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

Motion Planning with Secrecy

Anastasios Tsiamis, Andreea B. Alexandru, George J. Pappas

Abstract— **In this paper, we introduce the problem of motion planning with secrecy guarantees. A robot is tracking a desired trajectory, which is transmitted on-line by a planner, e.g. a base station or a mobile station. The communication between the robot and the planner is organized in packets and takes place over a wireless channel, which is susceptible to eavesdropping attacks. Our goal is to design secure communication codes in order to encode the trajectory information and hide it from any eavesdroppers. Meanwhile, the robot should be able to recover the trajectory and the planner should be able to estimate the robot's motion. We introduce a novel coding scheme that creates secrets between the robot and the planner based on i) the randomness of the robot's motion and ii) the imperfection of the communication channel. We show that every time the planner receives the corresponding packet while the eavesdropper misses it, a new secret is created, which can be used as a key to encode the information about the motion intent. If the motion-planning is random enough, one occurrence of this event makes the eavesdropper lose track of the trajectory; even if the eavesdropper has unlimited computational power. We apply our framework to the problem of way-point tracking, where the robot should visit some target positions one after the other. We illustrate the theoretical results in simulations.**

I. Introduction

Autonomous and remotely controlled agents have been developed for deployment in dangerous or mundane applications [1], [2]: from emergency response vehicles in hazardous environments to item sorting machines in industrial settings, and smart thermostats in Internet-of-Things applications. In many cases, an agent has to interact wirelessly with a supervisor, in order to exchange critical and confidential information about its operation. However, the wireless medium is vulnerable to eavesdropping attacks [3] and attacks that include tampering with the packets [4], such as denial-of-service attacks [5] and data-integrity attacks [6]–[10]. Here, we focus on the problem of confidentiality from eavesdroppers, which has been studied under various settings.

Encryption methods [11] offer confidentiality guarantees without requiring any mathematical model of the physical components, i.e. the source or the channel. However, their effectiveness is based on the assumption that the adversaries are computationally bounded, and introduce the issue of key management [12], [13]. On the other hand, physical layer security approaches [14]–[17] employ information theoretic tools, introduced in Shannon's seminal work [18], and exploit the channel model to develop codes in the physical layer of the wireless medium. The provided secrecy guarantees are provable, and independent of the eavesdropper's computational capability. However, constructing such codes is challenging and requires knowledge of the eavesdropper's channel model. In the case of packet erasure channels, more practical codes can be designed [19].

In this work, we introduce the problem of motion planning with secrecy guarantees. A robot is tracking a trajectory online as directed by a motion planner, e.g. the cloud or a base-station. This trajectory carries sensitive data about the robot's mission and should be kept confidential. Our goal is to design coding schemes to encode and hide the trajectory information from any eavesdroppers, while the robot is able to decode it. One of the challenges in such a problem is that a predictable motion cannot be hidden well; an eavesdropper can infer much information about the trajectory even under the most secure protocol [20]. The trajectory itself is an extra design parameter that affects secrecy.

A. Contributions

We propose a general coding scheme that exploits the inherent randomness of the robot's motion and the random erasures of the packet drop channels to create secrets between the robot and the planner. Contrary to the aforementioned approaches, we assume that the eavesdropper has unlimited computational power and there is no knowledge about the eavesdropper's channel. Based on this coding scheme, we explore how much information is leaked to the eavesdropper depending on the trajectory parameters. Then, we apply our framework to the problem of way-point tracking, where the robot should visit some target positions one after the other.

The main contributions are the following:

- We design a two-way coding scheme that creates new shared secrets between the mobile agent and the cloud planner, exploiting the inherent randomness in the robot and channel. It requires no previously shared secrets.
- Under our coding scheme, we study conditions on the trajectory such that its secrecy is protected.
- In the case of way-point tracking, under the condition that the way-points are independent, we guarantee secrecy for the desired reference sequence by using the proposed coding scheme.

Finally, we provide numerical simulations to exhibit the behavior of the proposed architecture and coding scheme and discuss the limitations and performance of such coding schemes. The proofs can be found in the Appendix.

This work was supported in part by ONR N00014-17-1-2012, and by NSF CNS-1505799 grant and the Intel-NSF Partnership for Cyber-Physical Systems Security and Privacy.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104. Emails: {atsiamis,aandreea,pappasg}@seas.upenn.edu

B. Related work

Several works considering the secrecy of the current state transmitted over lossy, eavesdropped channels, have been addressed in [21]–[26]. There, it is assumed that the system evolves around a set-point, which is known to all entities. In this work, we focus on achieving the confidentiality of not only the current state, but of the whole state trajectory. Also, we do not want the eavesdropper to learn about any set-points; the reference trajectory should be kept confidential.

Recently, distortion-based security was considered in [27] to protect the confidentiality of the whole state history of a mobile agent as it communicates with a legitimate receiver. The problem there is different since the receiver only estimates the state of the robot and does not transmit anything back to the robot, and the communication is assumed to be lossless. Using a pre-shared key, the robot confuses an eavesdropper by sending either the true trajectory or its mirror image with respect to an affine sub-space. The distortion guarantees are near-optimal but with high probability the eavesdropper knows the trajectory exactly under the mirroring scheme.

II. Problem formulation

The considered architecture with the robot, the channel, the eavesdropper and the planner is shown in Figure 1. First, we present the model for every component. Then, we present the goal of the paper–see Problem 1.

A. Robot model

The robot is modeled as a single integrator:

$$x_k = x_{k-1} + u_k + w_k,$$

where $x_k \in \mathbb{R}^n$ is the position vector, u_k is the velocity control input and $w_k \in \mathbb{R}^n$ is the process noise, modeled as i.i.d. Gaussian with zero mean and covariance Q. Suppose that the robot is required to track a reference trajectory r_k. In this case, it applies a control input of the form:

$$u_k = K(x_{k-1} - r_{k-1}) + r_k - r_{k-1}.$$

Hence, if we define $e_k = x_k - r_k$, the state equation of the robot becomes:

$$e_k = (I_n + K)e_{k-1} + w_k, \tag{1}$$

where I_n is the $n \times n$ identity matrix and the control gain K is designed such that the matrix $I_n + K$ has eigenvalues inside the unit circle. We assume that the initial state, trajectory and error are zero and known to all involved entities: $x_0 = r_0 = r_1 = e_0 = 0$. At the beginning of round k, the robot first applies the control u_k, then observes e_k. Then, based on e_k, it calculates an encoded message z_k and transmits it to the planner. At the end of round k, it receives and decodes an encoded message y_k from the planner, which contains information about the next reference r_{k+1}.

