Sequential Detection of Replay Attacks with a Parsimonious Watermarking Policy

Arunava Naha[1], André Teixeira[2], Anders Ahlén[1] and Subhrakanti Dey[3]

Abstract— In this paper, we have proposed a technique for Bayesian sequential detection of replay attacks on networked control systems with a constraint on the average number of watermarking (ANW) events used during normal system operations. Such a constraint limits the increase in the control cost due to watermarking. To determine the optimal sequence regarding the addition or otherwise of watermarking signals, first, we formulate an infinite horizon stochastic optimal control problem with a termination state. Then applying the value iteration approach, we find an optional policy that minimizes the average detection delay (ADD) for fixed upper bounds on the false alarm rate (FAR) and ANW. The optimal policy turns out to be a two thresholds policy on the posterior probability of attack. We derive approximate expressions of ADD and FAR as functions of the two derived thresholds and a few other parameters. A simulation study on a single-input single-output system illustrates that the proposed method improves the control cost considerably at the expense of small increases in ADD. We also perform simulation studies to validate the derived theoretical results.

I. INTRODUCTION

Safety and security issues associated with cyber-physical systems (CPS) must be addressed before the widespread adoption of CPS for safety-critical applications. CPS are vulnerable to adversarial attacks on both, the cyber-layer and the physical-layer [1]. Defence mechanisms employed to protect the cyber-layer against adversarial attacks may not be adequate to protect the system against attacks on the physical layer [2]. Attackers may take different strategies to launch adversarial attacks on CPS. In one approach, the attacker may jam the wireless channel of a networked control system (NCS), and prevent the controller from receiving the required observations. Such an attack is called denial of service (DoS) attack [3]. In another approach known as a deception attack, the attacker may feed wrong or fake information to the system to cause some damage [4]. One particular case of such attacks is the replay attack [5], [2], [6], where the attacker first records the true observation from the system, and then replaces the true observation with the recorded data at some later point in time. The statistical similarity of the recorded data with true observation makes the detection of the replay attack a challenging task. The attacker may inject the system with harmful exogenous inputs and remain stealthy during the replay attack. The use of off-the-shelf networking components, commodity software, etc. makes CPS vulnerable to such attacks [2]. One of the most studied incidents of replay attacks is the Stuxnet attack, which took place in a uranium enrichment plant in Iran [7]. Attackers exploited a particular vulnerability of a commodity software to alter the control inputs to increase the speed of the centrifuges, which in turn increased the pressure beyond the safety limits. Attackers also launched the replay attack during the alteration of the control inputs to remain stealthy. Attacks on CPS may cause a serious threat to the reliability and availability of such systems, which may also cause monetary loss and threat to human safety [8]. Therefore, the detection of such attacks as early as possible is of utmost importance.

The literature is quite rich with methods for replay attack detections. One widely adopted approach is to add watermarking either to the control inputs or to the observations before the transmission. The presence of the attack is checked by various statistical tests at the controller or receiver end. In [2], [5], [6], the watermarking was added to the control inputs and the χ^2 statistics is generated using the innovation signal from a Kalman state estimator to perform a threshold check for the replay attack detection. A similar watermarking scheme is followed in [9], [10], but test statistics are generated directly using observations. The addition of the watermarking increases the detectability of the attack, which in turn reduces the detection delay [2]. On the other hand, the addition of watermarking increases the control cost [2]. Since attacks are infrequent events, the watermarking scheme increases the total control cost significantly. In [11], a periodic watermarking scheme is designed by keeping a balance between the control cost and detection delay. In a different approach of adding watermarking to the observations, the added watermarking is filtered out before feeding the received observations to the controller. Therefore, such an approach does not increase the control cost. The added watermarking signal could be of different types, sinusoidal [12], time-varying sinusoidal [13], random noise [14], multiplicative type [15], etc. However, in the scenario, where the attacker hijacks the sensor node and feeds false or previously recorded observations before the addition of the watermarking, such detection mechanisms may fail.

In this paper, we have studied the problem of sequential detection of replay attacks with a constraint on the average number of watermarking (ANW) added before the attack.

During the normal system operation, the addition of watermarking increases the control cost, which is limited by the constraint on the ANW. We have explored the Bayesian approach of sequential change-point detection, which was first introduced by Shiryaev in 1963 [16], for the replay attack detection by assuming a prior distribution of the attack start point. The Shiryaev procedure is asymptotically optimum under certain conditions [17]. In [17], analytical approximate expressions of the detection delay and false alarm rate (FAR) are derived provided the prior distribution of the change point satisfies either of the following two conditions,

$$\lim_{k \to \infty} \frac{\log \mathrm{P}\{\Gamma \ge k+1\}}{k} = -c, \ c \ge 0, \tag{1}$$

where Γ is the change point. Similar to the several other literature on change-point detections [18], [19], we have also assumed the distribution of the change point Γ to be a geometric distribution with parameters ρ, which obeys the condition in (1). However, some other prior distribution satisfying (1) can also be used. We first formulate a stochastic optimal control problem with an infinite horizon cost function and a termination state. The optimal policy, that will minimize the average detection delay (ADD) with constraints on both, FAR and ANW, is derived by solving Bellman's equation using value iterations [20]. The optimal policy is found to be a two thresholds (Th^s and Th^d) policy on the posterior probability of attack, p_k. We derive approximate expressions of the ADD and FAR as functions of the thresholds and a few other parameters for the replay attack detection problem. We have performed simulations using a single-input single-output (SISO) NCS to illustrate the efficacy of the proposed method and the derived theories. Our work in this paper is inspired by the two other prior works, quickest intrusion detection using an optimal set of active sensors in [18], and quickest change detection using on-off observation control in [19].

The rest of the paper is organized as follows. Section II illustrates the system model and the replay attack mechanism assumed in this paper. Section III formulates and solves the stochastic optimal control problem. The approximate analytical expressions of ADD and FAR are derived in Section IV. Section V provides the simulation results and discusses them. Section VI concludes the paper.

II. SYSTEM MODEL AND REPLAY ATTACK STRATEGY

This section discusses the system model before and after the replay attack and the replay attack strategy adopted for this paper.

