order to find the maximum γ such that the constraints of Eq. (3) are satisfied. In particular, we verify that, for any set of inputs $\mathbf{u}(t)$ satisfying $\mathbf{u}(t) \in \mathcal{U}_\gamma(t)$ for all $t \in [0, T]$, the safety constraints are satisfied with probability $(1-\epsilon)$ and the reachability constraints are satisfied with probability $(1-\delta)$.

In order to verify the safety and reachability constraints, we use the barrier function method. The idea of the barrier function method is to construct a function $D(\mathbf{x})$ such that, for some L, $D(\mathbf{x}_0) \leq L$, $D(\mathbf{x}) > K$ for all $\mathbf{x} \in U$, and D is decreasing along any feasible trajectory of $\mathbf{x}(t)$. Hence, by construction $\mathbf{x}(t)$ does not enter the unsafe region U. Letting $\hat{\mathbf{u}}(t) = \mathbf{u}(t) - \mathbf{u}_\alpha(t)$, we have

$$\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}_\alpha(t) + B\hat{\mathbf{u}}(t) + \mathbf{w}(t),$$

where $\hat{\mathbf{u}}(t)$ satisfies $||\hat{\mathbf{u}}(t)||_2 \leq \gamma$. In order to guarantee safety under any false data injection strategy of the adversary, we assume that $\hat{\mathbf{u}}(t)$ is a disturbance that can take arbitrary values. Using the definition of $\mathbf{u}_\alpha(t)$, we can write an extended system as

$$\begin{aligned}
\dot{\mathbf{x}}(t) &= A\mathbf{x}(t) + BK(t)\hat{\mathbf{x}}_\alpha(t) - BR^{-1}B^T\mathbf{s}(t) \\
&\quad + B\hat{\mathbf{u}}(t) + \mathbf{w}(t) \\
\dot{\hat{\mathbf{x}}}_\alpha(t) &= A\hat{\mathbf{x}}_\alpha(t) + \Theta(t)C_\alpha\mathbf{x}(t) + \Theta(t)\mathbf{v}_\alpha(t) \\
&\quad - \Theta(t)C_\alpha\hat{\mathbf{x}}_\alpha(t) - BR^{-1}B^T\mathbf{s}(t) + B\hat{\mathbf{u}}(t) \\
&\quad + BK\hat{\mathbf{x}}_\alpha(t) \\
-\dot{P}(t) &= A^TP(t) + P(t)A - P(t)BR^{-1}B^TP(t) + Q \\
\dot{\Phi}(t) &= A\Phi(t) + \Phi(t)A^T + \Sigma_{\mathbf{w}} - \Theta(t)\Sigma_{\mathbf{v}_\alpha}\Theta(t)^T
\end{aligned}$$

where $\Theta(t) = \Phi(t)C_\alpha^T\Sigma_{\mathbf{v}_\alpha}^{-1}$ and $K(t) = -R^{-1}B^TP(t)$. In order to reduce the complexity of the computation, we assume as an approximation that the Kalman filter and LQR controller are in steady-state, so that the matrices Φ, K, P, and Θ are constant. The deviation between this approximation and the true value can be bounded using the exponential stability of the optimal LQR controller and Kalman filter. Introducing additional disturbance terms into the verification that capture these deviations is a potential direction for future work. Furthermore, we assume that $\mathbf{s}(t)$ can be approximated as a polynomial function of t.

Define $\overline{\mathbf{x}}(t) = (\mathbf{x}(t), \hat{\mathbf{x}}_\alpha(t))$ as an extended state vector, with

$$\dot{\overline{\mathbf{x}}}(t) = f(\mathbf{x}(t)) + \overline{B}\hat{\mathbf{u}}(t) + \overline{F}\mathbf{w}(t) + \overline{G}(t)\mathbf{v}(t),$$

where

$$f(\overline{\mathbf{x}}(t)) = \left(\begin{array}{cc} A & BK \\ \Phi C_\alpha^T \Sigma_{v_\alpha} C_\alpha & A - \Theta C_\alpha + BK \end{array}\right)\left(\begin{array}{c} \mathbf{x}(t) \\ \hat{\mathbf{x}}_\alpha(t) \end{array}\right)$$

and

$$\overline{B} = \left(\begin{array}{c} B \\ B \end{array}\right), \quad \overline{F} = \left(\begin{array}{c} I \\ 0 \end{array}\right), \quad \overline{G} = \left(\begin{array}{c} 0 \\ \Theta \end{array}\right).$$

Letting N_w and N_v be matrices that satisfy $N_wN_w^T = \Sigma_w$ and $N_vN_v^T = \Sigma_v^\alpha$, define

$$\Lambda = (\overline{F}N_w \quad \overline{G}N_v).$$

The following proposition describes the safety guarantees provided by the barrier method for this model.