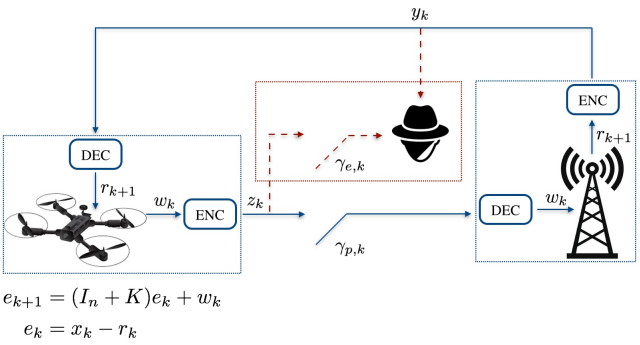

Fig. 1: The robot applies a control input u_k to track trajectory r_k and transmits a signal z_k containing information about e_k to the planner, over a packet-dropping channel. The planner computes the new reference r_{k+1} and sends its encoded version y_k back to the robot over a reliable channel. Meanwhile, an eavesdropper tries to intercept these messages.

B. Channel model

Communication between the robot and the planner follows the packet-based paradigm commonly used in networked control systems [28]–[30]. The robot transmits the messages z_k to the planner over a packet dropping channel with two outputs/receivers as shown in Figure 1. The first output $h_{p,k}$ is the authorized one to the planner, while the second $h_{e,k}$ is the unauthorized one to the eavesdropper. Communication with the planner is unreliable, i.e. packets might be dropped. Moreover, communication is not secure against the eavesdropper, i.e., the latter may intercept transmitted packets. We denote by $\gamma_{p,k} \in \{0, 1\}$ the outcome of the planner packet reception at time k, and by $\gamma_{e,k} \in \{0, 1\}$ the outcome of the eavesdropper's packet interception. If $\gamma_{p,k} = 1$ (or $\gamma_{e,k} = 1$), then the reception (interception, respectively) is successful. Otherwise, the respective packet is dropped. The outputs of the channel are modeled as:

$$h_{p,k} = \begin{cases} z_k, \text{ if } \gamma_{p,k} = 1 \\ \varepsilon, \text{ if } \gamma_{p,k} = 0 \end{cases}, \quad h_{e,k} = \begin{cases} z_k, \text{ if } \gamma_{e,k} = 1 \\ \varepsilon, \text{ if } \gamma_{e,k} = 0 \end{cases} \tag{2}$$

where symbol ε is used to represent the "no information" outcome. The channel outcomes $\{\gamma_{p,k}, \gamma_{e,k}\}_{k=0,1,\ldots}$ are random and assumed to be independent of the error e_k, and the reference r_k for $k = 0, 1, \ldots$. The joint distribution of the channel outcomes is assumed to be arbitrary.

In addition to the forward channel, the planner sends encoded trajectory information y_k back to the robot via the reverse channel. For clarity of exposition, we assume that the reverse channel is both reliable and totally unsecured, i.e. both the robot and the eavesdropper always receive y_k. The case of unreliable reverse channel is discussed in Section VI.

C. Trajectory planner

We assume that the planner updates the reference trajectory r_{k+1} in a linear way:

$$r_{k+1} = A\, r_k + B\, v_{k+1} \tag{3}$$

where $A, B \in \mathbb{R}^{n \times n}$ are matrices to be designed and $v_k \in \mathbb{V} := [-V_1, V_1] \times \ldots [-V_n, V_n]$ are the planner's control

signals, for some $V_i \in \mathbb{R}$, $i = 1, \ldots, n$. The bounds on V_i can represent the boundary of the workspace or constraints on how fast the robot's trajectory can change. The term $A\, r_k$ captures the dependency between the references across time and acts as a smoothing factor; it makes the motion less rough such that the robot can move without using large control effort. We assume that the inputs v_k, hence also r_k, are independent of the error e_k. For clarity of exposition, we postpone the details on how to incorporate dependency on e_k in Section VI. We assume that $\mathbb{V}$ is discretized in 2^N uniform regions so that the high-level input signals v_k can be encoded using N bits. The discretized space is denoted by $\mathbb{V}_d$. For simplicity, we ignore quantization errors in this paper and we assume that v_k are directly selected from $\mathbb{V}_d$. Finally, after updating the trajectory, the planner computes y_k, an encoded version of r_{k+1} and transmits it to the robot.

D. Eavesdropper

The eavesdropper is assumed to be powerful: it has unlimited computational power, knows all system parameters Q, K, A, B, all initial states $r_1 = r_0 = e_0 = 0$ and all the robot-planner channel outcomes $\gamma_{p,k}$. Furthermore, the eavesdropper is stealthy, i.e., the robot and the planner do not have any knowledge about the eavesdropper's channel and intercept successes $\gamma_{e,k}$.

E. Coding class

Here we define the class of the encoders which compute the encoded messages z_k and y_k. We will use the batch vector notation $a_{k_1:k_2} = \{a_{k_1}, \ldots, a_{k_2}\}$ to denote the collection of some vectors a_k for $k = k_1, \ldots, k_2$. We also use the notation

$$z_{0:k}(\gamma) = \{z_t, \text{ for all } \gamma_t = 1,\, t \leq k\},$$

where γ is a sequence in $\{0,1\}^{\mathbb{N}}$. We define the coding class of the robot encoder as follows:

$$\begin{aligned} z_k &= f_{enc,k}(e_k, z_{0:k-1}(\gamma_p), y_{0:k-1}) \\ e_k &= f_{dec,k}(z_k, z_{0:k-1}(\gamma_p), y_{0:k-1}), \end{aligned} \tag{4}$$

where $f_{enc,k}, f_{dec,k}$ are functions of appropriate dimensions. Similarly, we define the coding class of the planner encoder:

$$\begin{aligned} y_k &= g_{enc,k}(r_{k+1}, z_{0:k}(\gamma_p), y_{0:k-1}) \\ r_{k+1} &= g_{dec,k}(z_{0:k}(\gamma_p), y_{0:k}), \end{aligned} \tag{5}$$

where $g_{enc,k}, g_{dec,k}$ are functions of appropriate dimensions.