A. System model before the attack

Fig. 1: Block diagram of the system before the attack.

A block diagram of the NCS before the attack is shown in Fig. 1. We consider a linear time-invariant SISO system for illustrative purposes. The extension of the proposed method for the general multi-input multi-output (MIMO) system is being considered in an extended version of this paper.

The state and observation equations for the system are

$$x_k = Ax_{k-1} + Bu_{k-1} + w_{k-1}, \ \text{and} \tag{2}$$
$$y_k = Cx_k + v_k, \tag{3}$$

where $x_k \in \mathbb{R}$, $u_k \in \mathbb{R}$, and $y_k \in \mathbb{R}$ are the state variable, control input, and observation at the k-th instant in time, respectively. $w_k \in \mathbb{R}$ and $v_k \in \mathbb{R}$ are independent and identically distributed (iid) process and observation noise, respectively. Also, $w \sim \mathcal{N}(0, Q)$, where $Q \in \mathbb{R}$, and $w \sim \mathcal{N}(0, R)$, where $R \in \mathbb{R}$. $\mathcal{N}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2. It is assumed that w_k and v_k are uncorrelated with each other, and both are uncorrelated with the initial state x_0. $A \in \mathbb{R}$, $B \in \mathbb{R}$, and $C \in \mathbb{R}$. We estimate the state of the system using the Kalman estimator. The predicted state $\hat{x}_{k|k-1}$, and the filtered state $\hat{x}_{k|k}$ from the Kalman filter (KF) are given as,

$$\hat{x}_{k|k-1} \triangleq \mathrm{E}[x_k|\Psi_{k-1}] = A\hat{x}_{k-1|k-1} + Bu_{k-1}, \tag{4}$$
$$\hat{x}_{k|k} \triangleq \mathrm{E}[x_k|\Psi_k] = \hat{x}_{k|k-1} + K\gamma_k, \tag{5}$$

where Ψ_k is the set of all measurements up to time k, and $\mathrm{E}[\cdot]$ denotes the expected value. The innovation signal γ_k and the Kalman gain K take the following forms,

$$\gamma_k = y_k - C\hat{x}_{k|k-1}, \text{and} \tag{6}$$
$$K = CP\left(C^2P + R\right)^{-1}. \tag{7}$$

Here, $P = \mathrm{E}\left[(x_k - \hat{x}_{k|k-1})^2\right]$ is the steady state error covariance. We have assumed that the system is operational for a long time (say, $k = -\infty$) and the system is currently ($k \ge 0$) at steady-state. At steady-state, P becomes the solution to the following algebraic Riccati equation,

$$P = A^2P + Q - A^2C^2P^2\left(C^2P + R\right)^{-1}. \tag{8}$$

The control input is generated using a linear quadratic Gaussian (LQG) controller by minimizing the following infinite horizon average cost.

$$J_{lqg} = \lim_{T \to \infty} E\left[\frac{1}{2T+1}\sum_{k=-T}^{T}\left(Wx_k^2 + Uu_k^2\right)\right]. \tag{9}$$

Here, $W \in \mathbb{R} > 0$ and $U \in \mathbb{R} > 0$. The optimal input u_k^* becomes a fixed gain linear function of the estimated state as,

$$u_k^* = L\hat{x}_{k|k}, \text{ and} \tag{10}$$
$$L = -ABS/\left(B^2S + U\right). \tag{11}$$

Here, S is the solution to the following algebraic Riccati equation,

$$S = A^2S + W - A^2B^2S^2\left(B^2S + U\right)^{-1}. \tag{12}$$

B. System model during Replay attack

For the replay attack model considered in this paper, the attacker does not need to have any knowledge about the system or controller parameters. However, the attacker can hack the sensor nodes and can record and replace true observations, but the attacker does not know the instantaneous values of the adding watermarking signal. Therefore, to launch a replay attack, the attacker replaces the true observation y_k by the fake data $z_k = y_{k-k_0}$ at $k = \Gamma$, where k_0 represents the delay. As mentioned before, the attack start point Γ is assumed to have a geometric prior distribution with parameter ρ, and $0 < \rho < 1$. Therefore, the probability $\pi_k = P\{\Gamma = k\}$ takes the following form,

$$\pi_k = \rho (1 - \rho)^{k-1} \mathbb{1}_{\{k \geq 1\}}, \qquad (13)$$

where $\mathbb{1}$ is the indicator function. We also assume that the probability of the attack starting before the time instant $k = 1$ is zero. A block diagram of the system under the replay attack is shown in Fig. 2. A similar attack model is studied in many literature [2], [11], and also, there are reported incidents, such as the Stuxnet attack [7].


Fig. 2: Block diagram of the system during replay attack.

III. PROPOSED REPLAY ATTACK DETECTION MECHANISM

In this section, we first present the proposed replay attack detection scheme with parsimonious watermarking in Sub-section III-A and III-B. In the subsequent subsections, we discuss the key aspects of the proposed strategy.

A. Proposed Parsimonious Watermarking Policy

The formal definitions of ADD and FAR are given as follows [19],

$$ADD \triangleq E_1 [\tau - \Gamma | \tau \geq \Gamma], \qquad (14)$$

$$FAR \triangleq P_0 \{\tau < \Gamma\} \qquad (15)$$

Here, τ is the time instant when the attack is detected by some hypothesis testing. P_0 and P_1 denote the probability measures before and after the attack, respectively. $E_1[\cdot]$ denotes the expected value with respect to the probability measure P_1. As discussed in [2], the detectability increases, *i.e.*, ADD reduces if we add watermarking to the control input. However, the addition of watermarking increases the control cost. If an iid watermarking $e_k \sim \mathcal{N}\left(0, \sigma_e^2\right)$ is added to the optimal LQG control input u_k^* for all the time instants then the increase in the control cost during the normal system operation is given by [4],

$$\Delta LQG_A = \left(U + B^2 \left(W + L^2 U\right) \left[1 - (A + BL)^2\right]^{-1}\right) \sigma_e^2. \qquad (16)$$