Proposition 1: Suppose there exists a function D such that

$$D(\overline{\mathbf{x}}_0) \leq \epsilon \tag{7}$$

$$D(\overline{\mathbf{x}}) \geq 1 \ \forall \overline{\mathbf{x}} \in U \tag{8}$$

$$D(\overline{\mathbf{x}}) \geq 0 \ \forall \overline{\mathbf{x}} \tag{9}$$

$$\begin{aligned}\frac{\partial D}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}}) + \overline{B}\hat{\mathbf{u}}) + \frac{\partial D}{\partial t} + \frac{1}{2}\mathbf{tr}(\Lambda^T\frac{\partial^2 D}{\partial \overline{\mathbf{x}}^2}\Lambda) \\ \leq 0 \ \forall \overline{\mathbf{x}}, ||\hat{\mathbf{u}}||_2 \leq \gamma\end{aligned} \tag{10}$$

Then $Pr\left(\bigcup_{t \in [0,T]} \{\mathbf{x}(t) \in U\}\right) \leq \epsilon$.
The proof combines Proposition 2 and Theorem 15 of [19], and is included in the appendix for completeness.

We follow a standard procedure for computing barrier functions, namely, we look for functions of the form

$$D(\overline{\mathbf{x}}) = \sum_{j=1}^{L} c_j b_j(\overline{\mathbf{x}}),$$

where the b_j's are polynomials and c_j's are real coefficients. The problem is mapped to a sum-of-squares (SOS) optimization

$$\begin{array}{ll} \text{minimize} & \sum_{j=1}^{L} k_j c_j \\ \text{s.t.} & d_{i,0}(x) + \sum_{j=1}^{L} d_{i,j}(x) c_j \text{ is SOS },\ i = 1, \ldots, p \end{array}$$

where "is SOS" implies that the polynomial can be decomposed as a sum of squares of polynomials. Such SOS problems can be solved via semidefinite programming. Define

$$g_D^{\gamma}(\hat{\mathbf{u}}) = \gamma^2 - \sum_{i=1}^{m} \hat{u}_i^2,$$

so that $||\hat{u}||_2 \leq \gamma$ is equivalent to $g_D^{\gamma}(\hat{\mathbf{u}}) \geq 0$. The following proposition describes the problem mapping.

Proposition 2: Suppose that there exist polynomials $\lambda_U(\overline{\mathbf{x}})$, $\lambda_D(\overline{\mathbf{x}}, \hat{\mathbf{u}})$, and $D(\overline{\mathbf{x}})$ such that the following hold:

$$\begin{aligned} -D(\overline{\mathbf{x}}_0) + \epsilon &\geq 0 & (11) \\ D(\overline{\mathbf{x}}) - 1 - \lambda_U^T(\overline{\mathbf{x}}) g_U(\overline{\mathbf{x}}) &\geq 0 & (12) \\ D(\overline{\mathbf{x}}) &\geq 0 & (13) \\ -\frac{\partial D}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}}) + \overline{B}\hat{\mathbf{u}}) - \lambda_D^T(\overline{\mathbf{x}}, \hat{\mathbf{u}}) g_D^{\gamma}(\hat{\mathbf{u}}) & & \\ -\frac{\partial D}{\partial t} - \frac{1}{2}\mathbf{tr}(\Lambda^T \frac{\partial^2 D}{\partial \overline{\mathbf{x}}^2}\Lambda) &\geq 0 & (14) \\ \lambda_U(\overline{\mathbf{x}}) \geq 0, \lambda_D(\overline{\mathbf{x}}, \hat{\mathbf{u}}) &\geq 0 & (15) \end{aligned}$$

