

A Bayesian Nash equilibrium-based moving target defense against stealthy sensor attacks

David Umsonst, Serkan Sarıtaş, György Dán and Henrik Sandberg

Abstract— We present a moving target defense strategy to reduce the impact of stealthy sensor attacks on feedback systems. The defender periodically and randomly switches between thresholds from a discrete set to increase the uncertainty for the attacker and make stealthy attacks detectable. However, the defender does not know the exact goal of the attacker but only the prior of the possible attacker goals. Here, we model one period with a constant threshold as a Bayesian game and use the Bayesian Nash equilibrium concept to find the distribution for the choice of the threshold in that period, which takes the defender's uncertainty about the attacker into account. To obtain the equilibrium distribution, the defender minimizes its cost consisting of the cost for false alarms and the cost induced by the attack. We present a necessary and sufficient condition for the existence of a moving target defense and formulate a linear program to determine the moving target defense. Furthermore, we present a closed-form solution for the special case when the defender knows the attacker's goals. The results are numerically evaluated on a four-tank process.

Index Terms— Bayesian games, Cyber-physical security, Game theory, Optimization, Optimal Control, Moving Target Defense, Detection threshold, False data injection attacks

I. INTRODUCTION

Critical infrastructures and control systems are increasingly connected to public communication networks, such as the Internet, and constitute geographically distributed cyber-physical systems (CPS). The use of public network infrastructures can save costs, for example cabling, but also increase the performance of the CPS. However, this interconnection comes at the price of vulnerability to cyber-attacks, which are already impacting critical infrastructures, such as the Ukrainian power grid [1], as well as industrial control systems, such as a steel mill in Germany [2].

Date of Submission: February 25, 2021. This work was supported in part by the Swedish Research Council (grants 2016-00861 and 2020-03860), and the Swedish Civil Contingencies Agency through the CERCES2 project.

David Umsonst and Henrik Sandberg are with the Division of Decision and Control Systems in the School of Electrical Engineering and Computer Science at the KTH Royal Institute of Technology, 10044 Stockholm, Sweden. {umsonst, hsan}@kth.se

Serkan Sarıtaş is with the Department of Electrical and Electronics Engineering, Middle East Technical University, 06800, Ankara, Turkey. ssaritas@metu.edu.tr

György Dán is with the Division of Network and Systems Engineering in the School of Electrical Engineering and Computer Science at the KTH Royal Institute of Technology, 10044 Stockholm, Sweden. gyuri@kth.se

To improve CPS security and complement existing information technology (IT) security measures, such as encryption and authentication, a new branch of security measures based on control-theoretic methods has emerged over the last decade. These novel security measures are based on the physics of the CPS and use physical models to detect, isolate, and mitigate malicious attacks. Hence, the control-theoretic security measures are a complementary approach to the IT security measures. In the authors' opinion the main difference between IT security measures and control-theoretic security measures is that IT security measures consider cryptography and logical isolation, while control-theoretic security measures are based on the physical models of the closed-loop system and can verify its operational consistency. Consequently, control-theoretic security measures also allow for safety aspects to be taken into account. Thus, using both control-theoretic and IT security measures contributes directly to a defense-in-depth approach to security. An introduction to this topic can be found in the tutorial papers [3] and [4]. Since the attacker and the operator/defender are rational entities, their interaction is strategic and can thus be modeled using game-theoretic tools, see, for example, [5].

An emerging approach to detect attacks and to limit their impact, which can combine both physical models and game theory, are moving target defense (MTD) strategies [6] that induce controlled uncertainties into the CPS to confuse the attacker. Gairo *et al.* [7], for example, randomly switch between the sensors used to detect otherwise stealthy attacks, while in [8] a random watermarking signal is injected into the CPS to make stealthy attacks detectable. Furthermore, perturbations of power line impedances are analyzed in [9], where an in-depth analysis is conducted to determine when the MTD will be successful. Another system-switching approach is considered in [10], which considers both actuator and sensor attacks. However, the MTD strategies in [7]–[10] directly influence the closed-loop behavior of the CPS and can decrease its performance. Griffioen *et al.* [11] propose three different MTD schemes, where the first one is similar to [7] but it also switches the plant and input matrix and not only the measurements used. The second MTD of [11] introduces an auxiliary system to not influence the closed-loop behavior and simultaneously detect an attack, while the third MTD utilizes the nonlinearities in the measurements. In this work, we propose a moving target defense that is placed in the anomaly detector of the CPS, which is located outside of the control loop (see Figure 1). Therefore, the proposed MTD

neither influences the closed-loop performance directly, nor is there a need of introducing new auxiliary components to the system such that the controller and the moving target defense can be designed independently.

A. Contribution

When using an anomaly detector the defender faces a trade-off between the cost for false alarms and the cost for the impact of a stealthy attack, where ideally both costs should be as small as possible. However, fewer false alarms typically lead to a larger attack impact, and vice versa, such that we cannot minimize both costs at the same time. Therefore, we formulate a game where the defender periodically chooses a detector threshold at random to mitigate the trade-off between its cost for false alarms and its cost induced by the stealthy attack launched by the attacker. The goal of the attacker is to maximize its payoff, which can, for example, be characterized by an unsafe region in the system's state space, while the defender wants to minimize the cost induced by false alarms and the cost induced by the attack. However, the defender is uncertain about the payoff function the attacker tries to optimize and only has a belief of facing an attacker with a certain payoff function. Here, we present an initial analysis of this game. We consider a single period with a constant threshold and look at the threshold choice for this period. We show that there is an equivalent matrix game to analyse the equilibrium strategies of each player.

The matrix game formulation is used to provide a necessary and sufficient condition for when a Bayesian Nash equilibrium exists in which the defender's strategy is mixed and does not concentrate the whole probability on one action. The defender's equilibrium strategy is then a moving target defense strategy. Furthermore, by using the structure of the matrix game, we show that the Bayesian Nash equilibrium can be obtained by solving a linear program. For the special case where the defender knows the attacker's type, we provide a closed-form solution for the Nash equilibrium, which gives us insights about the equilibrium strategies of the defender and attacker. Finally, we numerically verify our results with a four-tank system.

B. Related Work

Since control systems are typically equipped with an anomaly detector to detect faults, several research groups have investigated how the choice and tuning of the anomaly detector threshold can help limiting the attack impact of stealthy attacks. When it comes to the choice of the detector, Murguia *et al.* [12] compare a χ^2 and a CUSUM detector and investigate which detector mitigates the impact of a sensor attack the most.

In the present work, we are interested in the case where the detector is already chosen and we want to define a way to choose the thresholds to limit the attack impact. Urbina *et al.* [13] point out that there will be a trade-off between the number of false alarms and the maximum impact of a stealthy attack when tuning the anomaly detector.

There are several other works that use the anomaly detector threshold to limit the attack impact or to detect attacks. Ghafouri *et al.* [14] propose a Stackelberg game framework for choosing the detector threshold. Both a static choice as well as a dynamic choice of the detector threshold are presented, but the attack is assumed to be detectable. The cost that the defender wants to minimize is composed of the cost of false alarms, the cost of the attack impact, and the cost for switching between thresholds.

In [15], we extend the static detector threshold choice of [14] to the case of stealthy sensor attacks and prove the existence of such a threshold and provide conditions for the uniqueness.

Niu *et al.* [16] formulated the detector threshold switching problem as a zero-sum Stackelberg game without considering the cost for false alarms.

In [17], we consider a similar game but there the defender exactly knows the attacker's objective. This assumption is relaxed in the present work since the defender only has a prior over possible attacker objectives. Therefore, we extend the results of [17] to a broader class of games, namely Bayesian games. Here, we also provide a closed-form solution to the special case considered in [17].