Therefore, we want to detect replay attacks with minimum ADD for a fixed upper bound on FAR and at a lower LQG

cost by a parsimonious watermarking policy. We consider a parsimonious watermarking policy based on the derived posterior probability of attack $p_k \triangleq P\{\Gamma \leq k | \mathcal{I}_k\}$, where $\mathcal{I}_k$ is the set of all available information up to the k-th instant of time. We define the variable average number of watermarking (ANW) used before the attack start point as follows,

$$ANW \triangleq E_0 [N_e], \qquad (17)$$

where $E_0 [\cdot]$ denotes the expectation with respect to the probability measure P_0, and N_e is the number of times the watermarking is added before the attack start point. Now, our final objective is to find a detection policy that will minimize the ADD for fixed upper bounds on FAR and ANW. The upper bound on the ANW will in turn limit the increase in the control cost due to the addition of the watermarking during the normal system operation. After the attack start point, our primary objective is to detect the attack as soon as possible to minimize the damage to the CPS, and we are not concerned about the increase in the control cost due to the watermarking. We define the control variable s_k, see (18), to control the watermarking addition process.

$$s_{k-1} = \begin{cases} 0, & \text{no watermarking added to } u_k^* \\ 1, & \text{watermarking added to } u_k^* \end{cases} \qquad (18)$$

The control input under the proposed watermarking scheme will be as follows,

$$u_k = u_k^* + s_{k-1} e_k. \qquad (19)$$

B. Proposed Replay Attack Detection Scheme

Figure 3 illustrates the proposed scheme in a block diagram. The method is provided systematically in the following steps.


Fig. 3: Schematic diagram of the system with the proposed watermarking scheme.

Step 1: We evaluate the value of the posterior probability of attack, p_k. Sub-section III-D discusses the derivation of p_k in detail.

Step 2: We compare p_k with two thresholds Th^s and Th^d, $Th^d \geq Th^s$, as follows.

    if $p_k < Th^s$ **then**

        No watermarking is added and Hypothesis H_0 selected.

    else if $p_k \geq Th^s$ & $p_k < Th^d$ **then**

        Added watermarking in next instant and Hypothesis H_0 selected.

    else if $p_k \geq Th^d$ **then**

        No watermarking is added and Hypothesis H_1 selected.

    end if

4870

Here the hypothesis H_0 denotes normal system operation, whereas H_1 denotes that the system is under replay attack. Thresholds Th^s and Th^d are evaluated by minimizing ADD for the fixed upper bounds on FAR and ANW. Sub-section III-C formulates the optimization problem, and the solution method is discussed in Sub-section III-E.

C. Optimization Problem Formulation

To formulate the stochastic optimal control problem, we define the following state variable θ_k and another control variable d_k.

$$\theta_k = \begin{cases} 0 & \text{Normal system operation,} \\ 1 & \text{System under replay attack,} \\ T_e & \text{Replay attack detected, process terminated.} \end{cases}$$

(20)

$$d_k = \begin{cases} 0, & \text{Hypothesis } H_0 \text{ selected} \\ 1, & \text{Hypothesis } H_1 \text{ selected.} \end{cases}$$

(21)

Now, we state the optimization problem as,

$$\min_{u_d \in \mathcal{U}} \ ADD$$
$$\text{s.t. } FAR \leq FAR_{th}$$
$$ANW \leq ANW_{th},$$

(22)

where FAR_{th} and ANW_{th} are the user defined thresholds. $u_{d,k}$ denotes the control action at the k-th time instant. The relationship between $u_{d,k}$, and s_k and d_k, and the the set of permissible actions, *i.e.*, $\mathcal{U}$ are given in Table I. To illustrate it further, for example, if the control action $u_{d,k} = 2$ is selected, then the corresponding decision variables will be $s_k = 1$ and $d_k = 0$.

TABLE I: $\mathcal{U}$

$u_{d,k}$	d_k	s_k
1	0	0
2	0	1
3	1	0

The constrained optimization problem of (22) is transformed into an unconstrained one using the Lagrangian multipliers λ_e and λ_f, see (23).

$$J^* = \min_{u_d \in \mathcal{U}} \ ADD + \lambda_f FAR + \lambda_e ANW.$$

(23)

ADD, FAR and ANW can be represented using the defined control and state variables as follows,

$$ADD = \mathrm{E}\left[\sum_{k=1}^{\tau} \mathbb{1}_{\{\theta_k=1\}} \mathbb{1}_{\{d_k=0\}}\right],$$

(24)

$$FAR = \mathrm{E}\left[\sum_{k=1}^{\tau} \mathbb{1}_{\{\theta_k=0\}} \mathbb{1}_{\{d_k=1\}}\right], \text{ and}$$

(25)

$$ANW = \mathrm{E}\left[\sum_{k=1}^{\tau} \mathbb{1}_{\{\theta_k=0\}} \mathbb{1}_{\{d_k=0\}} \mathbb{1}_{\{s_k=1\}}\right].$$

(26)

The posterior probability of attack, p_k, can also be represented as $p_k = \mathrm{E}\left[\mathbb{1}_{\{\theta_k=1\}} \mid \mathcal{I}_k\right]$ applying the Radon–Nikodym theorem [21]. Here, $\mathcal{I}_k$ denotes the set of all the required information upto the k-th time instant. Now, using (24)-(26) and p_k, we can rewrite the cost function of

(23) as the following infinite horizon cost function with a termination state [19].

$$J^* = \min_{u_d \in \mathcal{U}} \sum_{k=1}^{\tau} g_k\left(p_k, u_{d,k}\right),$$

(27)

where $g_k\left(p_k, u_{d,k}\right)$ is the expected per stage cost given as,

$$g_k\left(p_k, u_{d,k}\right) = p_k \mathbb{1}_{\{d_k=0\}} + \lambda_f\left(1 - p_k\right) \mathbb{1}_{\{d_k=1\}}$$
$$+ \lambda_e\left(1 - p_k\right) \mathbb{1}_{\{s_k=1\}} \mathbb{1}_{\{d_k=0\}}.$$

(28)

As per the theory studied in [22], the optimal policy found by solving (27) using dynamic programming value iterations will also be the solution to the original constrained problem (22).