F. Problem

The robot needs to know r_{k+1} at the end of every round k in order to apply its control law u_{k+1} at round $k+1$. Thus, the robot should be able to decode the message y_k. The planner's goal is to infer the error e_k of the robot as well as transmit the trajectory information r_{k+1} in a secure way at every round k. Meanwhile, the eavesdropper tries to infer trajectory r_{k+1} based on the intercepted messages.

The inference of every entity depends on the information that each one of them possesses. Fix a channel outcome sequence $\{\gamma_{p,0:\infty}, \gamma_{e,0:\infty}\} =: \gamma$. Based on this event, the robot's information $\mathcal{I}^{\gamma}_{r,k}$ at the end of round k includes past and present values of the error $e_{0:k}$, and encoded messages $z_{0:k}$, $y_{0:k}$:

$$\mathcal{I}^{\gamma}_{r,k} = \{e_{0:k}, y_{0:k}, z_{0:k}\}. \tag{6}$$

The information of the eavesdropper $\mathcal{I}^{\gamma}_{e,k}$ at the end of round k consists of the intercepted messages, the planner's channel outcomes $\gamma_{p,0:k}$, and the encoded messages $y_{0:k}$:

$$\mathcal{I}^{\gamma}_{e,k} = \{z_{0:k}(\gamma_e), \gamma_{p,0:k}, y_{0:k}\}. \tag{7}$$

Finally, the planner's information $\mathcal{I}^{\gamma}_{p,k}$ at the end of round k includes the received messages, the trajectory $r_{0:k+1}$, and the encoded messages $y_{0:k}$:

$$\mathcal{I}^{\gamma}_{p,k} = \{z_{0:k}(\gamma_p), r_{0:k+1}, y_{0:k}\}. \tag{8}$$

Ideally, the security goal is trajectory secrecy, i.e. the eavesdropper should not learn anything about r_k at any time $k \geq 0$. However, this definition of secrecy is too restrictive. For example, if the eavesdropper intercepts information about just one non-zero reference, e.g. r_2, then, since

$$r_k = A^{k-2} r_2 + \sum_{i=3}^{k} A^{k-i} B v_i,$$

some information about the future trajectories will be leaked unless $A = 0$. Also, most of the information about the goal of the planner is included in the second term $\sum_{i=3}^{k} A^{k-i} B v_i$. For this reason, we seek secrecy of the inputs v_k, which can be viewed as secrecy of the motion's intent. Besides, if A is stable, then input secrecy leads to trajectory secrecy in the long-term, see Corollary 1.

Conditioned on an arbitrary channel outcome γ, we will use the conditional entropy $H(v_k|\mathcal{I}^{\gamma}_{e,k})$[1] to quantify how much information is leaked to the eavesdropper.

Definition 1 (Input Secrecy): Fix a channel outcome sequence $\{\gamma_{p,0:\infty}, \gamma_{e,0:\infty}\} =: \gamma$. Input secrecy at time k is achieved if the eavesdropper does not gain any information about about the planner input v_k at any other time k':

$$H(v_k|\mathcal{I}^{\gamma}_{e,k'}) = H(v_k), \text{ for all } k' \geq 0.$$

Before we formally describe the problem that this paper addresses, we summarize the assumptions, which hold throughout the paper.

Assumption 1: Matrices K, Q, A, B and $e_0 = r_0 = r_1 = 0$ in (1), (3) are known to all entities, i.e. to the robot, planner and eavesdropper. The noise process w_k, $k \geq 0$ is independent of the planner's input sequence v_k, $k \geq 0$. $\diamond$

Assumption 2: The robot and the planner do not have any prior shared secret beforehand. $\diamond$

Assumption 2 allows us to cover a large range of applications where an a priori secure communication and key were not established. Furthermore, this confers flexibility to our scheme, since the secrets can be refreshed during the execution time, as described in Section III, rather than having the secret fixed from the beginning for all the duration of the

[1]The collection of elements in the set $\mathcal{I}^{\gamma}_{e,k}$ changes for different γ. Thus, $H(v_k|\mathcal{I}^{\gamma}_{e,k})$ also changes with γ.

problem. However, since the robot and planner do not have a pre-shared secret, we cannot guarantee input secrecy from the beginning of the motion planning $k \geq 0$, regardless of the coding scheme used – see Theorem 1-ii). For this reason, we require it to occur eventually from some point on $k \geq k_0$.

Problem 1: Given system (1), trajectory (3), with channel model (2) and under Assumptions 1, 2, design coding schemes for z_k, y_k in class (4), (5) and find conditions on v_k such that for any channel sequence $\{\gamma_{p,0:\infty}, \gamma_{e,0:\infty}\} = \gamma$:

i) the robot always knows r_{k+1} at the end of step k:

$$H(r_{k+1}|\mathcal{I}^{\gamma}_{r,k}) = 0;$$

ii) the planner knows e_k whenever the transmission is successful at time k:

$$H(e_k|\mathcal{I}^{\gamma}_{p,k}) = 0, \text{ if } \gamma_{p,k} = 1;$$

iii) input secrecy is achieved for $k \geq k_0$, for some k_0:

$$H(v_k|\mathcal{I}^{\gamma}_{e,k'}) = H(v_k), \text{ for all } k' \geq 0,\ k \geq k_0.$$

III. Coding scheme

In this section, we propose an encoding scheme that can achieve the secrecy goals defined in Problem 1 under certain conditions. The main idea is that the robot and the planner can create secrets between them by exploiting the random packet erasures of the channel as well as the process noise of the dynamical system. Then, they use these secrets to encode and decode the planner inputs v_k.

We define the *reference time* t_k to be the time of the most recent successful reception at the planner before time k:

$$t_k = \max\{t:\ 0 \leq t < k,\ \gamma_{p,t} = 1\}. \tag{9}$$

Until the first successful transmission, we define $t_k = 0$. At every time step, the robot transmits the noise w_k along with the necessary noise information from t_{k+1} up to $k-1$, if $t_k < k-1$:

$$z_k = \begin{cases} w_k, & \text{if } t_k = k-1, \\ \{w_k, e_{k-1} - A^{k-1-t_k} e_{t_k}\}, & \text{if } t_k < k-1. \end{cases} \tag{10}$$

Using the above coding scheme, the planner can always calculate e_k with zero error, when the transmission is successful. The reason why the robot sends w_k separately when $t_k < k-1$ is explained in Remark 1.