C. Notation

Let $x \in \mathbb{R}^n$ be an n -dimensional column vector and $A \in \mathbb{R}^{m \times n}$ be an m -by- n matrix. The i th element of x is denoted by x_i and A_{ij} corresponds to the element in the i th row and j th column of A . Further $x_{i:j}$ is the vector $[x_i, x_{i+1}, \dots, x_{j-1}, x_j]^T$, where $i \leq j$. If a random variable x has a Gaussian distribution with mean $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$, we denote it as $x \sim \mathcal{N}(\mu, \Sigma)$. The expected value of a random variable x is denoted by $\mathbb{E}\{x\}$. The n -by- n identity matrix is denoted by I_n and an n -dimensional column vector with all elements equal to one as $\mathbf{1}_n$, while the indicator function of an event D is represented by $\mathbf{1}_{\{D\}}$.

II. SYSTEM MODEL

In this section, we introduce the models for the plant, controller, and detector and present our assumptions on the attacker and the defender. Further, in Section II-E we will discuss the assumptions made on the system, the attacker, and the defender. Figure 1 shows a block diagram of the sensor attack scenario that we consider.

A. Plant and Controller Model

In our setup, the plant receives actuator signals and sends measurement signals over a network. We model the plant in Fig. 1 as a linear discrete-time system,

$$\begin{aligned} x(k+1) &= Ax(k) + B\tilde{u}(k) + w(k), \\ y(k) &= Cx(k) + v(k), \end{aligned}$$

where $x(k) \in \mathbb{R}^{n_x}$ is the plant's state, $\tilde{u}(k) \in \mathbb{R}^{n_u}$ is the actuator signal received over the network, $y(k) \in \mathbb{R}^{n_y}$ is the measurement signal, $w(k) \in \mathbb{R}^{n_x}$ is the process noise,

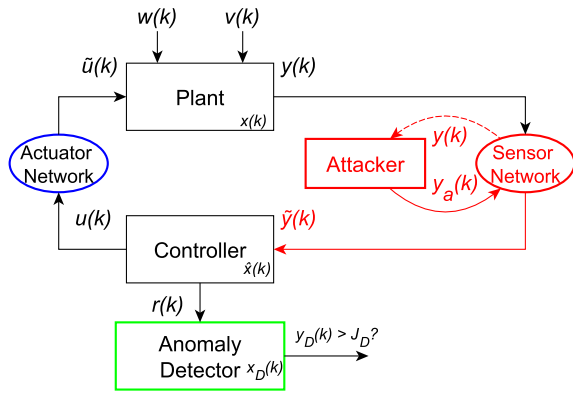


Fig. 1. Block diagram of the sensor attack scenario.

and $v(k) \in \mathbb{R}^{n_y}$ is the measurement noise. Both $w(k)$ and $v(k)$ have independent and identically distributed zero-mean multivariate Gaussian distributions with covariance matrices Σ_w and Σ_v , respectively. Further, $w(k)$ and $v(k)$ are independent processes. The system, input, and output matrices are $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$, and $C \in \mathbb{R}^{n_y \times n_x}$, respectively.

Assumption 1: (A, B) is stabilizable, (C, A) is detectable, and $(A, \Sigma_w^{-\frac{1}{2}})$ has no uncontrollable modes on the unit circle.

The plant is controlled using a Kalman filter-based observer, which estimates the plant's state as $\hat{x}(k) \in \mathbb{R}^{n_x}$. The dynamics of the controller are

$$\begin{aligned} \hat{x}(k+1) &= A\hat{x}(k) + Bu(k) + L(\tilde{y}(k) - C\hat{x}(k)), \\ u(k) &= -K\hat{x}(k), \end{aligned}$$

where $\tilde{y}(k)$ is the measurement signal received over the network, $u(k)$ is the actuator signal determined by the controller, and K and L are the controller gain and steady-state Kalman gain, respectively. Further, $L = APC^T(CPC^T + \Sigma_v)^{-1}$, where P is the stabilizing solution to the algebraic Riccati equation

$$P = APA^T + \Sigma_w - (APC^T)(CPC + \Sigma_v)^{-1}(APC^T)^T,$$

and P exists due to Assumption 1.

B. Detector Model

Since faults and/or malicious attacks can occur, the closed-loop system is equipped with an anomaly detector on the controller side, which has the possibly nonlinear dynamics

$$\begin{aligned} x_D(k+1) &= \theta(x_D(k), r(k)), \\ y_D(k+1) &= d(x_D(k), r(k)), \end{aligned} \quad (1)$$

where $x_D(k) \in \mathbb{R}^{n_D}$ is the detector's internal state, $y_D(k) \in \mathbb{R}_{\geq 0}$ is the detector output, and $r(k) \in \mathbb{R}^{n_y}$ is the detector input. The exact structure of $\theta(x_D(k), r(k))$ and $d(x_D(k), r(k))$ depends on the detector the defender will use. For example, in Section VII we consider the static χ^2 detector, i.e., $y_D(k+1) = \|r(k)\|_2^2$. More detector models can be found in [18], [19].

We define the input $r(k)$ to be the residual signal, which is the normalized difference between the received and the predicted measurements, i.e.,

$$r(k) = \Sigma_r^{-\frac{1}{2}} (\tilde{y}(k) - C\hat{x}(k)),$$

where $\Sigma_r = CPC^T + \Sigma_v$ is the steady state covariance matrix of $\tilde{r}(k) = \tilde{y}(k) - C\hat{x}(k)$ under nominal conditions (no faults, no attacks), i.e., $\tilde{u}(k) = u(k)$ and $\tilde{y}(k) = y(k)$ for all k .

Assumption 2: The detector dynamics (1) fulfill the subsequent three conditions:

- 1) $\theta(x_D(k), r(k))$ and $d(x_D(k), r(k))$ are continuous in $x_D(k)$ and $r(k)$;
- 2) $d(x_D(k), r(k))$ is coercive¹ in $x_D(k)$ and $r(k)$;
- 3) $d(0, 0) = 0$ and $\theta(0, 0) = 0$.

If the predictions are accurate, i.e., $r(k) \approx 0$, the detector output should be small. However, if the predictions are inaccurate, both the detector state and the detector output should increase. Furthermore, the detector triggers an alarm whenever the detector output $y_D(k)$ exceeds the detection threshold $J_D > 0$.

Since the detector input is a random variable under nominal conditions, that is, $r(k) \sim \mathcal{N}(0, I_{n_y})$ due to the Kalman filter, the detector output $y_D(k)$ is also a random variable. To avoid too frequent *false alarms*, i.e., alarms triggered under nominal conditions when there is *no attacker present*, the threshold J_D should be chosen large enough. However, if it is chosen too large, the detector might not be able to detect anomalies (missed detections). Hence, there is a trade-off between false alarms and missed detections when choosing J_D . Urbina *et al.* [13] further noted that there is also a trade-off between false alarms and the attack impact of a stealthy attack, which in our case corresponds to a missed detection. For example, a larger J_D reduces the frequency of false alarms but gives the attacker more space to remain stealthy while causing harm.

Since the amount of false alarms plays an important role in detector tuning, we denote by τ the *mean time between false alarms*. The larger time between false alarms we want to achieve, the larger the detector threshold has to be such that the following is a reasonable assumption.

Assumption 3: The detector threshold is a strictly increasing (possibly nonlinear) function of τ , i.e., $J_D(\tau_a) < J_D(\tau_b)$ if, and only if, $\tau_a < \tau_b$.

Instead of considering the threshold, J_D , we will consider the mean time between false alarms, τ , in the following, since there is a direct relation between J_D and τ . Further, the value of τ is more meaningful to the operator. To circumvent the trade-off between false alarms and missed detections (and the impact of stealthy attacks), the defender could periodically randomize the choice of the mean time between false alarms τ such that in one period the threshold reduces the number of false alarms while in another it limits the impact of a potential stealthy attack. That is, it chooses τ periodically from the fixed set $\{\tau_1, \dots, \tau_m\}$, pre-determined by the defender, with probability distribution p , where $1 \leq \tau_1 < \tau_2 < \dots < \tau_m$, $p_i \in [0, 1]$ is the probability of choosing τ_i , and $\sum_{i=1}^m p_i = 1$.

¹A function $g(x)$ is called coercive if $g(x) \rightarrow \infty$ as $\|x\| \rightarrow \infty$.

We make the following assumption about the random choice of τ .