D. Derivation of the posterior probability of attack, p_k

The posterior probability of attack, p_k, is derived using the joint distribution of the innovation signal and the watermarking signal, before and after the replay attack as given in Lemma 1.

Lemma 1: The posterior probability of the attack, p_k, for the proposed watermarking scheme follows the following recursion,

$$p_{k+1} = \begin{cases} \frac{T_m \mathrm{L}(\gamma_{k+1})}{T_m \mathrm{L}(\gamma_{k+1}) + 1 - T_m} & \text{if } s_k = 0 \\ \frac{T_m \mathrm{L}(\gamma_{k+1}, e_k)}{T_m \mathrm{L}(\gamma_{k+1}, e_k) + 1 - T_m} & \text{if } s_k = 1 \end{cases},$$

(29)

where $T_m = p_k + \left(1 - p_k\right) \rho$. $\mathrm{L}\left(\gamma_{k+1}\right)$ and $\mathrm{L}\left(\gamma_{k+1}, e_k\right)$ are the likelihood ratios as given below,

$$\mathrm{L}\left(\gamma_{k+1}\right) = \frac{\widetilde{f}\left(\gamma_{k+1}\right)}{f\left(\gamma_{k+1}\right)}, \text{ and}$$

(30)

$$\mathrm{L}\left(\gamma_{k+1}, e_k\right) = \frac{\widetilde{f}\left(\gamma_{k+1}, e_k\right)}{f\left(\gamma_{k+1}, e_k\right)},$$

(31)

where $f\left(\cdot\right)$ and $\widetilde{f}\left(\cdot\right)$ denote the likelihoods before and after the attack, respectively.

Proof: The proof makes use of the recursion formula for Shiraev statistics R_k (32) and the relation between p_k and R_k (33) [21],

$$R_k = (1 + R_{k-1})\mathcal{L}_k,$$

(32)

$$p_k = \frac{R_k}{R_k + 1/\rho},$$

(33)

where $\mathcal{L}_k$ denotes the likelihood ratio using the k-th data sample. Recalling that the likelihood ratios (30) and (31) correspond to $s_k = 0$ and $s_k = 1$, respectively, and applying (30)-(33), we obtain the recursion (29). ∎

The distributions of the test data, before and after the replay attack, remain zero-mean Gaussian. Therefore, we have only derived the required variances for $f\left(\cdot\right)$ and $\widetilde{f}\left(\cdot\right)$ as follows. The innovation signals before and after the replay attack take the following forms,

$$\gamma_k = \begin{cases} CA\left(x_{k-1} - \hat{x}_{k-1|k-1}\right) + Cw_{k-1} + v_k, & k < \Gamma \\ z_k - C\left(A + BL\right)\hat{x}_{k-1|k-1} - CBs_{k-2}e_{k-1}, & k \geq \Gamma \end{cases}$$

(34)

Therefore, the innovation signal before the attack is uncorrelated to the watermarking signal, but after the attack they

become correlated. The distributions of γ_k and e_{k-1}, before and after the replay attack are provided in the following lemma.

Lemma 2: For the system model and attack strategy given in Section II, the distributions of the innovation signal γ_k, and the joint distributions of the innovation signal γ_k and the watermarking signal e_{k-1}, before and after the replay attack, take the following forms,

$$\gamma_k \sim \begin{cases} f(\gamma_{k+1}) = \mathcal{N}(0, \sigma_\gamma^2), & k < \Gamma \\ \widetilde{f}(\gamma_{k+1}) = \mathcal{N}(0, \sigma_{\widetilde\gamma}^2), & k \geq \Gamma \end{cases}, \text{ and} \quad (35)$$

$$\{\gamma_k, e_{k-1}\} \sim \begin{cases} f(\gamma_{k+1}, e_k) = \mathcal{N}(\mathbf{0}, \Sigma_{\gamma_e}), & k < \Gamma \\ \widetilde{f}(\gamma_{k+1}, e_k) = \mathcal{N}(\mathbf{0}, \Sigma_{\widetilde\gamma_e}), & k \geq \Gamma \end{cases}, \quad (36)$$

where

$$\Sigma_{\gamma_e} = \begin{bmatrix} \sigma_\gamma^2 & 0 \\ 0 & \sigma_e^2 \end{bmatrix}, \quad (37)$$

$$\Sigma_{\widetilde\gamma_e} = \begin{bmatrix} \sigma_\gamma^2 & -BC\sigma_e^2 \\ -BC\sigma_e^2 & \sigma_e^2 \end{bmatrix}, \quad (38)$$

$$\sigma_\gamma^2 = C^2 P + R, \quad (39)$$

$$\sigma_{\widetilde\gamma}^2 = \left(1 + \frac{C^2 K^2 (A+BL)^2}{1-\mathcal{A}^2}\right)\sigma_y^2 + \left(\frac{2K\mathcal{A}C^2(A+BL)^2}{1-\mathcal{A}^2} - 2C(A+BL)\right)E_{\hat{x}z}(-1)$$
$$+ \left(\frac{B^2(1-KC)^2\widetilde\mu_s}{1-\mathcal{A}^2} + C^2 B^2\right)\sigma_e^2, \text{ and} \quad (40)$$

$$\sigma_y^2 = C^2 P + R + \frac{C^2 K^2 (A+BL)^2 (C^2 P + R)}{1-(A+BL)^2}$$
$$+ \left(C^2 B^2 \mu_s + \frac{B^2 \mu_s}{1-(A+BL)^2}\right)\sigma_e^2. \quad (41)$$

Here, $\mathcal{A} = (1-CK)(A+BL)$, $\mu_s = \mathrm{E}_0[s_k]$, and $\widetilde\mu_s = \mathrm{E}_1[s_k]$. The expression of $E_{\hat{x}z}(-1)$ is provided in Lemma 3.

Proof: The proof of Lemma 2 is provided in Appendix I. ∎

Remark 1: Since there is no constraint on the usage of watermarking after the attack start point, we assume $\widetilde\mu_s$ to be 1. On the other hand, there is an upper bound on the ANW, *i.e.*, ANW_{th}. Therefore, the range of μ_s will be $0 \leq \mu_s \leq \rho ANW_{th}$. The expectation of the attack start point Γ with respect to the prior distribution is ρ^{-1}.