Notice that under Assumption 1, the noises w_k are i.i.d. and assumed independent of the planner's inputs v_k. Hence, if the eavesdropper misses some of the packets, then those values can potentially be used as secret keys between the robot and the planner. Also, even if the eavesdropper intercepts information about e_k, this does not reveal anything about r_k, according to (1). However, we have to solve two problems: i) w_k are not uniformly distributed and ii) we do not exactly know which packets the eavesdropper has missed.

To solve the former problem, the planner applies a function $F : \mathbb{R}^n \to \{0, \ldots, 2^N - 1\}$, which is publicly known, such that $\bar{w}_k = F(w_k)$ is quantized and uniformly distributed on $\{0, \ldots, 2^N - 1\}$. For example, for $w \in \mathbb{R}^n$ Gaussian with zero mean and covariance Q, we can define:

$$\bar{w} = F(w) \triangleq \mathcal{Q}_N \left[\Phi \left(1/\sqrt{n} \sum_{i=1}^{n} \left[Q^{-1/2} w \right]_i \right) \right], \tag{11}$$

where Φ is the standard normal cumulative distribution function and $\mathcal{Q}_N : [0,1] \to \{0, \ldots, 2^N - 1\}$ is a uniform quantization function. Notice that the variable:

$$\tilde{w} := 1/\sqrt{n} \sum_{i=1}^{n} [Q^{-1/2} w]_i$$

has the standard normal distribution $\mathcal{N}(0,1)$. Thus, $\Phi(\tilde{w})$ is uniformly distributed on $[0,1]$ (see [31], exercise 1.2.4). As a result, applying $\mathcal{Q}_N$ on $\tilde{w}$ yields a $\bar{w}$ which is uniformly distributed on $\{0, \ldots, 2^N - 1\}$.

After computing $\bar{w}_k$, the planner calculates a value s_k at time k by XORing the binary representation of all the previously received quantized noises:

$$s_k = \bigoplus_{t \in P_k} F(w_t), \text{ with } P_k = \{t \leq k : \gamma_{p,k} = 1\}, \tag{12}$$

where $\oplus$ is the XOR operator and P_k is a set containing the times of successful reception up to time k. Until the first successful transmission (P_k is empty), we define $s_k = 0$. Finally, the planner uses s_k as a key to encode the trajectory information. In particular, it XORs the new input v_{k+1} with s_k. It also appends the information about the channel outcome $\gamma_{p,k}$ using an extra bit; this guarantees that the robot and the planner agree on the correct value of t_k.

$$y_k = \{v_{k+1} \oplus s_k, \gamma_{p,k}\}. \tag{13}$$

Upon receiving y_k, the robot uses the knowledge about $\gamma_{p,k}$ and the same quantizer F to compute s_k and recover v_{k+1}. Since matrices A, B are public, the robot has all the necessary information to compute the reference at time $k+1$, where r'_k is the previously decoded reference value and $y_k(1)$ is the first part of y_k (without the acknowledgement bit):

$$r'_{k+1} = A r'_k + B \Big[y_k(1) \oplus \underbrace{\Big(\bigoplus_{t \in P_k} \bar{w}_t \Big)}_{s_k} \Big],$$

As long as the eavesdropper keeps intercepting all the packets that the planner receives, it can compute s_k and recover v_{k+1}. However, due to the channel randomness, eventually the eavesdropper will miss some of those packets. Suppose this happens at time k_0, with $\gamma_{p,k_0} = 1$, $\gamma_{e,k_0} = 0$. Then, s_{k_0} will act as a secret between the planner and the robot, in fact, a one-time pad [11, Ch. 2] for the message v_{k+1}. The eavesdropper will not be able to compute the secret s_{k_0} in (12), since it has missed w_{k_0} permanently; from (10), the noise w_{k_0} does not appear again in any of the future messages z_k once the planner receives z_{k_0}. Under some independence conditions on the inputs v_k, the future messages y_k, z_k will not leak any information about the secret s_{k_0}, despite the reuse of the one-time pad, and the eavesdropper will be unable to decode v_{k+1}, for $k \geq k_0$.

This is why we call the event $\gamma_{p,k_0}=1$, $\gamma_{e,k_0}=0$ (when the user receives the packet while the eavesdropper misses it) a *critical event at time* k_0. Every time it occurs, a new secret between the planner and the robot is created.

Remark 1: When $t_k < k-1$, we send w_k along with $x_{k-1}-A^{k-1-t_k}x_{t_k}$. The planner has to know w_k separately for the secret creation procedure. Meanwhile, it also needs the additional information from time t_k+1 up to $k-1$ in order to estimate e_k accurately. This introduces an additional communication overhead when $t_k < k-1$. If the robot would send only w_k every time and omit the information from time t_k+1 up to $k-1$, this would impair the planner's estimation scheme. However, since the matrix I_n+K is stable, if the channel is good ($\gamma_{p,k}=1$ occurs often), then the estimation error of the planner would be nonzero but small if that information was omitted. ◇

The next theorem states that input secrecy can be achieved if the planner's inputs v_k are independent over time, under just one occurrence of the critical event.

Theorem 1 (Input secrecy): Consider system (1), trajectory (3) with channel model (2) and coding scheme (10) and (13). Let γ be a fixed channel outcome sequence.

i) The robot decodes the reference r_{k+1} at every time step k, while the planner accurately estimates e_k, when $\gamma_{p,k}=1$:

$$H(r_{k+1}|\mathcal{I}^{\gamma}_{r,k})=0 \tag{14}$$
$$H(e_k|\mathcal{I}^{\gamma}_{p,k})=0, \text{ for } \gamma_{p,k}=1 \tag{15}$$

ii) Let k_0 be the time that the first critical event occurs. No code within the classes (4), (5) can achieve input secrecy before k_0.

iii) Assume that the planner's inputs v_k are i.i.d., uniformly distributed on $\mathbb{V}_d\subset\mathbb{V}$, with $|\mathbb{V}_d|=2^N$. If the critical event occurs at some k_0 then input secrecy is achieved for all $k\geq k_0+1$:

$$H(v_k|\mathcal{I}^{\gamma}_{e,k'})=H(v_k), \text{ for all } k'\geq 0. \tag{16}$$

◇

Remark 2: The result of Theorem 1 implies that we can enhance security not only by designing secure communication protocols but also by making the trajectory more unpredictable, i.e. have i.i.d. inputs or at least design them to look like i.i.d. from the perspective of the eavesdropper.