Assumption 4: At the beginning of each period, τ is drawn from the probability distribution p , independent from previous realizations.

Our definition of a moving target defense is stated next.

Definition 1: A probability distribution $p \in \mathbb{R}^m$ over a fixed set of mean times between false alarms $\{\tau_1, \dots, \tau_m\}$ is a *moving target defense* if p does not have singleton support, i.e., the probability of choosing τ_i fulfills $p_i \in [0, 1)$ for all $i \in \{1, \dots, m\}$.

C. Attacker Model

In this paper, our focus lies on sensor attacks.

Assumption 5: The measurement signals are subject to an additive attack $y_a(k)$ chosen by the attacker, and the actuator signals are transmitted fault/attack-free, i.e., $\tilde{y}(k) = y(k) + y_a(k)$ and $\tilde{u}(k) = u(k)$.

Furthermore, we make the following assumption on the attacker's model knowledge.

Assumption 6: The attacker knows the closed-loop system matrices, A, B, C, L, K , the noise statistics Σ_w, Σ_v , and the detector dynamics. The attack starts at time $k = 0$ and has a length of N time steps. The length N of the attack is such that the attacker is able to complete the attack before the next threshold switch. The attacker knows both $x_D(0)$ and $\hat{x}(0)$, and has access to the measurements $y(k)$. Moreover, the attacker knows the function $J_D(\tau)$ and the set $\{\tau_1, \dots, \tau_m\}$ but *not* the exact value of τ .

An attacker according to Assumption 6 can launch an attack of the form (see [12])

$$y_a(k) = -y(k) + C\hat{x}(k) + \Sigma_f^{\frac{1}{2}}a(k),$$

which gives the attacker complete control over the detector input, i.e., $r(k) = a(k)$. This attack is a *closed-loop* attack since it uses the measurements $y(k)$, whereas $a(k)$ can be interpreted as the attacker's reference signal. The attacker can define the set of attacks that do not trigger an alarm for a given τ as

$$\mathbb{A}(\tau) := \left\{ \{a, x_D(0)\} \left| \begin{array}{l} x_D(k+1) = \theta(x_D(k), a(k)) \\ y_D(k+1) = d(x_D(k), a(k)) \leq J_D(\tau) \\ k \in [0, N-1] \end{array} \right. \right\},$$

where $a = [a(0)^T, \dots, a(N-1)^T]^T$ is the complete attack trajectory during the attack. Here, the set $\mathbb{A}(\tau)$ constrains the size of $y_a(k)$ as well by constraining $a(k)$.

Remark 1: Note that $\mathbb{A}(\tau_a) \subset \mathbb{A}(\tau_b)$ if $\tau_a < \tau_b$ due to Assumption 3. Thus, the threshold chosen for setting the false alarm rate under nominal conditions does affect the set of attacks that do not trigger an alarm, i.e., the false negative rate. Furthermore, where appropriate we will use $a \in \mathbb{A}(\tau)$ instead of $\{a, x_D(0)\} \in \mathbb{A}(\tau)$ for the sake of readability.

If the attacker manages to choose a such that $\{a, x_D(0)\} \in \mathbb{A}(\tau)$ in the current period with a constant threshold, the attacker remains stealthy, which is the main constraint of the attacker as described below.

Assumption 7: The attacker has one of n_ϕ different types, which determine the objective of the attacker. An attacker of type ϕ wants to maximize its expected payoff characterized by $f_\phi(a)$. Further, if the attack is detected, the attacker receives no payoff and, therefore, it wants to remain stealthy, i.e., $y_D(k+1) \leq J_D(\tau)$ for $k \in [0, N-1]$.

Next, we define the attacker's expected payoff for a given attacker type,

$$\mathfrak{p}(\tau, a|\phi) := \mathbf{1}_{\{a, x_D(0)\} \in \mathbb{A}(\tau)} f_\phi(a), \quad (2)$$

where we use the indicator function to take into account that the attacker will not get any payoff if it is detected.

Assumption 8: For a given attacker type ϕ , the corresponding expected attacker payoff $f_\phi(a)$ is continuous and fulfills $\max_{a \in \mathbb{A}(\tau_a)} f_\phi(a) < \max_{a \in \mathbb{A}(\tau_b)} f_\phi(a)$ if $\tau_a < \tau_b$ except when $f_\phi(a) = 0$ for all a .

Remark 2: If $f_\phi(a)$ is a continuous, convex function and $\mathbb{A}(\tau)$ is a closed convex set, then Assumption 8 is fulfilled, since then $\max_{a \in \mathbb{A}(\tau)} f_\phi(a)$ is equivalent to a concave minimization problem, whose optimizers are the extreme points of $\mathbb{A}(\tau)$ (see [20]). If we use a vector norm-based stateless detector, such as the χ^2 detector, $\mathbb{A}(\tau)$ is a closed convex set.

D. Defender model

Next, we describe our defender model. When choosing τ the defender needs to take into account the expected cost that is induced by the false alarms in the nominal case, but also the expected cost of an undetectable attack. This leads to the following cost function for the defender assuming an attacker of type ϕ ,

$$\mathfrak{c}(\tau, a|\phi) := \frac{c_F}{\tau} + \mathfrak{p}(\tau, a|\phi), \quad (3)$$

where $c_F > 0$ is the cost factor for false alarms. Note that while the attacker has one of n_ϕ possible types, the defender has only one type, but its cost function is influenced by the attacker type.

Remark 3: Since the defender's cost (3) is always influenced by the attacker's payoff, it is reasonable to introduce an attacker type with zero payoff, i.e., $f_\phi(a) = 0$ for all a . This means that the case of there not being an attacker present in the system is modeled as well in our moving target defense framework.

Next, let us make the following assumption about the knowledge the attacker and defender have about each other.

Assumption 9: The defender knows (2) for each possible ϕ and it has a prior π_ϕ of facing an attacker of type ϕ , where $\pi_\phi \in [0, 1]$ and $\sum_{\phi=1}^{n_\phi} \pi_\phi = 1$. The attacker, in addition to Assumption 6, knows the defender's cost function (3), its own type $\phi \in \{1, \dots, n_\phi\}$, and the defender's prior π_ϕ for each attacker type.

E. Discussion of the system model

In this section, we discuss the assumptions made during the setup of the model.

First, we discuss the assumption about the iid choice of τ (Assumption 4) and the attack length (Assumption 6). Since

the values of τ are realizations of iid random variables, the current value of τ is independent of its previous values, observing the system does hence not reveal information about τ beyond its distribution, which the attacker can deduce from its system knowledge. The attacker will be able to estimate the distribution p under the iid choice if it has access to previous values of τ . This case is already taken into account in our MTD framework. A change of threshold implies a reconfiguration of the system, which can be costly for the operator of a safety-critical large-scale infrastructure. Therefore, the operator does not switch the thresholds too frequently. Hence, it is not unreasonable to analyse the case where the attack is carried out during a fixed, but random, configuration. Note that an approach to consider the cost of a finite amount of switches can be found in [14]. Furthermore, for industrial processes, where a product is produced in batches, an iid choice of the threshold between different batches is also a reasonable assumption.

The analysis of an attacker that experiences threshold switches during the attack is similar to the analysis we present in the subsequent sections, because we can determine the probability of first choosing τ_i and then τ_j due to the iid choice in Assumption 4. However, the notation would become more involved. Therefore, these assumptions simplify the problem formulation so that it becomes mathematically more tractable.

Next, we justify the assumptions on the attacker's knowledge and goals. According to [21], one should design the plant for the worst-case attacker knowledge, because, given enough time, an attacker may be able to obtain a perfect model of the plant, the controller, and the detector. For example, the plant and controller could be estimated through system identification techniques from the observed sensor data, while the detector model could be obtained from leaked documentation of the system. Hence, the extensive knowledge of the attacker about the closed-loop system and detector according to Assumption 6 is in line with [21]. In our previous work [22], we showed how the attacker can obtain the internal states of both the controller and the detector in an experimental setup. Hence, assuming that the attacker has knowledge of the controller and detector states is not unreasonable. Further, the knowledge of $x_D(0)$ and $\hat{x}(0)$ can be interpreted as an opportunistic attacker choosing to attack at the best time instant, which we define to occur without loss of generality at $k = 0$. In contrast, the choice of τ is not visible in the sensor data observed by the attacker. In addition, the attacker knowledge in Assumption 6 together with the assumption that the attacker will maximize its objective function $f_\phi(a)$ (Assumption 7) results in a worst-case scenario for the defender under the given assumption.