Lemma 3: For the system model and attack strategy given in Section II, the correlation between the estimated state $\hat{x}_{k-1|k-1}$ after the replay attack and the attack signal z_k will be as follows,

$$E_{\hat{x}z}(-1) = \sum_{i=0}^{\infty} \mathcal{A}^i K C_a A_a^{i+1} E_{x_a}(0) C_a^T, \quad (42)$$

where $E_{x_a}(0)$ is the solution to the following Lyapunov equation,

$$E_{x_a}(0) = A_a E_{x_a}(0) A_a^T + Q_a, \text{ and} \quad (43)$$

$$A_a = \begin{bmatrix} A+BLKC & BL(1-KC) & BLK \\ (A+BL)KC & (A+BL)(1-KC) & (A+BL)K \\ 0 & 0 & 0 \end{bmatrix} \quad (44)$$

$$C_a = \begin{bmatrix} C & 0 & 1 \end{bmatrix} \quad (45)$$

$$Q_a = \begin{bmatrix} B^2\sigma_e^2\mu_s + Q & B^2\sigma_e^2\mu_s & 0 \\ B^2\sigma_e^2\mu_s & B^2\sigma_e^2\mu_s & 0 \\ 0 & 0 & R \end{bmatrix} \quad (46)$$

Proof: The proof of Lemma 3 is provided in Appendix II. ∎

Remark 2: $\mathcal{A}_a$ and A_a both are assumed to be strictly stable. Therefore, the summation of (42) will converge to a finite number as $i \to \infty$. We estimate the value of $E_{\hat{x}z}(-1)$ by taking a large i for which the rest of the terms in the summation become negligible.

E. Deriving The Optimal Policy, u_d^*

We have solved the optimization problem of (27) applying value iteration [20] in the following steps.

Step-1: p_k, λ_f, λ_e, and μ_s are discretized, see Table II.

TABLE II: Discretization

Name	Range	Discrete levels
p_k	$0 \leq p_k \leq 1$	N_p
λ_f	$0 < \lambda_f \leq \lambda_{f,max}$	N_{λ_f}
λ_e	$0 < \lambda_e \leq \lambda_{e,max}$	N_{λ_e}
μ_s	$0 \leq \mu_s \leq \rho ANW_{th}$	N_{μ_s}

Step-2: Using Monte-Carlo (MC) simulations, the state transition matrices $\mathbf{P}_{ne}$ and $\mathbf{P}_e$ are estimated offline, where $[\mathbf{P}_{ne}]_{ij} = \mathrm{P}\{p_k = j | p_k = i, s_{k-1} = 0\}$, $[\mathbf{P}_e]_{ij} = \mathrm{P}\{p_k = j | p_k = i, s_{k-1} = 1\}$, and $i, j = 1, 2, \cdots, N_p$. Here $[\cdot]_{ij}$ denotes the i-th row and j-th column element in a matrix.

Step-3: For each grid point in the combined search space of λ_e, λ_f and μ_s, the optimal policy u_d^* is derived using the following value iteration.

$$\mathrm{T}^{k+1}\mathbf{J} = \min_{u_d \in \mathcal{U}} \Big[\mathbf{g}(u_{d,k}) + \mathbf{P}_{ne}\left[\mathrm{T}^k\mathbf{J}\right]\mathbb{1}_{\{d_k=0\}}\mathbb{1}_{\{s_k=0\}} + \mathbf{P}_e\left[\mathrm{T}^k\mathbf{J}\right]\mathbb{1}_{\{d_k=0\}}\mathbb{1}_{\{s_k=1\}}\Big], \quad (47)$$

where T denotes the transformation operator. **J** and $\mathbf{g}(u_{d,k})$ are as follows

$$\mathbf{J} = \begin{bmatrix} J(1) & J(2) & \cdots & J(N_p) \end{bmatrix}^T, \quad (48)$$

$$\mathbf{g}(u_{d,k}) = \begin{bmatrix} g(1, u_{d,k}) & g(2, u_{d,k}) & \cdots & g(N_p, u_{d,k}) \end{bmatrix}^T. \quad (49)$$

$J(i)$ represents the cost function value when the discretized initial state p_0 is i. $g(i, u_{d,k})$ is evaluated using (28) after replacing the state i with a corresponding value of $p_k \in \mathbb{R}$.

Step-4: ADD, FAR and ANW are estimated for every grid point in the search space with the corresponding optimal policy u_d^* using MC simulations, and the suitable values for λ_f, λ_e, and μ_s are selected which satisfy the given constraints, *i.e.*, $FAR \leq FAR_{th}$ and $ANW \leq ANW_{th}$.

4872

Remark 3: We have studied several SISO models and found that the optimal policy u_d^* is a two thresholds policy, Th^s and Th^d, $Th^s \leq Th^d$, on p_k. According to the policy, $s_k = 1$ is selected if $p_k \geq Th^s$, and $d_k = 1$ is decided if $p_k \geq Th^d$.

Remark 4: A more systematic way of finding the optimal values of λ_e, λ_e and μ_s for the given FAR_{th} and ANW_{th} is currently being investigated.