Remark 3: The i.i.d. uniform requirement restricts the trajectories the robot can follow. If it is not satisfied, then the eavesdropper might be able to infer the secrets s_k. In [27], although a different coding scheme is used, the main result (Theorem 4.4) also requires the inputs to be independent (since they are independent of the trajectory). This is an inherent limitation of the problem if information theoretic secrecy is sought; the independence and uniform distribution of v_k are crucial for the secrecy of a scheme that reuses one-time pads. In fact, this can also be seen as the messages $v_1,\dots,v_k$ acting as one-time pads for the key s_k. ◇

In the following section, we show how this result can still be useful even if we have some mild dependence; the idea is to only encode the new random part of v_k. The main computational burden is computing the quantized noise at every time step $\bar{w}_k$. Nevertheless, for Gaussian noise, as in our model, this computation is cheap. If the inputs follow another known distribution, then we can make them uniform using a similar technique as in (11). However, transforming a multivariate distribution which is not already Gaussian or uniform can be computationally heavy.

The following corollary shows that in the long-term, input secrecy makes the eavesdropper lose track of the trajectory. Its minimum mean square error (mmse) estimate $\mathbb{E}\left\{r_k|\mathcal{I}^{\gamma}_{e,k}\right\}$ converges to the trivial estimate $\mathbb{E}\{r_k\}$, when the eavesdropper receives no information.

Corollary 1: Under the conditions of Theorem 1-iii), if matrix A is stable we obtain:

$$\left\|\mathbb{E}\left\{r_k|\mathcal{I}^{\gamma}_{e,k}\right\}-\mathbb{E}\{r_k\}\right\|\overset{a.s.}{\to}0$$

exponentially fast with rate $\rho(A)$, where $\rho(A)$ represents the spectral radius of A. Here a.s. denotes almost sure convergence with respect to v_k, w_k. ◇

If A is smaller the motion becomes more rough and unpredictable and the security guarantees are achieved with a faster rate. Moreover, if $A=0$, then trajectory secrecy is achieved immediately: $H\left(r_k|\mathcal{I}^{\gamma}_{e,k'}\right)=H(r_k)$, for all k,k' such that $k\geq k_0+1$, $k'>0$.

IV. Way-point tracking

In this section we show how to employ our coding scheme in the case of way-point tracking, defined as follows. Assume that there is a set of points $\left\{x_d^1,x_d^2\dots,x_d^M\right\}\in\mathbb{V}$, for some $M\geq 1$, that the robot should visit successively. For this reason, the planner uses a stable matrix A and $B=I_n-A$ as parameters and applies:

$$v_{k+1}=x_d^i, \text{ if } (i-1)T\leq k<iT. \tag{17}$$

In other words, we want the robot to visit a different point every T time steps. Notice that the inputs are piecewise dependent, so the previous result does not apply directly. However, we can reduce it to the previous case if the points $\left\{x_d^1,x_d^2\dots,x_d^M\right\}\in\mathbb{V}$ are independent. In order for the robot to be able to receive the relevant reference information, it is sufficient and necessary to encode and send only $v_{(i-1)T+1}=x_d^i$ for every i, instead of transmitting v_k every time:

$$\begin{aligned} y_k&=\{v_{k+1}\oplus s_k,\gamma_{p,k}\}, \quad \text{if } k=(i-1)T\\ y_k&=\gamma_{p,k}, \quad \text{otherwise.}\end{aligned} \tag{18}$$

This guarantees that the eavesdropper is not able to decode the key. Otherwise, if $v_{k+1}\oplus s_k$ is sent every time, security would be compromised. For example, if the critical event occurs at $(i-1)T<k_0<iT$ and the eavesdropper receives both $s_{k_0-1}\oplus x_d^i$, $s_{k_0}\oplus x_d^i$, this is equivalent to knowing $s_{k_0-1}\oplus s_{k_0}=F(w_{k_0})$. In other words, s_{k_0} would not be a secret despite the critical event occurring. The next theorem states that if the points are independent, then their secrecy is achieved after the critical event under the modified coding

scheme (18). The proof is omitted as it is similar to the one of Theorem 1.

Theorem 2 (Point stabilization): Consider the sequential point stabilization problem of system (1), with channel model (2), trajectory (3) with (17), and coding scheme (10), (18). Assume that the desired points x_d^i are i.i.d., uniformly distributed on $\mathbb{V}_d \subset \mathbb{V}$, with $|\mathbb{V}_d| = 2^N$. If the critical event occurs at some k_0, then input secrecy is achieved for $k \geq (i_0-1)T$, with $i_0 = \lceil k_0/T \rceil + 1$. Moreover,

$$H(x_d^i | \mathcal{I}_{e,k'}) = H(x_d^i), \quad \text{for all } i \geq i_0, k' \geq 0. \qquad \diamond$$

Notice that there is some trade-off between control performance and security. The planner's control v_k changes every T time steps. We just need the critical event to occur within T instead of one step to protect the next updated input. In this sense, the larger the T, the higher the chance that the sequence of points will be secure. For instance, if the channel outcomes are Bernoulli i.i.d. with $p_{1,0} = \mathbb{P}(\gamma_{p,k} = 1, \gamma_{e,k} = 0)$, then the probability that point x_d^i is protected is:

$$1 - (1 - p_{1,0})^{T(i-1)}.$$

A larger T makes the probability of leakage smaller. At the same time, a larger T also makes the motion slower.

Remark 4: It is possible to incorporate other types of dependency between the desired points as well, by decoding only the random part of the new information. For example, suppose that two consecutive points should be close: $\left\| x_d^i - x_d^{i-1} \right\|_\infty \leq R$, for some radius R. Then, if x_d^i is uniformly distributed on the box $B(x_d^{i-1}, R)$, instead of sending x_d^i, we can just encode $x_d^i - x_d^{i-1}$. It is subject of future work to generalize this idea to other types of dependencies, i.e. Markov chains.

V. Simulations

Let us consider an instance of the way-point stabilization problem with a sequence of 5 points to visit in a box with bounds $V_1 = V_2 = 10$, shown in Figure 2. The trajectory parameters used are $A = 0.98\, I_2$, $B = I_2 - A$. Let the channel outcomes be Bernoulli i.i.d., characterized by the following probabilities:

$p_{1,1} = 0.9$,	both receive the packet
$p_{0,0} = 0.05$,	both drop the packet
$p_{1,0} = 0.025$,	only the planner receives the packet
$p_{0,0} = 0.025$,	only the eavesdropper receives the packet.