While both $x_D(0)$ and $\hat{x}(0)$ depend on the measurements and could, therefore, be estimated by the attacker, τ is chosen randomly from $\{\tau_1, \dots, \tau_m\}$ (see Assumption 4) such that the attacker cannot know the exact value of τ . Further, τ does not directly influence the system variables, such that the attacker is also not able to estimate τ from the measurements.

In Assumption 7, we introduce attacker types. A given attacker type, ϕ , describes the target of the attacker through its objective, $f_\phi(a)$. Since the attacker's target is often not known to the defender, having different attacker types gives the defender the possibility to distinguish between different targets

while using our sensor attack model, and also incorporate the case of no attacker being present (Remark 3).

The assumption that the attacker will not get any payoff when detected (Assumption 7) is a strong assumption on both the attacker and defender, which is mostly beneficial for the defender. However, if we consider critical infrastructures, such as the power grid, an operator has to mitigate the attack quickly when detected to prevent harm. Furthermore, we can also imagine that the attacker has made a significant investment to obtain its system knowledge and infiltrate the system. Hence, the attacker wants to remain undetected in order to not risk losing its investment. This kind of attacker has similarities to an advanced persistent threat, which is an attacker with knowledge about the system and that targets specific parts of the system while remaining stealthy (see, for example, [23]). It is important to point out that due to the random choice of τ (Assumption 4) it is more difficult for the attacker to remain stealthy but at the same time obtain a large payoff.

The defender will rarely know the intentions of the attacker. To obtain information about potential targets of the attack, the defender can conduct a risk assessment [24] of the system. By conducting a risk assessment, the defender determines the vulnerabilities in its system, the likelihood of an attacker exploiting a vulnerability, and the potential impact of a successful attack. A vulnerability could be an unsafe region in the system's state space, e.g., the overpressure region for a tank, such that the attacker's objective would be to bring the system into this unsafe region. The different vulnerabilities can be interpreted as attacker types ϕ , the prior π_ϕ as the likelihood of an attacker exploiting a vulnerability, and the impacts are reflected by the attacker's payoff $f_\phi(a)$, which directly influence the defender's cost (3). Hence, the defender's knowledge about possible attack objectives, their prior and their impact, as assumed in Assumption 9, can be interpreted as the outcome of a risk assessment conducted by the defender before implementing the proposed MTD. Therefore, we can interpret the Bayesian moving target defense framework as a tool to enhance the security for the defender similar to the ARMOR framework deployed at LAX [25], which makes use of the outcomes of the risk assessment.

III. PROBLEM FORMULATION

Now we formulate the problem of finding a moving target defense strategy as a game between the defender and the attacker, where the defender's goal is to choose τ to minimize the expected value of (3) with respect to the prior of the attacker types while the attacker chooses a to maximize (2). Due to Assumption 6, we can focus on the game over one period with a constant threshold. This focus on one only period can also be interpreted as a repeated game with memoryless players, which has been considered in [26].

The game has both *imperfect* and *incomplete* information. The information is imperfect because neither player observes the action taken by the other player. The information is incomplete because the defender does not know which type of attacker it faces. The defender believes that with probability π_ϕ it will play the game with an attacker of type ϕ . The

imperfect information lets us interpret the game as a game with simultaneous moves, while the incomplete information results in a Bayesian game framework. Therefore, we define the moving target defense game $\mathcal{M} = \langle \mathcal{P}, \mathcal{A}, \mathcal{T}, \Pi, \mathcal{U} \rangle$, where $\mathcal{P} = \{\text{Defender}, \text{Attacker}\}$ is the set of players, $\mathcal{A} = \{\tau_1, \dots, \tau_m\} \times \mathbb{R}^{N_{ny} + n_D}$ is the action set, $\mathcal{T} = \{1\} \times \{1, \dots, n_\phi\}$ is the set of player types, $\Pi = \{1\} \times \{\pi_1, \dots, \pi_{n_\phi}\}$ is the prior, and $\mathcal{U} = (c(\tau, a|\phi), p(\tau, a|\phi))$ contains the cost and payoff functions of each player. For the analysis, we also define the game $\mathcal{M}_\phi = \langle \mathcal{P}, \mathcal{A}, \mathcal{U} \rangle$, where the defender is certain about the attacker type it faces, that is, $\pi_\phi = 1$ for some $\phi \in \{1, \dots, n_\phi\}$ in \mathcal{M} .

The Bayesian game framework together with the simultaneous choice of actions lead us to the *Bayesian Nash equilibrium* as the solution concept. To define the Bayesian Nash equilibrium we introduce the (possibly mixed) strategies of the defender and attacker. Let Δ_p be the set of probability distributions over the defender's actions. Then $p \in \Delta_p$ is a discrete probability distribution, where the i th element, p_i , is the probability that the defender chooses τ_i . For a given attacker type ϕ , $\Delta_q(\phi)$ is the set of probability distributions over the attacker's action set. Since the attacker, to obtain a non-zero payoff, chooses a trajectory a from $\mathbb{A}(\tau)$, which is typically not a discrete set, $q_\phi \in \Delta_q(\phi)$ may represent a continuous probability distribution. We call both p and q_ϕ a *mixed* strategy, if it does *not* concentrate the whole probability on one action. Otherwise, we call it a *pure* strategy.

Since both the attacker and defender might use mixed strategies, we investigate the average cost of the defender

$$\bar{c}_\phi(p, q_\phi) = \int \sum_{i=1}^m p_i c(\tau_i, a|\phi) q_\phi(a) da \quad (4)$$

and the average payoff of the attacker

$$\bar{p}_\phi(p, q_\phi) = \int \sum_{i=1}^m p_i p(\tau_i, a|\phi) q_\phi(a) da \quad (5)$$

for a given attacker type ϕ . Hence, (4) and (5) represent the average cost and payoff of the players in the game \mathcal{M}_ϕ .

A *mixed strategy Bayesian Nash equilibrium*, $p^* \in \Delta_p$ and $q_\phi^* \in \Delta_q(\phi)$, fulfills

$$\sum_{\phi'=1}^{n_\phi} \pi_{\phi'} \bar{c}_{\phi'}(p^*, q_{\phi'}^*) \leq \sum_{\phi'=1}^{n_\phi} \pi_{\phi'} \bar{c}_{\phi'}(p, q_{\phi'}^*), \quad (6)$$

$$\bar{p}_\phi(p^*, q_\phi^*) \geq \bar{p}_\phi(p^*, q_\phi)$$

for all $p \in \Delta_p$, $q_\phi \in \Delta_q(\phi)$, and $\phi \in \{1, \dots, n_\phi\}$.

In the Bayesian Nash equilibrium, a change from p^* to another $p \in \Delta_p$ does not lead to a decrease in the cost for the defender, and, similarly, a change from q_ϕ^* to another $q_\phi \in \Delta_q(\phi)$ does not lead to an increase in payoff for an attacker of type ϕ . Hence, neither the defender nor the attacker want to deviate from their equilibrium strategies. Here, the defender needs to consider all possible attacker types, which results in averaging of the costs of each game \mathcal{M}_ϕ over the prior, while the attacker needs to have an equilibrium strategy for each type. This is because the attacker knows its own type,

which the defender does not know, while the defender has only one type, which is known to both the defender and attacker.

Equipped with the definition of both the MTD and the Bayesian Nash equilibrium, we formulate the two problems we investigate in the remainder of this paper.

Problem 1: Characterize when a Bayesian Nash equilibrium (6) representing a MTD (Definition 1) exists.

Problem 2: If a Bayesian Nash equilibrium representing a MTD exists, compute an equilibrium strategy p^* .