Then $Pr\left(\bigcup_{t \in [0,T]} \{\mathbf{x}(t) \in U\}\right) \leq \epsilon$

Proof: Eq. (11) and (13) clearly imply Eqs. (7) and (9). If (12) and (15) hold, then

$$D(\overline{\mathbf{x}}) - 1 \geq \lambda_U^T(\overline{\mathbf{x}}) g_U(\overline{\mathbf{x}}) \geq 0$$

for all $\overline{\mathbf{x}} \in U$, and hence (8) holds. Finally, if (14) and (15) hold, then

$$\begin{aligned} &\frac{\partial D}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}}) + \overline{B}\hat{\mathbf{u}}) + \frac{\partial D}{\partial t} + \frac{1}{2}\mathbf{tr}\left(\Lambda^T \frac{\partial^2 D}{\partial \overline{\mathbf{x}}^2}\Lambda\right) \\ &\qquad \leq -\lambda_D^T(\mathbf{x}, \hat{\mathbf{u}}) g_D^{\gamma}(\hat{\mathbf{u}}) \leq 0 \end{aligned}$$

when $||\hat{\mathbf{u}}||_2 \leq \gamma$, implying (10). Since the conditions of Proposition 1 hold, the probability of violating the safety constraint is bounded above by ϵ. ■

Proposition 2 implies that, for a given γ, the safety criterion can be checked by solving a sum-of-squares optimization problem, since each of the constraints (11)–(15) is a polynomial SOS constraint. Algorithm 1 gives a procedure for selecting γ. The procedure assumes existence of a function SOS_Feasible that takes a set of SOS constraints as input and returns a 1 if there exist polynomials satisfying the constraints and 0 otherwise.

In the case where the safety constraint is not satisfied with sufficient proability even when $\gamma = 0$, i.e., even when the optimal control law is followed based on the secure sensors alone, then the algorithm does not return a control policy.

The above approach can also be modified to ensure that the reachability constraint is satisfied with the desired probability $(1 - \delta)$. In this case, we incorporate time as an additional "state variable", i.e., $\dot{t} = 1$ and $t(0) = 0$. The extended state space is then given by $\mathbb{R}^n \times [0, \infty)$, while the unsafe region consists of $(V \setminus R) \times \{T\}$. The barrier function equations are given by

Algorithm 1 Algorithm for computing the maximum parameter γ that ensures safety.

```
1:  procedure SAFETY_SOS
2:      $\underline{\gamma} \leftarrow 0$, $\overline{\gamma} \leftarrow \gamma_{max}$
3:      while $|\underline{\gamma} - \overline{\gamma}| > \rho$ do
4:          $\gamma \leftarrow (\underline{\gamma} + \overline{\gamma})/2$
5:          $q \leftarrow$ SOS_Feasible(Eq. (11), Eq. (12), Eq. (13), Eq. (14), Eq. (15))
6:          if $q == 0$ then
7:              $\overline{\gamma} \leftarrow \gamma$
8:          else
9:              $\underline{\gamma} \leftarrow \gamma$
10:         end if
11:     end while
12:     return $\underline{\gamma}$
13: end procedure
```

$$\begin{aligned} D_1(\overline{\mathbf{x}}_0, 0) &\leq \delta & (16) \\ D_1(\overline{x}, T) &\geq 1\ \forall \overline{\mathbf{x}} \in V \setminus R & (17) \\ D_1(\overline{\mathbf{x}}, t) &\geq 0 & (18) \\ \frac{\partial D_1}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}}) + \overline{B}\hat{\mathbf{u}}) + \frac{\partial D_1}{\partial t} & & (19) \\ + \frac{1}{2}\mathbf{tr}(\Lambda^T \frac{\partial^2 D_1}{\partial \overline{\mathbf{x}}^2}\Lambda) &\leq 0\ \forall \overline{\mathbf{x}}, ||\hat{\mathbf{u}}||_2 \leq \gamma & (20) \end{aligned}$$

by the same argument as Proposition 1. The corresponding SOS constraints are then defined by