As depicted in Figure 2a, a robot (black) moving according to (1) follows the reference dictated by the planner (blue) according to (3), (17). Due to the channel unreliability, critical events occur for the eavesdropper (the planner receives the packet from the robot, but the eavesdropper drops it), and their occurrence is depicted with red. The small probability of the user receiving the packet and the eavesdropper dropping it means that critical events are less likely to occur. However, as proved in Theorems 1 and 2, only one critical event is sufficient for the eavesdropper to lose track of the true reference.

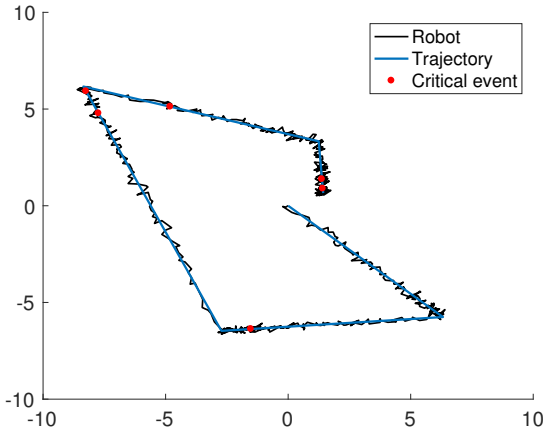

(a) Desired trajectory, actual trajectory and occurrence of critical events for the eavesdropper.

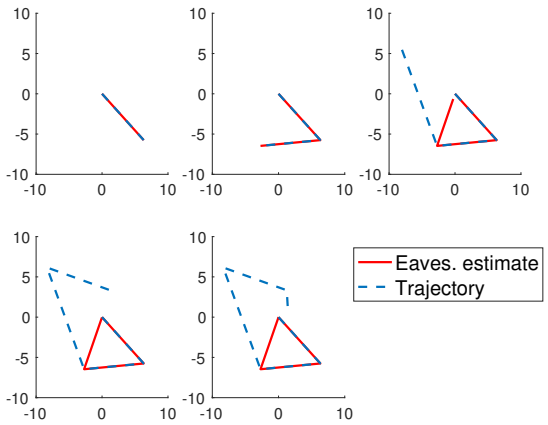

(b) Comparison between the true reference sequence and the eavesdropper's estimated reference sequence.

Fig. 2: Example of a sequence of reference points imposed by a planner for a robot to follow, under a coding scheme that ensures the confidentiality of the planner's sequence from a computationally unbounded eavesdropper.

We implemented the coding scheme from (10), (18) for the robot and the planner, with $N = 16$ bits, $T = 114$ time steps. We select the latter such that $\rho(A)^T < 0.1$; this guarantees that r_{iT} is close to x_d^i. In Figure 2b, the true reference sequence (blue), imposed by the planner, is compared to the eavesdropper's minimum mean square error estimate $r_{e,k} = \mathbb{E}\left\{ r_k | \mathcal{I}_{e,5T}^{\gamma} \right\}$ (red). Notice that the first critical event occurs in the second leg of the reference sequence in Figure 2a. The eavesdropper receives all the messages in the first leg, and can determine x_d^1. The eavesdropper also receives some messages in the second leg and it can infer the second reference point. However, after the critical event, the eavesdropper is not capable of estimating the rest of the trajectory and its estimation goes to zero.

The previous results were for a specific sample of channel outcome and a specific set of desired points. Next, we perform Monte Carlo simulation for 10000 samples of channel outcomes and sets of points. We simulate the eavesdropper's minimum mean square error (mmse) esti-

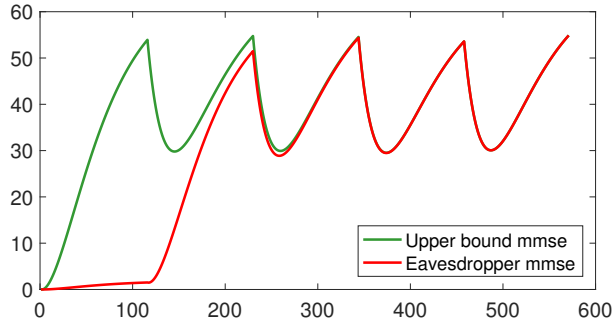

Fig. 3: The simulated mmse of the eavesdropper converges to the simulated upper-bound mmse, i.e. when the eavesdropper receives no information. For most channel outcomes, the critical event occurs within the first T time steps. Hence, the eavesdropper's simulated mmse starts increasing after $k = T$.

mate $\mathbb{E}\left\{\|r_k - r_{e,k}\|_2^2\right\}$ for every k. We compare it with $\mathbb{E}\left\{\|r_k\|_2^2\right\}$, which is the largest error for the eavesdropper, achieved when it has no information–see Figure 3. The comparison shows that the eavesdropper's simulated mmse converges to the largest possible value.

VI. Extensions and Future Work

In previous sections, we used some simplifying models for the clarity of the exposition. In this part, we provide details about how the simplifying assumptions can be dropped.

A. Unreliable acknowledgments

The planner should always agree with the robot about the messages they share, hence, they should agree on t_k. Otherwise, a critical event might occur for the planner too. This implies that we should deal with the case of unreliable acknowledgments. Suppose that the reverse channel is also a packet dropping one. Let $\gamma_{r,k} \in \{0, 1\}$ denote the reverse channel outcome at time k. If $\gamma_{r,k} = 1$, then the robot successfully receives the respective packet, otherwise it does not. Then, we can redefine the reference time to be:

$$\bar{t}_k = \max\{t : \ 0 \leq t < k, \, \gamma_{p,t}\gamma_{r,t} = 1\},$$

where we require both transmissions to be successful in order to update the reference time. To make sure that the user knows $\bar{t}_k$, the robot should also transmit $\bar{t}_k$ at every time step. Assume also that the trajectory r_{k+1} is updated only if the previous y_{k-1} was received by the robot. Then, if we define $\bar{\gamma}_{p,k} := \gamma_{p,k}\gamma_{r,k}$, the results of this paper still hold if we replace $\gamma_{p,k}$ by $\bar{\gamma}_{p,k}$.