IV. MATRIX GAME FORMULATION

Recall that the defender plays against one of n_ϕ adversaries, but it does not know which adversary it is facing. Furthermore, while the defender has a finite set of actions, i.e., $\{\tau_1, \dots, \tau_m\}$, the attacker's action set, $\mathbb{R}^{N_{ny} + n_D}$, is a continuum. This makes finding Bayesian Nash equilibrium strategies challenging. In this section, we will show that each game \mathcal{M}_ϕ can be reformulated into a strategically equivalent game $\tilde{\mathcal{M}}_\phi$, where the attacker's action set is finite too.

We begin by recalling that for a given τ the attacker will only receive a non-zero payoff if $\{a, x_D(0)\} \in \mathbb{A}(\tau)$. Hence, for a given τ the attacker will always choose its attack trajectory such that $\{a, x_D(0)\} \in \mathbb{A}(\tau)$. Due to the discrete set of actions for the defender, we can separate the continuous action space of the attacker into $m+1$ sets in the game \mathcal{M}_ϕ as shown in Table I. The set $\mathbb{A}(\tau_i) \setminus \mathbb{A}(\tau_{i-1})$ contains all attack trajectories that are stealthy for τ_i excluding the ones that are stealthy for τ_{i-1} . Hence, if $\{a, x_D(0)\} \in \mathbb{A}(\tau_i) \setminus \mathbb{A}(\tau_{i-1})$ then the attack will be detected if the defender chooses τ_{i-1} , but not if it chooses τ_i . We can remove the last column from Table I, because $a \notin \mathbb{A}(\tau_m)$ results in zero payoff for the attacker.

We define the maximum payoff for a given τ_i and a given attacker type ϕ as

$$\mathcal{I}_i^\phi := \max_{a, x_D(0)} \mathbf{1}_{\{\{a, x_D(0)\} \in \mathbb{A}(\tau_i)\}} f_\phi(a) = \max_{\{a, x_D(0)\} \in \mathbb{A}(\tau_i)} f_\phi(a).$$

Note that we also optimize over $x_D(0)$, which the attacker has normally no influence over. We do that to obtain the maximum possible payoff an attacker could achieve, which goes along with the scenario of the worst-case attacker and the interpretation that the attacker waits for the optimal time to attack. We can show the following for the maximum payoff.

Lemma 1: Consider an attacker of type ϕ . For a given τ_i , the maximum payoff exists and $\mathcal{I}_i^\phi := \max_{\{a, x_D(0)\} \in \mathbb{A}(\tau_i)} f_\phi(a) = \max_{\{a, x_D(0)\} \in \mathbb{A}(\tau_i) \setminus \mathbb{A}(\tau_{i-1})} f_\phi(a)$.

Proof: We begin by proving the first part of the lemma.

The first and second condition in Assumption 2 guarantee that $\mathbb{A}(\tau_i)$ is a compact set for a given τ_i (see Theorem 7.1 in [27]). Hence, by the extreme value theorem, we know that \mathcal{I}_i^ϕ always exists for a given τ_i and ϕ . The second part of the lemma readily follows from Assumption 8. ■

With the maximum payoff for a given τ_i , we can formulate a finite matrix game $\tilde{\mathcal{M}}_\phi = \langle \mathcal{P}, \tilde{\mathcal{A}}_\phi, \tilde{\mathcal{U}}_\phi \rangle$, where \mathcal{P} is defined as in \mathcal{M} , $\tilde{\mathcal{A}}_\phi = \{\tau_1, \dots, \tau_m\} \times \{\mathcal{I}_1^\phi, \dots, \mathcal{I}_m^\phi\}$, and $\tilde{\mathcal{U}}_\phi = (\frac{cE}{\tau_i} + \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi, \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi)$, where $i, j \in \{1, \dots, m\}$. Since both the attacker and the defender have finite actions sets in $\tilde{\mathcal{M}}_\phi$, we formulate $\tilde{\mathcal{M}}_\phi$ as the matrix game shown in Table II.

TABLE I
THE GAME \mathcal{M}_ϕ WITH DISJOINT SETS FOR THE ATTACKER'S ACTION

	$a_1 \in \mathbb{A}(\tau_1)$	$a_2 \in \mathbb{A}(\tau_2) \setminus \mathbb{A}(\tau_1)$	\dots	$a_m \in \mathbb{A}(\tau_m) \setminus \mathbb{A}(\tau_{m-1})$	$a \notin \mathbb{A}(\tau_m)$
τ_1	$\frac{c_F}{\tau_1} + f_\phi(a_1), f_\phi(a_1)$	$\frac{c_F}{\tau_1}, 0$	\dots	$\frac{c_F}{\tau_1}, 0$	$\frac{c_F}{\tau_1}, 0$
τ_2	$\frac{c_F}{\tau_2} + f_\phi(a_1), f_\phi(a_1)$	$\frac{c_F}{\tau_2} + f_\phi(a_2), f_\phi(a_2)$	\dots	$\frac{c_F}{\tau_2}, 0$	$\frac{c_F}{\tau_2}, 0$
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
τ_m	$\frac{c_F}{\tau_m} + f_\phi(a_1), f_\phi(a_1)$	$\frac{c_F}{\tau_m} + f_\phi(a_2), f_\phi(a_2)$	\dots	$\frac{c_F}{\tau_m} + f_\phi(a_m), f_\phi(a_m)$	$\frac{c_F}{\tau_m}, 0$

Further, we define the $m \times m$ matrix $\Omega(\phi)$ as the defender's cost matrix with elements $\Omega_{i,j}(\phi) = \frac{c_F}{\tau_i} + \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi$, and $\Upsilon(\phi)$ as the $m \times m$ matrix that has the attacker's payoff matrix with elements $\Upsilon_{i,j}(\phi) = \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi$.

Proposition 1: The finite game $\widetilde{\mathcal{M}}_\phi$ in Table II is strategically equivalent to the game \mathcal{M}_ϕ in Table I.

Proof: Since the attacker's objective is to maximize its payoff (2), it always chooses the trajectory that maximizes its payoff. From Lemma 1 we know there exists a maximum payoff trajectory for each of the columns in Table I. Hence, choosing the maximum payoff is strategically equivalent to choosing an attack trajectory that yields the maximum payoff. ■

TABLE II

FINITE $m \times m$ MATRIX GAME $\widetilde{\mathcal{M}}_\phi$ FOR THRESHOLD SWITCHING

	\mathcal{I}_1^ϕ	\mathcal{I}_2^ϕ	\dots	\mathcal{I}_m^ϕ
τ_1	$\frac{c_F}{\tau_1} + \mathcal{I}_1^\phi, \mathcal{I}_1^\phi$	$\frac{c_F}{\tau_1}, 0$	\dots	$\frac{c_F}{\tau_1}, 0$
τ_2	$\frac{c_F}{\tau_2} + \mathcal{I}_1^\phi, \mathcal{I}_1^\phi$	$\frac{c_F}{\tau_2} + \mathcal{I}_2^\phi, \mathcal{I}_2^\phi$	\dots	$\frac{c_F}{\tau_2}, 0$
\vdots	\vdots	\vdots	\ddots	\vdots
τ_m	$\frac{c_F}{\tau_m} + \mathcal{I}_1^\phi, \mathcal{I}_1^\phi$	$\frac{c_F}{\tau_m} + \mathcal{I}_2^\phi, \mathcal{I}_2^\phi$	\dots	$\frac{c_F}{\tau_m} + \mathcal{I}_m^\phi, \mathcal{I}_m^\phi$

By using the equivalent game in Table II, we can simplify the average cost functions, (4) and (5) of the game \mathcal{M}_ϕ , used in the Bayesian Nash equilibrium (6) to bilinear functions of p and q_ϕ , which helps us to solve both Problem 1 and Problem 2.