$$\begin{aligned} -D_1(\overline{\mathbf{x}}_0, 0) + \delta &\geq 0 & (21) \\ D_1(\overline{\mathbf{x}}, T) - 1 + \lambda_R^T(\mathbf{x}) g_R(\overline{\mathbf{x}}) &\geq 0 & (22) \\ D_1(\overline{\mathbf{x}}, t) &\geq 0 & (23) \\ \frac{\partial D_1}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}}) + \overline{B}\hat{\mathbf{u}}) + \frac{\partial D_1}{dt} & & \\ + \frac{1}{2}\mathbf{tr}\left(\Lambda^T \frac{\partial^2 D_1}{\partial \overline{\mathbf{x}}^2}\Lambda\right) - \lambda_D^T(\overline{\mathbf{x}}, \hat{\mathbf{u}}, t) g_D^{\gamma}(\hat{\mathbf{u}}) &\geq 0 & (24) \\ \lambda_R(\overline{\mathbf{x}}) \geq 0, \lambda_D(\mathbf{x}, \hat{\mathbf{u}}, t) &\geq 0 & (25) \end{aligned}$$

Proposition 3: Suppose that Eqs. (21)–(25) hold. Then $Pr(\mathbf{x}(T) \notin R) \leq \delta$.

The proof is analogous to that of Proposition 2. An algorithm for choosing the maximum parameter γ satisfying the reachability constraint can be obtained by starting with Algorithm 1 and replacing Line 5 with

$q \leftarrow$
SOS_Feasible(Eq. (21), Eq. (22), Eq. (23),
Eq. (24), Eq. (25)).

V. Analysis

In this section, we analyze the optimality of our proposed approach. We observe that if the optimal control action $\mathbf{u}^*(t)$ of (3) with the safety and reachability constraints removed is in $\mathcal{U}_{\gamma}(t)$ at each time t, then our proposed approach provides the same utility as the best possible control when no adversary is present, and hence is optimal. In what follows, we derive a bound on the probability that $\mathbf{u}^{*}(t) \in \mathcal{U}_{\gamma}(t)$ for all $t \in [0,T]$. As a preliminary, we define the concept of a martingale as follows.

Definition 1: A continuous random process (X_t) is a *martingale* if $\mathbf{E}(X_s|X_t) = X_t$ for all $s \geq t$. A *supermartingale* is a random process such that $\mathbf{E}(X_s|X_t) \leq X_t$ for all $s \geq t$. A *submartingale* is a random process such that $\mathbf{E}(X_s|X_t) \geq X_t$ for all $s \geq t$.

The probability that a submartingale crosses a particular bound is bounded as follows.

Lemma 1 (Doob's Martingale Inequality [20]): Let X_t be a nonnegative submartingale. Then for any $T > 0$ and constant θ,

$$Pr\left(\sup_{0 \leq t \leq T} X_t \geq \theta\right) \leq \frac{\mathbf{E}(X_T)}{\theta}.$$

We now state the following result bounding the probability of optimality for our approach. First, let $\overline{K} = \sup\{||K(t)||_2 : t \in [0,T]\}$.

Theorem 1: Suppose that (A,C) and (A,C_{α}) are both observable. Define

$$\lambda^{*} = \sup_{t}\{\lambda_{max}(\Sigma(t))\} \tag{26}$$

$$\lambda_{\alpha}^{*} = \sup_{t}\{\lambda_{max}(\Sigma_{\alpha}(t))\} \tag{27}$$

where $\Sigma(t)$ and $\Sigma_{\alpha}(t)$ are the covariance matrices of $(\mathbf{x}(t)-\hat{\mathbf{x}}(t))$ and $(\mathbf{x}(t)-\hat{\mathbf{x}}_{\alpha}(t))$, respectively, and $\lambda_{max}(\cdot)$ denotes the maximum eigenvalue of a matrix. Then we have

$$Pr(\mathbf{u}^{*}(t) \in \mathcal{U}_{\gamma}(t)\ \forall t \in [0,T]) \geq \frac{4(\lambda^{*}+\lambda_{\alpha}^{*})n\overline{K}^{2}}{\gamma^{2}} \tag{28}$$