IV. APPROXIMATE ANALYTICAL EXPRESSIONS OF ADD AND FAR

To derive asymptotic approximate expressions of ADD and FAR for the replay attack detection problem under study, p_k is transformed to Z_k [19], see (50), so that when $p_k \to 1$, $Z_k \to \infty$.

$$Z_k = \log \frac{p_k}{1 - p_k} \tag{50}$$

Similarly, $Th^S = \log \frac{Th^s}{1 - Th^s}$ and $Th^D = \log \frac{Th^d}{1 - Th^d}$. Using (29) and (50), the recursion equation of Z_k can be written as

$$Z_{k+1} = Z_k + |\log(1 - \rho)| + \log(1 + \rho \exp(-Z_k)) \tag{51}$$
$$+ \mathbb{1}_{\{Z_k < Th^S\}} \log L(\gamma_{k+1}) + \mathbb{1}_{\{Z_k \geq Th^S\}} \log L(\gamma_{k+1}, e_k)$$

Lemma 4: The variable Z_n, i.e., $Z_n = \log \frac{p_n}{1 - p_n}$, can be written as the summation of two terms, S_n and l_n, see (52), where S_n is a ladder variable, see (53), and l_n is a slowly changing variable in the sense defined in [23], see (54).

$$Z_n = S_n + l_n \tag{52}$$

$$S_n = \sum_{k=1}^{n} \log L(\gamma_{k+1}, e_k) + n|\log(1 - \rho)| \tag{53}$$

$$l_n = \sum_{k=1}^{n-1} \log(1 + \rho \exp(-Z_k)) + \sum_{k=1}^{n} \mathbb{1}_{\{Z_k < Th^S\}} \log L(\gamma_{k+1})$$
$$- \sum_{k=1}^{n} \mathbb{1}_{\{Z_k < Th^S\}} \log L(\gamma_{k+1}, e_k) + \log(\exp(Z_0) + \rho) \tag{54}$$

Proof: The proof of Lemma 4 is provided in Appendix III. ∎

Exploiting the special structure of Z_n as given in Lemma 4, and applying the analyses from [17], [19], we can get the following approximate asymptotic expressions of ADD and FAR.

$$ADD \to \frac{Th^D + \bar{r} - \bar{l}}{D\left(\widetilde{f}(\gamma_{k+1}, e_k), f(\gamma_{k+1}, e_k)\right) + |\log(1 - \rho)|}, \tag{55}$$

and $$FAR \to \frac{\xi}{\exp(Th^D)}, \text{ as } Th^D \to \infty, \tag{56}$$

where r is the overshoot, $r \triangleq S_\eta - Th^D$. η is the stopping time, $\eta \triangleq \inf\{n \geq 1 : S_n \geq Th^D\}$. $\bar{r} = E_1[r]$, $\bar{l} = \lim_{n \to \infty} E_1[l_n]$, and $\xi = E_1[\exp(-r)]$. $D\left(\widetilde{f}(\gamma_{k+1}, e_k), f(\gamma_{k+1}, e_k)\right)$ denotes the Kullback Leibler divergence (KLD) measure between the distributions after and before the replay attack. For the current problem

$D\left(\widetilde{f}(\gamma_{k+1}, e_k), f(\gamma_{k+1}, e_k)\right)$ will take the following form [4],

$$D\left(\widetilde{f}(\gamma_{k+1}, e_k), f(\gamma_{k+1}, e_k)\right)$$
$$= 0.5\left(\frac{\sigma_{\widetilde{\gamma}}^2}{\sigma_\gamma^2} - 1 - \log\left(\frac{\sigma_{\widetilde{\gamma}}^2 - C^2 B^2 \sigma_2^2}{\sigma_\gamma^2}\right)\right) \tag{57}$$

Remark 5: There are no closed-form analytical expressions of $\bar{r}$, $\bar{l}$ and ξ available for the replay attack detection problem under study. We have estimated those values by MC simulations. ADD expression in (55) can be further approximated by ignoring the terms $\bar{r}$ and $\bar{l}$ as,

$$ADD \approx \frac{Th^D}{D\left(\widetilde{f}(\gamma_{k+1}, e_k), f(\gamma_{k+1}, e_k)\right) + |\log(1 - \rho)|}. \tag{58}$$

V. NUMERICAL RESULTS

We have performed simulation studies to illustrate the proposed method and validate the theoretical results. An open-loop unstable SISO system is used for the simulation studies with the following parameters, $A = 1.1$, $B = C = R = Q = W = 1$, $U = 0.4$, $\rho = 0.01$, $N_p = 50$, $N_{\lambda_f} = 100$, $N_{\lambda_e} = 100$, $N_{\mu_s} = 10$, $\lambda_{f,max} = 1000$, and $\lambda_{e,max} = 1$. The proposed method is compared with a method where watermarking is always present, which is referred to as Method-A.

Figure 4a shows a plot of the optimal policy u_d^* with respect to p_k for a set of fixed values of λ_f, λ_e, and μ_s. So, for the example test case of Fig. 4a, the optimal policy is a two threshold policy on p_k.

Figure 4b plots p_k, and the two control variables s_k and d_k from a test run with the optimal policy given in Fig. 4a. It can be observed that the watermarking is added only for a few instances when $p_k \geq Th^s$. But after the attack point, the watermarking is added more frequently compared to the time before the attack.


Fig. 4: Trial run, (a) Optimal policy u_d^* vs. p_k plot, (b) Plot of s_k, d_k and p_k for a random test run.

Figure 5 plots the ADD and FAR (%) vs. σ_e^2 for the proposed method. ADD and FAR are derived from MC simulations. The plots are shown for two different values of λ_e while keeping λ_f and μ_s fixed. Figure 6a and Figure 6b plot the same variables as in Fig 5 for the case where λ_f and μ_s are changed, respectively, keeping the other two variables fixed.

Fig. 5: ADD and FAR (%) vs. σ_e^2 plots for two different λ_e. $\lambda_f = 100$ and $\mu_s = 0.05$.

Fig. 6: ADD and FAR (%) vs. σ_e^2 plots , (a) two different λ_f. $\lambda_e = 0.3$ and $\mu_s = 0.05$, (b) two different μ_s. $\lambda_e = 0.3$ and $\lambda_f = 200$.

From Fig. 5 - Fig. 6b, we can have the following observations. ADD and FAR increases with the increase in λ_e. If we increase λ_f then the FAR reduces at the expense of ADD. It seems, μ_s does not have any major effect on ADD and FAR as long as it is small. By a grid search, we can find proper values of λ_e for each σ_e^2 that will give the ANW which will match with the prior assumption of μ_s. Figure 7 plots the ADD and FAR (%) vs. σ_e^2 for the proposed method, where ADD and FAR are estimated form MC simulations and derived approximate expressions (55), (58) and (56). The approximation of (58) gives poor accuracy compared to (55), but it does not need the $\bar{r}$ and $\bar{l}$ values.