Corollary 1: In the strategically equivalent finite game $\widetilde{\mathcal{M}}_\phi$, the average cost of the defender is given by

$$\bar{c}_\phi(p, q_\phi) = p^T \Omega(\phi) q_\phi$$

and the average payoff of the attacker is given by

$$\bar{p}_\phi(p, q_\phi) = p^T \Upsilon(\phi) q_\phi$$

for each attacker type, where the i th element, $q_{\phi,i}$, of q_ϕ is the probability of choosing an attack trajectory that leads to the maximum payoff \mathcal{I}_i^ϕ .

Proof: Since the attacker has a finite set of actions in $\widetilde{\mathcal{M}}_\phi$, its mixed strategy q_ϕ is a discrete probability distribution. This leads directly to bilinear functions of p and q_ϕ for the average cost and payoff, respectively. ■

Remark 4: We would like to point out that the matrix game in Table II depends only on the worst-case impact of the attacker type. Hence, while the focus of this paper is on sensor attacks, one could also focus on attacks targeting other system components, such as actuator attacks, to determine the matrix game in Table II.

V. BAYESIAN NASH EQUILIBRIUM-BASED MTD

In the previous section, we showed that for any particular ϕ the corresponding game \mathcal{M}_ϕ is strategically equivalent to a finite matrix game $\widetilde{\mathcal{M}}_\phi$. This means that the Bayesian game \mathcal{M} is strategically equivalent to a finite Bayesian game, denoted by $\widetilde{\mathcal{M}}$ and its equilibria can be found by formulating an induced matrix game [28], obtained by combining the games $\widetilde{\mathcal{M}}_\phi$ with respect to the prior. In what follows, we first illustrate the procedure and we then use the induced game to give a necessary and sufficient condition for the existence of a Bayesian Nash equilibrium that is a moving target defense according to Definition 1.

A. An illustrative example

We start with an illustrative example, where each player has two actions to choose from. The attacker is assumed to have type 1 with probability π_1 and type 2 with probability $\pi_2 = 1 - \pi_1$. Hence, the finite game \mathcal{M}_ϕ corresponding to attacker type ϕ is as shown in Table III, where $\phi \in \{1, 2\}$.

TABLE III

2 × 2 MATRIX GAME EXAMPLE

	\mathcal{I}_1^ϕ	\mathcal{I}_2^ϕ
τ_1	$\frac{c_F}{\tau_1} + \mathcal{I}_1^\phi, \mathcal{I}_1^\phi$	$\frac{c_F}{\tau_1}, 0$
τ_2	$\frac{c_F}{\tau_2} + \mathcal{I}_1^\phi, \mathcal{I}_1^\phi$	$\frac{c_F}{\tau_2} + \mathcal{I}_2^\phi, \mathcal{I}_2^\phi$

To find the Bayesian Nash equilibrium, we can formulate an induced matrix game (see [28]) and find the Nash equilibria of that induced matrix game, which correspond to the Bayesian Nash equilibria of the original game \mathcal{M} . In the induced game, we combine the matrix games $\widetilde{\mathcal{M}}_1$ and $\widetilde{\mathcal{M}}_2$ into one game. The actions of the defender in the induced game are the same as in the games $\widetilde{\mathcal{M}}_1$ and $\widetilde{\mathcal{M}}_2$, that is, it can choose τ_1 or τ_2 as its action. The attacker, however, has the actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2$, where i_1 and i_2 are in $\{1, 2\}$. Hence, the attacker in the induced game is a combination of the attackers in the games $\widetilde{\mathcal{M}}_1$ and $\widetilde{\mathcal{M}}_2$ and its payoff is the expected value over the attacker types given the defender's prior $[\pi_1, \pi_2]$. The induced game is illustrated in Table IV. If the attacker chooses $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2$ in the induced game, then in $\widetilde{\mathcal{M}}_1$ the action of the attacker is its i_1 th action, i.e., $\mathcal{I}_{i_1}^1$, and in $\widetilde{\mathcal{M}}_2$ the action of the attacker is its i_2 th action, i.e., $\mathcal{I}_{i_2}^2$.

From Table IV, we observe that the defender prefers τ_2 over

TABLE IV
INDUCED MATRIX GAME FOR THE 2 × 2 GAME EXAMPLE

	$\mathcal{I}_1^1 \mathcal{I}_2^2$	$\mathcal{I}_1^1 \mathcal{I}_2^1$	$\mathcal{I}_2^1 \mathcal{I}_2^2$	$\mathcal{I}_2^1 \mathcal{I}_2^1$
τ_1	$\frac{c_F}{\tau_1} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2, \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2$	$\frac{c_F}{\tau_1} + \pi_1 \mathcal{I}_1^1, \pi_1 \mathcal{I}_1^1$	$\frac{c_F}{\tau_1} + \pi_2 \mathcal{I}_1^2, \pi_2 \mathcal{I}_1^2$	$\frac{c_F}{\tau_1}, 0$
τ_2	$\frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2, \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2$	$\frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_2^2, \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_2^2$	$\frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_1^1, \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_1^1$	$\frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_2^2, \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_2^2$

τ_1 if the following conditions hold

$$\begin{aligned} \frac{c_F}{\tau_1} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2 &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_1^2, \\ \frac{c_F}{\tau_1} + \pi_1 \mathcal{I}_1^1 &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_1^1 + \pi_2 \mathcal{I}_2^2, \\ \frac{c_F}{\tau_1} + \pi_2 \mathcal{I}_1^2 &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_1^1, \\ \frac{c_F}{\tau_1} &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^2 + \pi_2 \mathcal{I}_2^2, \end{aligned}$$

which are equivalent to

$$\begin{aligned} \frac{c_F}{\tau_1} &> \frac{c_F}{\tau_2}, \\ \frac{c_F}{\tau_1} &> \frac{c_F}{\tau_2} + \pi_2 \mathcal{I}_2^2, \\ \frac{c_F}{\tau_1} &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^1, \\ \frac{c_F}{\tau_1} &> \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^1 + \pi_2 \mathcal{I}_2^2. \end{aligned}$$

Note that the first inequality always holds, while the second and third inequalities hold if the last inequality holds.

Hence, we see that the defender prefers to play τ_2 over playing τ_1 if $\frac{c_F}{\tau_1} > \frac{c_F}{\tau_2} + \pi_1 \mathcal{I}_2^1 + \pi_2 \mathcal{I}_2^2$. In this case, the defender will play τ_2 independent of the attacker's action, such that the attacker will always play the action that maximizes its payoff, i.e., $\mathcal{I}_2^1 \mathcal{I}_2^2$. Therefore, there exists only a pure Bayesian Nash equilibrium strategy, which is *not* a moving target defense. For this simple example, we determined a sufficient condition for when an MTD does not exist. However, this is a simple example where the induced matrix game has a manageable size and we can calculate the Bayesian Nash equilibrium by hand. Assume now that the defender has $m > 2$ actions, while the attacker has m actions and $n_\phi > 1$ types. Then the induced matrix game is an $m \times m^{n_\phi}$ matrix game, whose size becomes unmanageable as either m , n_ϕ , or both, grow.

B. Best responses and strictly dominated actions

In the induced matrix game, the actions of the attacker are $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$, where $i_\phi \in \{1, \dots, m\}$ and $\phi \in \{1, \dots, n_\phi\}$, while the defender chooses τ_l . This leads to the attacker payoff

$$p^{\text{ind}}(\tau_l, \mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi}) = \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j$$

and the defender cost

$$c^{\text{ind}}(\tau_l, \mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi}) = \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j$$

in the induced matrix game, which we can use to characterize the best responses of the players.

Lemma 2: The best response of the attacker to a given action τ_l is

$$b_A(\tau_l) = \{\mathcal{I}_l^1 \mathcal{I}_l^2 \dots \mathcal{I}_l^{n_\phi}\}, \quad (7)$$

and the best response of the defender to a given action $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$ is

$$b_D(\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi}) = \{\tau_l | l \in \arg \min_{l \in \{1, \dots, m\}} \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j\}. \quad (8)$$

Proof: We start by investigating the best response of the attacker. For a given τ_l , the payoff $\mathcal{I}_{i_j}^j$ of type j with $i_j > l$ is zero since it is detected. Hence, $i_j \leq l$ needs to be fulfilled for each attacker type if it wants to get a payoff. Recall that $\mathcal{I}_i^j < \mathcal{I}_\eta^j$ for all $i < \eta$. Hence, to obtain the maximum payoff for a given τ_l , the attacker has a unique best response given by (7) in the induced matrix game.