Proof: We have that $\mathbf{u}^{*}(t) = \frac{1}{2}K(t)\hat{\mathbf{x}}(t) - \frac{1}{2}R^{-1}B^{T}\mathbf{s}(t)$ and $\mathbf{u}_{\alpha}(t) = \frac{1}{2}K(t)\hat{\mathbf{x}}_{\alpha}(t) - \frac{1}{2}R^{-1}B^{T}\mathbf{s}(t)$. Hence

$$Pr\left(\sup_{t\in[0,T]} ||\mathbf{u}^{*}(t)-\mathbf{u}_{\alpha}(t)||_2 \geq \gamma\right) \leq Pr\left(\sup_{t\in[0,T]} ||\hat{\mathbf{x}}(t)-\hat{\mathbf{x}}_{\alpha}(t)||_2 \geq \frac{\gamma}{\overline{K}}\right)$$

Therefore, it suffices to bound the maximum deviation between the estimate $\hat{\mathbf{x}}(t)$ computed using all of the sensors and the estimate $\hat{\mathbf{x}}_{\alpha}(t)$ computed using only the secure sensors. By the triangle inequality, we have

$$\begin{aligned} Pr\left(\sup_{t\in[0,T]} ||\hat{\mathbf{x}}(t)-\hat{\mathbf{x}}_{\alpha}(t)||_2 \geq \frac{\gamma}{\overline{K}}\right) &\leq Pr\left(\sup_{t\in[0,T]} ||\hat{\mathbf{x}}(t)-\mathbf{x}(t)||_2 \right. \\ &\quad \left. + \sup_{t\in[0,T]} ||\hat{\mathbf{x}}_{\alpha}(t)-\mathbf{x}(t)||_2 \geq \frac{\gamma}{\overline{K}}\right). \end{aligned}$$

In order for

$$\sup_{t\in[0,T]} ||\hat{\mathbf{x}}(t)-\mathbf{x}(t)||_2 + \sup_{t\in[0,T]} ||\hat{\mathbf{x}}_{\alpha}(t)-\mathbf{x}(t)||_2 \geq \frac{\gamma}{\overline{K}}$$

to hold, we must have $||\mathbf{x}(t)-\hat{\mathbf{x}}(t)||_2 \geq \frac{\gamma}{2\overline{K}}$ or $||\mathbf{x}(t)-\hat{\mathbf{x}}_{\alpha}(t)||_2 \geq \frac{\gamma}{2\overline{K}}$ for some $t \in [0,T]$. Taking a union bound, we obtain

$$\begin{aligned} Pr\Bigg(\sup_{t\in[0,T]} ||\hat{\mathbf{x}}(t)-\mathbf{x}(t)||_2 & \\ + \sup_{t\in[0,T]} ||\hat{\mathbf{x}}_{\alpha}(t)-\mathbf{x}(t)||_2 &\geq \frac{\gamma}{\overline{K}}\Bigg) \\ \leq Pr\left(\sup_{t\in[0,T]} ||\mathbf{x}(t)-\hat{\mathbf{x}}(t)||_2 \geq \frac{\gamma}{2\overline{K}}\right) & \\ + Pr\left(\sup_{t\in[0,T]} ||\mathbf{x}(t)-\hat{\mathbf{x}}_{\alpha}(t)||_2 \geq \frac{\gamma}{2\overline{K}}\right) & \end{aligned} \tag{29}$$

Considering each term of the right-hand side separately, we have

$$\begin{aligned} &Pr\left(\sup_{t\in[0,T]} ||\mathbf{x}(t)-\hat{\mathbf{x}}(t)||_2 \geq \frac{\gamma}{2\overline{K}}\right) \\ &= Pr\left(\sup_{t\in[0,T]} (\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}(\mathbf{x}(t)-\hat{\mathbf{x}}(t)) \geq \frac{\gamma^{2}}{4\overline{K}^{2}}\right) \\ &\leq Pr\left(\sup_{t\in[0,T]} (\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}\frac{1}{\lambda^{*}}(\mathbf{x}(t)-\hat{\mathbf{x}}(t)) \geq \frac{\gamma^{2}}{4\overline{K}^{2}\lambda^{*}}\right) \\ &\leq Pr\left(\sup_{t\in[0,T]} (\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}\Sigma(t)^{-1}(\mathbf{x}(t)-\hat{\mathbf{x}}(t)) \geq \frac{\gamma^{2}}{4\overline{K}^{2}\lambda^{*}}\right) \end{aligned}$$