B. Dependence of reference on the state error

In some cases, the reference r_{k+1} can incorporate a feedback from the state error e_k. We can depict this dependency in a linear way and assume that:

$$r_{k+1} = A\, r_k + B\, v_{k+1} + C\, e_k, \tag{19}$$

where again v_k are assumed independent of w_k. In our scheme, the eavesdropper has knowledge about $C\, e_k$, since it can compute e_k before the critical event occurs, and after the critical event k_0, it can compute e_{k_0+1} up to w_{k_0} uncertainty. However, this quantity acts as a disturbance to the motion planning. The information that the eavesdropper is interested in is still included in v_k as stated in Problem 1 and the eavesdropper cannot successfully use $C\, e_k$ to retrieve r_{k+1}.

C. Future work

Our coding scheme offers a way to create secrets between the robot and the planner. Those secrets can be used to hide the trajectory as long as the planner's new control input can be encoded independently of the previous ones, as a result of the limitation of reusing one-time pads. It is also required that the model (3) of the trajectory is known to the robot. In future work, it would be interesting to consider the case where the model of the trajectory, e.g. A, B and the distribution of v_k, is unknown to both the robot and the eavesdropper.

Appendix

Lemma 1: Assume $X, Y_1, \ldots, Y_k$ are i.i.d. uniformly distributed on $\{0, \ldots, 2^N - 1\}$ for some N. Then, the same is true for $(X, X \oplus Y_1, \ldots, X \oplus Y_k)$ and $(Y_j, X \oplus Y_1, \ldots, X \oplus Y_k)$ for all $j = 1, \ldots, k$. ◇

Proof: The proof is standard but we include it for completeness. We show that the joint distribution $(X, X \oplus Y_1, \ldots, X \oplus Y_k)$ is uniform on $\{0, \ldots, 2^N - 1\}^{k+1}$. The proof for the other cases is identical. Let $j, i_1, \ldots, i_k$ be arbitrary values $\in \{0, \ldots, 2^N - 1\}$. Then

$$\begin{aligned}
&\mathbb{P}(X = j, X \oplus Y_1 = i_1, \ldots, X \oplus Y_k = i_k) = \\
&\mathbb{P}(X = j, Y_1 = i_1 \oplus j, \ldots, Y_k = i_k \oplus j) = \\
&\mathbb{P}(X = j)\mathbb{P}(Y_1 = i_1 \oplus j) \ldots \mathbb{P}(Y_k = i_k \oplus j) = 2^{-N(k+1)}.
\end{aligned}$$

■

The next elementary information theoretic inequality can be found in [32], p.23, 29.

Lemma 2: Assume X, Y, Z are discrete random variables. Then:

$$H(X|Z) - H(X|Y, Z) = I(X; Y|Z) \geq 0. \qquad \diamond$$

Proof of Theorem 1

Proof of i) It follows by construction of the code. Since the robot receives $\gamma_{p,k}$, it knows P_k at every time k, where P_k is defined in (12). The robot's information also includes $w_{0:k}$. Thus, the secrets $s_{0:k}$ and consequently also $v_{0:k+1}$ are functions of its information $\mathcal{I}_{r,k}^{\gamma}$. Finally, since $r_0 = 0$, the trajectory r_{k+1} is a function of $v_{0:k+1}$. Thus, r_{k+1} can be computed from $\mathcal{I}_{r,k}^{\gamma}$ and (14) is satisfied.

Equation (15) follows by induction. Assume that the planner knows e_k, $\gamma_{p,k} = 1$. Let k' be the time that the next reception occurs: $t_{k'} = k$, $\gamma_{p,k'} = 1$. If $k' = k + 1$, the planner knows e_k and can compute e_k from (1). If $k' > k+1$, the planner can first compute $e_{k'-1}$ by adding $A^{k'-1-k}e_k$. The induction base case is similar, since $e_0 = 0$ is known.

Proof of ii) If $k < k_0$, then $z_{0:k}(\gamma_p) \subseteq z_{0:k}(\gamma_e)$, i.e. the eavesdropper has received all the packets that the planner

has. Furthermore, the eavesdropper knows all messages $y_{0:k}$. Thus, from (5), we obtain that $H\left(v_{k+1}|\mathcal{I}^{\gamma}_{e,k}\right)=0$.

Proof of iii) We analyze the case $k'\geq k_0$. The other case follows from the independence properties of the processes $w_{0:k}$, $v_{0:k}$. Fix a $k\geq k_0+1$. Recall the definition of $\bar{w}_{k'}=F\left(w_{k'}\right)$ in (11). Since the critical event occurred at k_0, the eavesdropper does not know w_{k_0}: $w_{k_0},\bar{w}_{k_0}\notin\mathcal{I}^{\gamma}_{e,k'}$, for any k'. Let the augmented information of the eavesdropper be:

$$\bar{\mathcal{I}}^{\gamma}_{e,k'}=\mathcal{I}^{\gamma}_{e,k'}\cup\left\{w_1,\dots,w_{k_0-1},w_{k_0+1},\dots,w_{k'}\right\},$$

where we add the information about every other noise w_t. By two applications of Lemma 2, we obtain:

$$H(v_k|\bar{\mathcal{I}}^{\gamma}_{e,k'})\leq H(v_k|\mathcal{I}^{\gamma}_{e,k'})\leq H(v_k).$$

It is sufficient to show that $H(v_k|\bar{\mathcal{I}}^{\gamma}_{e,k'})=H(v_k)$. We have:

$$\begin{aligned}&H(v_k|\bar{\mathcal{I}}^{\gamma}_{e,k'})\overset{a)}{=}\\&H(v_k|w_{1:k_0-1},w_{k_0+1:k'},v_{1:k_0},\bar{w}_{k_0}\oplus v_{k_0+1},\dots\bar{w}_{k_0}\oplus v_{k'+1})\\&\overset{b)}{=}H(v_k|\bar{w}_{k_0}\oplus v_{k_0+1},\dots,\bar{w}_{k_0}\oplus v_{k'+1}),\end{aligned}$$

where a) follows from the fact that $P_{k'}$ is known to the eavesdropper and $(s_t\oplus r_{t+1},w_{1:k_0-1},w_{k_0+1:k'})$ is equivalent to $(w_{k_0}\oplus r_{t+1},w_{1:k_0-1},w_{k_0+1:k'})$, for $t\geq k_0$. Further, b) follows from eliminating the independent variables.