For a given attack action, the defender's best response is to choose τ_l to minimize its cost, which results in the set of best responses given in (8). ■

While Lemma 2 provides the unique best response of the attacker, the defender might have several best responses. For example, it could be best to choose τ_m for a given attacker action to minimize the cost for false alarms. By choosing a smaller τ , even though it increases the false alarm cost, it makes more attacks detectable, which in turn decreases the attack cost. Hence, the best response depends on many factors.

Now that we looked at best responses, we will investigate when actions are strictly dominated in the induced matrix game. For the defender, an action τ_l strictly dominates τ_η if the cost for τ_l is strictly lower than the cost for τ_η for all possible actions of the attacker. Strict dominance of one attacker action over another can be defined similarly. With the best responses and the strictly dominated actions, we are then equipped to prove the existence of moving target defenses according to Definition 1. Recall that by eliminating strictly dominated actions, we do not change the set of the Nash equilibria of the induced game and, therefore, neither the Bayesian Nash equilibria of the original game \mathcal{M} .

Lemma 3: Assume that there exists $l > \eta$ such that

$$\frac{c_F}{\tau_\eta} > \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_l^j, \quad (9)$$

then we can eliminate the defender actions τ_1, \dots, τ_η and the attacker actions for which $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \dots \mathcal{I}_{i_{n_\phi}}^{n_\phi} \neq \mathcal{I}_{\eta+1}^1 \mathcal{I}_{\eta+1}^2 \dots \mathcal{I}_{\eta+1}^{n_\phi}$ holds for $i_j \in \{1, \dots, \eta + 1\}$, and obtain a reduced $(m - \eta) \times (m^{n_\phi} - (\eta + 1)^{n_\phi} + 1)$ induced matrix game.

Proof: We start with the strict dominance of the rows.

First note that τ_l strictly dominates τ_η if

$$\frac{c_F}{\tau_\eta} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq \eta\}} \pi_j \mathcal{I}_{i_j}^j > \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j \quad (10)$$

holds for all possible attacker actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$. We can split this condition into three cases:

- 1) All attacker types use attacks that are stealthy for τ_η , i.e., actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$ with $\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq \eta\}} = n_\phi$. Here, we see that the terms related to the attacker payoff on both sides of the inequality in (10) are the same, such that (10) simplifies to $\frac{c_F}{\tau_\eta} < \frac{c_F}{\tau_l}$. Hence, $\eta < l$ is necessary for the strict dominance of τ_l over τ_η .
- 2) All attacker types use attacks that will be detected for τ_l , i.e., actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$ with $\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j > l\}} = n_\phi$. Since the attacks are detected by τ_l , they will also be detected by τ_η such that the terms related to the attacker payoff on both sides of the inequality in (10) disappear. Therefore, (10) simplifies again to $\frac{c_F}{\tau_\eta} < \frac{c_F}{\tau_l}$.
- 3) There exists at least one attacker type that uses a strategy that is stealthy for τ_l but not for τ_η , i.e., actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi}$ with $\sum_{j=1}^{n_\phi} \mathbf{1}_{\{\eta+1 \leq i_j \leq l\}} > 0$.

Subtracting the terms related to the attacker payoff on the left side of (10) from the inequality itself leads to

$$\frac{c_F}{\tau_\eta} > \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{\eta+1 \leq i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j. \quad (11)$$

The first two cases show that we need $\tau_l > \tau_\eta$ or equivalently $l > \eta$ for τ_l to strictly dominate τ_η . For the third case, since

$$\frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_l^j \geq \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{\eta+1 \leq i_j \leq l\}} \pi_j \mathcal{I}_{i_j}^j$$

is always correct, (11) holds if (9) holds.

Further, since the following inequalities

$$\frac{c_F}{\tau_\nu} + \sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j \leq \nu\}} \pi_j \mathcal{I}_{i_j}^j > \frac{c_F}{\tau_\nu} > \frac{c_F}{\tau_\eta} > \frac{c_F}{\tau_l} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_l^j$$

hold for all $i_j \in \{1, \dots, m\}$ and $j \in \{1, \dots, n_\phi\}$, we see that if (9) holds τ_l does not only strictly dominate τ_η , but all τ_ν with $\nu \in \{1, \dots, \eta\}$.

Therefore, if (9) holds we can remove the first η rows of the induced matrix game. With the first η rows removed it follows that $\mathcal{I}_{\eta+1}^1 \mathcal{I}_{\eta+1}^2 \cdots \mathcal{I}_{\eta+1}^{n_\phi}$ strictly dominates all actions where the attacker of type j chooses $i_j \in \{1, \dots, \eta+1\}$ with $j \in \{1, \dots, n_\phi\}$ such that $\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j = \eta+1\}} \neq n_\phi$ holds, i.e., all actions $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi} \neq \mathcal{I}_{\eta+1}^1 \mathcal{I}_{\eta+1}^2 \cdots \mathcal{I}_{\eta+1}^{n_\phi}$ for $i_j \in \{1, \dots, \eta+1\}$. Hence, we can additionally remove the $(\eta+1)^{n_\phi} - 1$ columns corresponding to these actions to obtain the reduced matrix game. ■

C. Existence of a MTD strategy (Problem 1)

We now formulate a necessary and sufficient condition for the existence of a MTD strategy for the defender in the Bayesian game \mathcal{M} according to Definition 1.

Theorem 1: A moving target defense strategy exists if, and only if,

$$\frac{c_F}{\tau_{m-1}} \leq \frac{c_F}{\tau_m} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_m^j. \quad (12)$$

Proof: First, assume that (12) does not hold. Then (9) is fulfilled with $l = m$ and $\eta = m-1$. Hence, we can reduce the induced matrix game to a 1×1 matrix game (see Lemma 3), which has a pure strategy equilibrium. Therefore, the original Bayesian game, \mathcal{M} , has a unique and pure Bayesian Nash equilibrium, such that no MTD strategy exists.

Next, we show that there exists at least one Nash equilibrium where the defender plays an MTD strategy if (12) holds. Since the induced matrix game is a finite matrix game, we know that there exists at least one Nash equilibrium and equivalently at least one Bayesian Nash equilibrium for the original game, \mathcal{M} . For the Nash equilibrium to be a pure strategy Nash equilibrium, each player needs to play a best response to the other player's best response. Assume that $(\tau_l, \mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi})$ is a pure strategy Nash equilibrium, then according to Lemma 2 the following needs to be fulfilled

$$\begin{aligned} \mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi} &\in b_A(\tau_l), \\ \tau_l &\in b_D(\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_{n_\phi}}^{n_\phi}), \end{aligned}$$

i.e., each player's action is a best response to the other player's best response. Comparing the first equation with the attacker's best response (7), we see that in a pure Nash equilibrium $i_1 = i_2 = \dots = i_{n_\phi} = l$. With (8), we determine that the best response of the defender is

$$b_D(\mathcal{I}_l^1 \mathcal{I}_l^2 \cdots \mathcal{I}_l^{n_\phi}) = \begin{cases} \{\tau_{l-1}\} & \text{if } \frac{c_F}{\tau_{l-1}} < \frac{c_F}{\tau_m} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_l^j, \\ \{\tau_{l-1}, \tau_m\} & \text{if } \frac{c_F}{\tau_{l-1}} = \frac{c_F}{\tau_m} + \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_l^j, \\ \{\tau_m\} & \text{otherwise.} \end{cases} \quad (13)$$

To have a pure Nash equilibrium we need $l = m$. We observe that there cannot be a pure Nash equilibrium if (12) holds with inequality such that all equilibria are moving target defenses. However, if (12) holds with equality, the best response of the defender to $\mathcal{I}_m^1 \mathcal{I}_m^2 \cdots \mathcal{I}_m^{n_\phi}$ can be both τ_{m-1} and τ_m . Hence, in this case there exists a pure strategy Nash equilibrium in the induced matrix game. Next, we show that a moving target defense equilibrium strategy exists as well in this case.