For an observable system, the function

$$V(\mathbf{x}(t),t) = (\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}\Sigma(t)^{-1}(\mathbf{x}(t)-\hat{\mathbf{x}}(t))$$

is known to have a differential generator that is strictly decreasing [21], and hence is a submartingale. Doob's martingale inequality then implies that

$$\begin{aligned} &Pr\left(\sup_{t\in[0,T]} (\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}\Sigma(t)^{-1}(\mathbf{x}(t)-\hat{\mathbf{x}}(t)) \geq \frac{\gamma^{2}}{4\overline{K}^{2}\lambda^{*}}\right) \\ &\leq \left(\frac{\gamma^{2}}{4\overline{K}^{2}\lambda^{*}}\right)^{-1} \lim_{t\rightarrow\infty}\left[\mathbf{E}((\mathbf{x}(t)-\hat{\mathbf{x}}(t))^{T}\Sigma(t)^{-1}(\mathbf{x}(t)-\hat{\mathbf{x}}(t)))\right] \\ &= \frac{4\lambda^{*}n\overline{K}^{2}}{\gamma^{2}}. \end{aligned}$$

Following a similar procedure for the second term of (29) and combining inequalities yields the desired result. ■

The results of Proposition 2 and Theorem 1 imply that the control strategy chosen using our approach satisfies the safety and reachability properties in the presence of the adversary with a desired probability, while also achieving optimality in the absence of the adversary with probability characterized by Theorem 1.

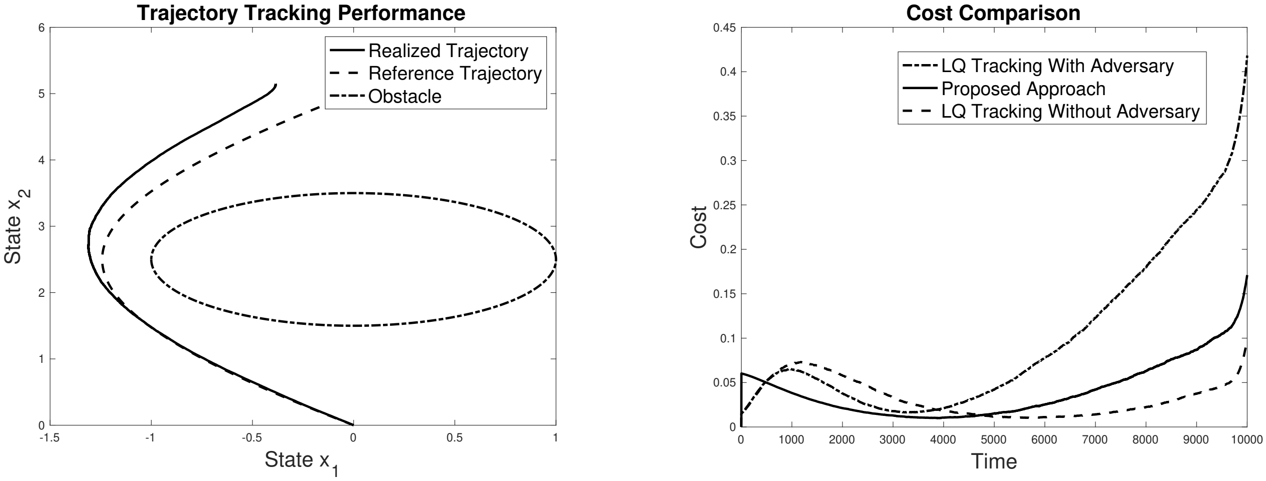

(a)

(b)

Fig. 2: Numerical evaluation of proposed approach. (a) Comparison of reference trajectory with realized trajectory using our proposed approach. The realized trajectory slightly deviates from the reference due to the impact of the attack in order to avoid reaching the obstacle. (b) Comparison of the cost function for LQ tracking controller, our proposed approach, and the performance of LQ tracking when no adversary is present. Our proposed approach experiences a slight increase in error compared to the non-adversarial case, but greatly reduces the cost when an adversary is present.