Fig. 7: Comparison between the estimated values and theoretical values. ADD and FAR (%) vs. σ_e^2 plot. $\lambda_f = 1000$. $\lambda_e = 0.3$ and $\mu_s = 0.05$

Figure 8a plots the ADD and FAR (%) vs. σ_e^2 for the proposed method and Method-A for a fixed set of values for λ_f, λ_e and μ_s. ADD and FAR are estimated using MC simulations. The ADD is high for the proposed method compared to Method-A. However, the difference is small, *i.e.*, 36% (approx.) for the example model used in this paper. On the other hand, there is not much difference between the two methods for the FAR.

Figure 8b shows the ΔLQG vs σ_e^2 plots for the proposed

(a)

(b)

Fig. 8: Comparison between proposed method and Method-A. (a) ADD and FAR (%) vs. σ_e^2 plot. $\lambda_f = 200$. $\lambda_e = 0.3$ and $\mu_s = 0.05$, (b) ΔLQG vs. σ_e^2 plot for two different λ_e. $\lambda_f = 100$ and $\mu_s = 0.05$.

method and Method-A, for two different values of λ_e, when λ_f and μ_s are kept fixed. ΔLQG values are estimated by MC simulations. We can observe a large improvement in ΔLQG (80% approx.) for the proposed method when compared to Method-A. Also, a large λ_e gives a better improvement in ΔLQG, but at the same time, it also increases ADD.

VI. CONCLUSION

We have proposed and demonstrated a sequential method for detecting replay attacks by parsimoniously adding watermarking to the control input. The proposed scheme reduces the overall control cost during the normal operation of the system at the expense of a small increase in ADD. We have also derived a few theoretical results, such as the asymptotic approximate expressions of ADD and FAR for the proposed replay attack detection scheme. The numerical simulations illustrate the efficacy of the proposed method and validate the derived theoretical results. In our future work, the proposed method will be extended for more generalized MIMO system models. One can also find an analytical expression for the ΔLQG for the proposed method.

APPENDIX I
PROOF OF LEMMA 2

The correlation between γ_k and e_{k-1}, *i.e.*, $-BC\sigma_e^2$, and σ_γ^2 are derived following the same same steps as given in [4]. Using (34) for the case $k \geq \Gamma$, $\sigma_{\widetilde{\gamma}}^2$ is derived as

$$\sigma_{\widetilde{\gamma}}^2 = \sigma_z^2 - 2C(A + BL)E_{\hat{x}z}(-1) + C^2(A + BL)^2\sigma_{\hat{x}}^2 + C^2B^2\widetilde{\mu}_s\sigma_e^2, \tag{59}$$

where $\sigma_z^2 = E_0[y_{k-k_0}^2] = \sigma_y^2$, $E_{\hat{x}z}(-1) = E_1[\hat{x}_{k-1|k-1}z_k]$, $\sigma_{\hat{x}}^2 = E_1[\hat{x}_{k|k}^2]$, and $\widetilde{\mu}_s = E_1[s_k]$. e_{k-1} is uncorrelated to $\hat{x}_{k-1|k-1}$ and z_k. Due to underlying stationarity and ergodicity, we assume that all the random variables converge to some steady-state distributions asymptotically, and only use these steady state distributions in the subsequent analysis. Combining (4)-(6) for $k \geq \Gamma$, and replacing y_k by z_k, we get

$$\hat{x}_{k|k} = Kz_k + \mathcal{A}\hat{x}_{k-1|k-1} + B(1 - CK)e_{k-1}s_{k-2}, \tag{60}$$

where $\mathcal{A} = (1 - CK)(A + BL)$. Multiplying (60) with itself, taking expectations on both sides, and finally rear-

ranging the terms we get,

$$\sigma_{\hat{x}}^2 = \frac{K^2}{1-\mathcal{A}^2}\sigma_y^2 + \frac{2\mathcal{A}K}{1-\mathcal{A}^2}E_{\hat{x}z}(-1) + \frac{B^2(1-CK)^2\widetilde{\mu}_s}{1-\mathcal{A}^2}\sigma_e^2. \tag{61}$$

Using (61) in (59), and rearranging the terms we get (40). To derive σ_y^2 (41), the following expression of y_k for the case before the attack is used.

$$y_k = \gamma_k + C(A+BL)\hat{x}_{k-1|k-1} + CBe_{k-1}s_{k-2}. \tag{62}$$

From (62), σ_y^2 is derived as

$$\sigma_y^2 = C^2P + R + C^2(A+BL)^2\sigma_{\hat{x}}^2 + C^2B^2\mu_s\sigma_e^2, \tag{63}$$

where $\mu_s = \mathrm{E}_0[s_k]$. $\sigma_{\hat{x}}^2$ is derived using the similar steps as before for the case $k < \Gamma$ as

$$\sigma_{\hat{x}}^2 = \frac{K^2(C^2P+R) + B^2\mu_s\sigma_e^2}{1-(A+BL)^2}. \tag{64}$$

Using (64) in (62), we get (41).

APPENDIX II
PROOF OF LEMMA 3

To derive the expression of $E_{\hat{x}z}(-1)$, we first assume that the fake observation z_k is generated form the following partially observed Gaussian Markov process (GMP) as follows,

$$\mathbf{x}_{a,k} = \mathbf{A}_a\mathbf{x}_{a,k-1} + \mathbf{w}_{a,k-1}, \tag{65}$$

$$z_k = \mathbf{C}_a\mathbf{x}_{a,k}. \tag{66}$$

Now, for the replay attack, where $z_k = y_{k-k_0}$, the elements of the above GMP will take the following forms,

$$\mathbf{x}_{a,k} = \begin{bmatrix} x_{k-k_0} & \hat{x}_{k-k_0|k-k_0-1} & v_{k-k_0} \end{bmatrix}^T, \tag{67}$$

$$\mathbf{w}_{a,k} = \begin{bmatrix} Be_{k-k_0} + w_{k-k_0} & Be_{k-k_0} & v_{k-k_0+1} \end{bmatrix}^T, \tag{68}$$

$$\mathbf{A}_a = \tag{69}$$

$$\begin{bmatrix} A+BLKC & BL(1-KC) & BLK \\ (A+BL)KC & (A+BL)(1-KC) & (A+BL)K \\ 0 & 0 & 0 \end{bmatrix},$$

$$\mathbf{C}_a = \begin{bmatrix} C & 0 & 1 \end{bmatrix}, \tag{70}$$

$$\mathbf{Q}_a = \begin{bmatrix} B^2\mu_s\sigma_e^2 + Q & B^2\mu_s\sigma_e^2 & 0 \\ B^2\mu_s\sigma_e^2 & B^2\mu_s\sigma_e^2 & 0 \\ 0 & 0 & R \end{bmatrix}. \tag{71}$$

Using the given expressions (67) - (71), we can derive (42) and (43) using the same steps given in [24].