From Lemma 1, we obtain that the sequence $(v_k,\bar{w}_{k_0}\oplus v_{k_0+1},\dots,\bar{w}_{k_0}\oplus v_{k'+1})$ is i.i.d. uniformly distributed on $\left\{0,\dots,2^N-1\right\}$, which implies:

$$H(v_k|\bar{w}_{k_0}\oplus v_{k_0+1},\dots,\bar{w}_{k_0}\oplus v_{k'+1})=H(v_k).\qquad\blacksquare$$

Proof of Corollary 1

From Theorem 1-iii), we have that after k_0+1, it holds that $H(v_k|\mathcal{I}^{\gamma}_{e,k})=H(v_k)$. Hence, $\mathbb{E}\left\{v_{k_0+1:k}|\mathcal{I}^{\gamma}_{e,k}\right\}=\mathbb{E}\left\{v_{k_0+1:k}\right\}$, equivalent to:

$$\mathbb{E}\left\{r_k|\mathcal{I}^{\gamma}_{e,k}\right\}-\mathbb{E}\left\{r_k\right\}=A^{k-k_0}\left(\mathbb{E}\left\{r_{k_0}|\mathcal{I}^{\gamma}_{e,k}\right\}-\mathbb{E}\left\{r_{k_0}\right\}\right)$$

Since the inputs v_k are bounded ($\mathbb{V}$ is bounded), r_{k_0} will be almost surely bounded. Thus, $\left\|\mathbb{E}\left\{r_k|\mathcal{I}^{\gamma}_{e,k}\right\}-\mathbb{E}\left\{r_k\right\}\right\|\overset{a.s.}{\to}0$ exponentially fast with rate $\rho(A)$. $\blacksquare$

REFERENCES

[1] P. Maes, *Designing autonomous agents: Theory and practice from biology to engineering and back*. MIT press, 1990.

[2] K. Goldberg and R. Siegwart, *Beyond Webcams: an introduction to online robots*. MIT press, 2002.

[3] Y. Zou, J. Zhu, X. Wang, and L. Hanzo, "A survey on wireless security: Technical challenges, recent advances, and future trends," *Proceedings of the IEEE*, vol. 104, no. 9, pp. 1727–1765, Sept 2016.

[4] H. Sandberg, S. Amin, and Johansson, K.H. (Organizers), "Cyberphysical Security in Networked Control Systems [Special Issue]," *IEEE Control Systems*, vol. 35, no. 1, 2015.

[5] A. Gupta, C. Langbort, and T. Başar, "Optimal control in the presence of an intelligent jammer with limited actions," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 1096–1101.

[6] Y. Mo, J. P. Hespanha, and B. Sinopoli, "Resilient detection in the presence of integrity attacks," *IEEE Transactions on Signal Processing*, vol. 62, no. 1, pp. 31–43, 2014.

[7] 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.

[8] 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.

[9] F. Pasqualetti, F. Dorfler, and F. Bullo, "Control-theoretic methods for cyberphysical security: Geometric principles for optimal cross-layer resilient control systems," *IEEE Control Systems*, vol. 35, no. 1, pp. 110–127, 2015.

[10] V. Lesi, I. Jovanov, and M. Pajic, "Security-aware scheduling of embedded control tasks," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 16, no. 5s, p. 188, 2017.

[11] J. Katz and Y. Lindell, *Introduction to modern cryptography*. CRC press, 2014.

[12] A. Perrig, J. Stankovic, and D. Wagner, "Security in wireless sensor networks," *Communications of the ACM*, vol. 47, no. 6, pp. 53–57, 2004.

[13] R. C. Merkle, "Secure communications over insecure channels," *Communications of the ACM*, vol. 21, no. 4, pp. 294–299, 1978.

[14] P. A. Regalia, A. Khisti, Y. Liang, and Tomasin, S. (Eds.), "Secure Communications via Physical-Layer and Information-Theoretic Techniques [Special Issue]," *Proceedings of the IEEE*, vol. 103, no. 10, 2015.

[15] A. D. Wyner, "The wire-tap channel," *The Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, Oct 1975.

[16] Y. Liang, H. V. Poor, and S. Shamai, "Secure communication over fading channels," *IEEE Transactions on Information Theory*, vol. 54, no. 6, pp. 2470–2492, 2008.

[17] H. Li, L. Lai, and W. Zhang, "Communication requirement for reliable and secure state estimation and control in smart grid," *IEEE Transactions on Smart Grid*, vol. 2, no. 3, pp. 476–486, 2011.

[18] C. E. Shannon, "Communication theory of secrecy systems," *Bell system technical journal*, vol. 28, no. 4, pp. 656–715, 1949.

[19] I. Safaka, L. Czap, K. Argyraki, and C. Fragouli, "Creating secrets out of packet erasures," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 6, pp. 1177–1191, 2016.

[20] J. N. Tsitsiklis and K. Xu, "Delay-predictability trade-offs in reaching a secret goal," *Operations Research*, vol. 66, no. 2, pp. 587–596, 2018.

[21] M. Wiese, K. H. Johansson, T. J. Oechtering, P. Papadimitratos, H. Sandberg, and M. Skoglund, "Secure estimation for unstable systems," in *IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 5059–5064.

[22] A. Tsiamis, K. Gatsis, and G. J. Pappas, "State estimation codes for perfect secrecy," in *IEEE 56th Conference on Decision and Control (CDC)*, 2017.

[23] ——, "State-secrecy codes for stable systems," in *American Control Conference (ACC), 2018*.

[24] A. S. Leong, D. E. Quevedo, D. Dolz, and S. Dey, "On remote state estimation in the presence of an eavesdropper," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 7339–7344, 2017.

[25] A. S. Leong, D. E. Quevedo, and S. Dey, "State estimation over markovian packet dropping links in the presence of an eavesdropper," in *IEEE 56th Conference on Decision and Control (CDC)*, 2017.

[26] A. Tsiamis, K. Gatsis, and G. J. Pappas, "State estimation with secrecy against eavesdroppers," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 8385–8392, 2017.

[27] G. Agarwal, M. Karmoose, C. Fragouli, S. Diggavi, and P. Tabuada, "Distorting an adversary's view in cyber-physical systems," in *IEEE Control and Decision Conference (CDC)*, 2018.

[28] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1453–1464, 2004.

[29] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proceedings of the IEEE*, vol. 95, no. 1, p. 138, 2007.

[30] K. Gatsis, A. Ribeiro, and G. J. Pappas, "Optimal power management in wireless control systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1495–1510, 2014.

[31] R. Durrett, *Probability: theory and examples*. Cambridge university press, 2010.

[32] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.