VI. Simulation Study

In this section, we evaluate the proposed approach using a numerical case study. We consider a system with $n = 2$ states and $m = 2$ inputs equipped with $p = 2$ sensors. The system matrices A and B are set to be identity matrices with proper dimensions. Matrix C is designed as

$$C = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

The cost matrices are defined as $Q = I$ and $R = I/1000$ with proper dimensions. The noises $\mathbf{w}$ and $\mathbf{v}$ are generated using zero-mean Gaussian distributions. The covariance matrices $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{v}}$ are set as identity matrices.

The system tracks a reference trajectory $\mathbf{r}(t)$ in the presence of an adversary who can corrupt the measurement of the second sensor by injecting arbitrary false data. We simulate an adversary who injects noise chosen uniformly from $[0, 1]$ at each time step. The reference trajectory is designed as a hyperbola, which is characterized as

$$x_1 = 3\cosh(t) - 3\sqrt{2}, \quad x_2 = 2.5\sinh(t) + 2.5.$$

Moreover, the system is required to avoid the unsafe states $\{\bar{\mathbf{x}} | g_U(\bar{\mathbf{x}}) \geq 0\}$, where

$$g_U(\bar{\mathbf{x}}) = 1 - x_1^2 - (x_2 - 2.5)^2$$

defines a circle with radius 1 whose center located at $(0, 2.5)$. The worst case probability of reaching unsafe state was $\epsilon = 0.05$. The value of γ was chosen as 9.9219 using Algorithm 1.

First, we present the reference tracking performance in Fig. 2(a). We observe that the trajectory generated using the proposed approach tracks the reference trajectory well at the beginning. The trajectory then slightly deviates from the reference trajectory due to the safety requirement and the impact of the noise injected by the adversary.

The cost incurred by our approach, measured by the quadratic objective function with matrices Q and R as defined above, is shown in Fig. 2(b). The figure compares three scenarios: optimal LQG controller when there is an adversary injecting a noise signal as described above, our proposed approach when there is an adversary injecting the same noise signal, and an optimal LQG controller when no adversary is present. We found that our proposed approach significantly reduces the cost compared to the LQG controller that is corrupted by false data, and only slightly increases the cost compared to an LQG controller with no adversary. Characterizing the gap between the performance of our approach and the optimal LQG control when no adversary is present is a direction of future work.

VII. Conclusions

This paper considered the problem of reference tracking in the presence of an adversary under safety and reachability constraints. We modeled a system in which a subset of sensors can be compromised by an adversary, while the remaining sensors are assumed to be secure. We proposed a control strategy in which an optimal LQG control input is computed based on the secure sensors, and the selected control input is constrained to within a given bound of this nominal input value. We proposed a barrier certificate approach for selecting the bound and proving both safety and reachability with a desired probability. This approach requires only a Kalman filter and real-time solution of quadratically-constrained quadratic programs. Furthermore, we analyzed the probability that our proposed policy is optimal when no adversary is present. Our approach was evaluated via simulation study.

REFERENCES

[1] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned aircraft capture and control via GPS spoofing," *Journal of Field Robotics*, vol. 31, no. 4, pp. 617–636, 2014.

[2] M. Conti, R. Di Pietro, L. V. Mancini, and A. Mei, "Emergent properties: detection of the node-capture attack in mobile wireless sensor networks," in *Proceedings of the first ACM conference on Wireless network security*. ACM, 2008, pp. 214–219.

[3] S. Karnouskos, "Stuxnet worm impact on industrial cyber-physical system security," in *IECON 2011-37th Annual Conference on IEEE Industrial Electronics Society*. IEEE, 2011, pp. 4490–4494.

[4] J. Petit and S. E. Shladover, "Potential cyberattacks on automated vehicles." *IEEE Trans. Intelligent Transportation Systems*, vol. 16, no. 2, pp. 546–556, 2015.

[5] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 1, p. 13, 2011.

[6] A. Teixeira, D. Pérez, H. Sandberg, and K. H. Johansson, "Attack models and scenarios for networked control systems," in *Proceedings of the 1st international conference on High Confidence Networked Systems*. ACM, 2012, pp. 55–64.

[7] Y. Mo, E. Garone, A. Casavola, and B. Sinopoli, "False data injection attacks against state estimation in wireless sensor networks," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 5967–5972.