APPENDIX III
PROOF OF LEMMA 4

Using the recursion of Z_k (51) and starting from the initial value Z_0, (52)-(54) can be derived directly. After the attack start point, $\exp(-Z_k) \to 0$ as $k \to \infty$, and $Z_k < Th^S$ will be true for a very short interval of time. Therefore, l_n will converge to a finite value as $n \to \infty$, which is a property of a slowly changing variable as defined in [23]. However, a formal proof is not provided due to space constraints and will be provided in the future extended version of the work.

REFERENCES

[1] D. Ding, Q. L. Han, Y. Xiang, X. Ge, and X. M. Zhang, "A survey on security control and attack detection for industrial cyber-physical systems," *Neurocomputing*, vol. 275, pp. 1674–1683, 2018.

[2] Y. Mo, S. Weerakkody, and B. Sinopoli, "Physical authentication of control systems: Designing watermarked control inputs to detect counterfeit sensor outputs," *IEEE Control Syst.*, vol. 35, no. 1, pp. 93–109, jan 2015.

[3] C. Peng and H. Sun, "Switching-like event-triggered control for networked control systems under malicious denial of service attacks," *IEEE Transactions on Automatic Control*, vol. 65, no. 9, pp. 3943–3949, 2020.

[4] S. Salimi, S. Dey, and A. Ahlen, "Sequential detection of deception attacks in networked control systems with watermarking," *2019 18th Eur. Control Conf. ECC 2019*, pp. 883–890, 2019.

[5] Y. Zhao and C. Smidts, "A control-theoretic approach to detecting and distinguishing replay attacks from other anomalies in nuclear power plants," *Prog. Nucl. Energy*, vol. 123, no. March, p. 103315, 2020.

[6] M. Hosseinzadeh, B. Sinopoli, and E. Garone, "Feasibility and Detection of Replay Attack in Networked Constrained Cyber-Physical Systems," *2019 57th Annu. Allert. Conf. Commun. Control. Comput. Allert. 2019*, pp. 712–717, 2019.

[7] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *IEEE Secur. Priv.*, vol. 9, no. 3, pp. 49–51, 2011.

[8] R. Alguliyev, Y. Imamverdiyev, and L. Sukhostat, "Cyber-physical systems and their security issues," *Comput. Ind.*, vol. 100, no. July 2017, pp. 212–223, 2018.

[9] B. Satchidanandan and P. R. Kumar, "Dynamic Watermarking: Active Defense of Networked Cyber–Physical Systems," *Proc. IEEE*, vol. 105, no. 2, pp. 219–240, feb 2017.

[10] L. Zhai and K. G. Vamvoudakis, "A data-based private learning framework for enhanced security against replay attacks in cyber-physical systems," *Int. J. Robust Nonlinear Control*, no. January, pp. 1–17, 2020.

[11] C. Fang, Y. Qi, P. Cheng, and W. X. Zheng, "Optimal periodic watermarking schedule for replay attack detection in cyber–physical systems," *Automatica*, vol. 112, 2020.

[12] R. M. Ferrari and A. M. Teixeira, "Detection and Isolation of Replay Attacks through Sensor Watermarking," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 7363–7368, 2017.

[13] H. S. Sánchez, D. Rotondo, T. Escobet, V. Puig, J. Saludes, and J. Quevedo, "Detection of replay attacks in cyber-physical systems using a frequency-based signature," *J. Franklin Inst.*, vol. 356, no. 5, pp. 2798–2824, 2019.

[14] D. Ye, T. Y. Zhang, and G. Guo, "Stochastic coding detection scheme in cyber-physical systems against replay attack," *Inf. Sci. (Ny).*, vol. 481, no. 61773097, pp. 432–444, 2019.

[15] C. Trapiello, D. Rotondo, H. Sanchez, and V. Puig, "Detection of replay attacks in CPSs using observer-based signature compensation," *2019 6th Int. Conf. Control. Decis. Inf. Technol. CoDIT 2019*, pp. 1–6, 2019.

[16] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory Probab. Its Appl.*, vol. 8, no. 1, pp. 22–46, 1963.

[17] A. G. Tartakovsky and V. V. Veeravalli, "General asymptotic Bayesian theory of quickest change detection," *Theory of Probability and its Applications*, vol. 49, no. 3, pp. 458–497, 2005.

[18] K. Premkumar and A. Kumar, "Optimal sleep-wake scheduling for quickest intrusion detection using sensor networks," in *Proc. - IEEE INFOCOM*, no. ii, 2008, pp. 2074–2082.

[19] T. Banerjee and V. V. Veeravalli, "Data-Efficient Quickest Change Detection with On-Off Observation Control," *Seq. Anal.*, vol. 31, no. 1, pp. 40–77, 2012.

[20] D. P. Bertsekas, *Dynamic programming and optimal control (vol 1 and 2)*. Athena scientific Belmont, MA, 1995.

[21] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*, 2014.

[22] F. J. Beutler and K. W. Ross, "Optimal policies for controlled markov chains with a constraint," *Journal of mathematical analysis and applications*, vol. 112, no. 1, pp. 236–252, 1985.

[23] D. Siegmund, *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media, 2013.

[24] A. Naha, A. Teixeira, A. Ahlen, and S. Dey, "Sequential detection of replay attacks," *arXiv preprint arXiv:2012.10748*, 2020.