[8] R. Zhang and P. Venkitasubramaniam, "A game theoretic approach to analyze false data injection and detection in LQG system," in *Communications and Network Security (CNS), 2017 IEEE Conference on*. IEEE, 2017, pp. 427–431.

[9] H. Fawzi, P. Tabuada, and S. Diggavi, "Secure estimation and control for cyber-physical systems under adversarial attacks," *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1454–1467, 2014.

[10] Y. Shoukry, M. Chong, M. Wakaiki, P. Nuzzo, A. Sangiovanni-Vincentelli, S. A. Seshia, J. P. Hespanha, and P. Tabuada, "SMT-based observer design for cyber-physical systems under sensor attacks," *ACM Transactions on Cyber-Physical Systems*, vol. 2, no. 1, p. 5, 2018.

[11] M. Pajic, I. Lee, and G. J. Pappas, "Attack-resilient state estimation for noisy dynamical systems," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 82–92, 2017.

[12] A. Clark and L. Niu, "Linear quadratic Gaussian control under false data injection attacks," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 5737–5743.

[13] "Iran's alleged drone hack: Tough, but possible," https://www.wired.com/2011/12/iran-drone-hack-gps/.

[14] "Hackers fool tesla s's autopilot to hide and spoof obstacles," https://www.wired.com/2016/08/hackers-fool-tesla-ss-autopilot-hide-spoof-obstacles/.

[15] B. D. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*. Courier Corporation, 2007.

[16] A. P. Aguiar, J. P. Hespanha, and P. V. Kokotović, "Performance limitations in reference tracking and path following for nonlinear systems," *Automatica*, vol. 44, no. 3, pp. 598–610, 2008.

[17] E. Gilbert and I. Kolmanovsky, "Nonlinear tracking control in the presence of state and control constraints: a generalized reference governor," *Automatica*, vol. 38, no. 12, pp. 2063–2073, 2002.

[18] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2004, pp. 477–492.

[19] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.

[20] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*. Springer Science & Business Media, 2012, vol. 113.

[21] K. Reif, S. Gunther, E. Yaz, and R. Unbehauen, "Stochastic stability of the continuous-time extended kalman filter," *IEE Proceedings-Control Theory and Applications*, vol. 147, no. 1, pp. 45–52, 2000.

APPENDIX

Proof: [Proof of Proposition 1] By (8),

$$Pr\left(\bigcup_{t\in[0,T]}\{\mathbf{x}(t)\in\mathcal{U}\}\right) \leq Pr\left(\bigcup_{t\in[0,T]}\{D(\mathbf{x}(t))\geq 1\}\right)$$
$$= Pr\left(\sup_{t\in[0,T]} D(\mathbf{x}(t))\geq 1\right).$$

Our approach will be to show that $D(\mathbf{x}(t))$ is a supermartingale and then use Doob's martingale inequality (Lemma 1). By Dynkin's formula, for $s \leq t$,

$$\mathbf{E}(D(\mathbf{x}(t))|\mathbf{x}(s)) = D(\mathbf{x}(s)) + \mathbf{E}\left[\int_s^t AB(D(\mathbf{x}(\tau)))\ d\tau|\mathbf{x}(s)\right],$$

where $AB(\cdot)$ is the differential generator. In this case,

$$AB(D(\mathbf{x}(t))) = \frac{\partial D}{\partial \overline{\mathbf{x}}}(f(\overline{\mathbf{x}})+\overline{B}\hat{\mathbf{u}})+\frac{\partial D}{\partial t}+\frac{1}{2}\mathbf{tr}\left(\Lambda^T\frac{\partial^2 D}{\partial \overline{\mathbf{x}}^2}\Lambda\right) \quad (30)$$

Since the right-hand side of (30) is bounded above by 0 due to the constraint (10), $\mathbf{E}(D(\mathbf{x}(t))|\mathbf{x}(s)) \leq \mathbf{E}(D(\mathbf{x}(s)))$ and hence $D(\mathbf{x}(t))$ is a supermartingale.

By Doob's martingale inequality, we then have

$$Pr\left(\sup_{t\in[0,T]} D(\mathbf{x}(t))\geq 1\right) \leq \frac{D(\mathbf{x}(0))}{1} \leq \epsilon,$$

where the latter inequality is due to (7), thus completing the proof. ∎