First, note that if (12) holds with equality then τ_i is strictly dominated by τ_m for all $i \in \{1, \dots, m-2\}$, such that we can reduce the induced matrix game to a $2 \times (m^{n_\phi} - (m-1)^{n_\phi} + 1)$ matrix game. Further, from (13) we see that any distribution over τ_{m-1} and τ_m is a best response to the attack strategy $\mathcal{I}_m^1 \mathcal{I}_m^2 \cdots \mathcal{I}_m^{n_\phi}$. If we can show that $\mathcal{I}_m^1 \mathcal{I}_m^2 \cdots \mathcal{I}_m^{n_\phi}$ is also a best response to at least one distribution over τ_{m-1} and τ_m then we have found a Bayesian Nash equilibrium, which fulfills Definition 1. By multiplying the attacker's payoff matrix in the reduced matrix game from the left with the distribution over τ_{m-1} and τ_m , we determine that the expected payoff for playing $\mathcal{I}_{i_1}^1 \mathcal{I}_{i_2}^2 \cdots \mathcal{I}_{i_\phi}^{n_\phi}$ is

$$\sum_{j=1}^{n_\phi} \pi_j \mathbf{1}_{\{i_j \leq m-1\}} \mathcal{I}_{i_j}^j + p_m \sum_{j=1}^{n_\phi} \pi_j \mathbf{1}_{\{i_j = m\}} \mathcal{I}_m^j,$$

where $i_j \in \{1, \dots, m\}$ for all $j \in \{1, \dots, n_\phi\}$ and p_m is the probability of choosing τ_m . Note that $\mathcal{I}_m^1 \mathcal{I}_m^2 \dots \mathcal{I}_m^{n_\phi}$ is a best response to the mixed strategy of the defender, if the expected payoff for choosing $\mathcal{I}_m^1 \mathcal{I}_m^2 \dots \mathcal{I}_m^{n_\phi}$ is greater than or equal to all other expected payoffs the attacker could receive, i.e.,

$$\sum_{j=1}^{n_\phi} \pi_j \mathbf{1}_{\{i_j \leq m-1\}} \mathcal{I}_{i_j}^j + p_m \sum_{j=1}^{n_\phi} \pi_j \mathbf{1}_{\{i_j = m\}} \mathcal{I}_m^j \leq p_m \sum_{j=1}^{n_\phi} \pi_j \mathcal{I}_m^j$$

for all i_j such that $\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j = m\}} < n_\phi$. Hence, the attacker prefers to play $\mathcal{I}_m^1 \mathcal{I}_m^2 \dots \mathcal{I}_m^{n_\phi}$ if the defender chooses

$$p_m \in \left[\max_{i_1, \dots, i_{n_\phi}} \frac{\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j = m-1\}} \pi_j \mathcal{I}_{m-1}^j}{\sum_{j=1}^{n_\phi} \mathbf{1}_{\{i_j = m-1\}} \pi_j \mathcal{I}_m^j}, 1 \right),$$

where we used that $\mathcal{I}_{i_j} < \mathcal{I}_{m-1}$ for all $i_j < m-1$. This shows us that there are infinitely many Nash equilibria in the induced matrix game where the defender uses a MTD strategy according to Definition 1 if (12) holds with equality.

Therefore, we conclude that a moving target defense strategy according to Definition 1 exists if, and only if, (12) holds. ■

Remark 5: If $\pi_i = 1$ for some i , i.e., we have only one attacker type, the condition in Theorem 1 simplifies to the condition for the existence of a MTD from [17].

Remark 6: In case (12) holds with equality, we can introduce an attacker type that has zero payoff as mentioned in Remark 3 with a prior of $\epsilon > 0$ and subtract ϵ from one of the priors π_j . This will lead to (12) holding with inequality. Hence, (12) can always be turned into an inequality by an arbitrarily small change in the priors.

VI. COMPUTING A MTD STRATEGY (PROBLEM 2)

In this section, we look into computing a MTD strategy. First, we investigate the general case and formulate a linear program to compute MTD strategies. Second, we investigate the special case $n_\phi = 1$ and provide a closed-form solution for computing a MTD strategy.

A. General case

Finding Nash equilibria of a finite matrix game leads to a bilinear optimization problem as shown in [29]. For Bayesian Nash equilibria, we can adopt the optimization problem in [30] to obtain the following bilinear optimization problem

$$\begin{aligned} \min_{p, q_\phi, \bar{c}, \bar{p}(\phi)} \quad & p^T \left(\sum_{\phi=1}^{n_\phi} \pi_\phi (\Omega(\phi) - \Upsilon(\phi)) q_\phi \right) + \bar{c} - \sum_{\phi=1}^{n_\phi} \pi_\phi \bar{p}(\phi) \\ \text{s.t.} \quad & p^T \mathbf{1}_m = 1, \quad q_\phi^T \mathbf{1}_m = 1, \quad p \geq 0, \quad q_\phi \geq 0, \\ & \sum_{\phi=1}^{n_\phi} \pi_\phi \Omega(\phi) q_\phi \geq -\bar{c} \mathbf{1}_m, \\ & -\Upsilon^T(\phi) p \geq \bar{p}(\phi) \mathbf{1}_m, \quad \phi \in \{1, \dots, n_\phi\}. \end{aligned} \quad (14)$$

Recall that the elements of the matrices $\Omega(\phi)$ and $\Upsilon(\phi)$ are $\Omega_{i,j}(\phi) = \frac{\epsilon}{\tau_i} + \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi$ and $\Upsilon_{i,j}(\phi) = \mathbf{1}_{\{j \leq i\}} \mathcal{I}_j^\phi$, respectively. Here, q_ϕ is the mixed strategy for the attacker with type ϕ and $\bar{p}(\phi)$ is its average payoff, while p is the mixed strategy of the defender and \bar{c} is its average cost.

Proposition 2: Assume that the condition of Theorem 1 holds, and thus a MTD exists. A MTD strategy can then be computed by solving the linear program,

$$\begin{aligned} \min_{p, q_\phi, \bar{c}, \bar{p}(\phi)} \quad & p^T \gamma + \bar{c} - \sum_{\phi=1}^{n_\phi} \pi_\phi \bar{p}(\phi) \\ \text{s.t.} \quad & p^T \mathbf{1}_m = 1, \quad q_\phi^T \mathbf{1}_m = 1, \quad p \geq 0, \quad q_\phi \geq 0, \\ & \gamma + \sum_{\phi=1}^{n_\phi} \pi_\phi \Upsilon(\phi) q_\phi \geq -\bar{c} \mathbf{1}_m, \\ & -\Upsilon^T(\phi) p \geq \bar{p}(\phi) \mathbf{1}_m, \quad \phi \in \{1, \dots, n_\phi\}, \end{aligned} \quad (15)$$

where the l th element of the m -dimensional vector γ is $\frac{\epsilon}{\tau_l}$.

Proof: Due to the special structure of $\Omega(\phi)$, we note that the l th element of $\Omega(\phi) q_\phi$ equals $\frac{\epsilon}{\tau_l} + \sum_{j=1}^l q_{\phi,j} \mathcal{I}_j^\phi$, such that $\Omega(\phi) q_\phi = \gamma + \Upsilon(\phi) q_\phi$. Inserting that in the objective function and the constraints of (14) leads to the optimization problem in (15), where we further used that $\sum_{\phi=1}^{n_\phi} \pi_\phi = 1$. ■

The optimization problem in (15) is a convex linear program and therefore, we are guaranteed to find the global optimum. This has an advantage over directly solving (14), where we may get stuck in a local optimum.

Remark 7: Note that the computational complexity of (15) does not depend on the complexity of the system, for which we determine the MTD, because the number of optimization variables and constraints only depend on m and n_ϕ . However, determining the maximum attacker payoff to setup the matrix game in Table II for each attacker type depends on the complexity of the system, such as the number of inputs, outputs and states, and the expected attacker payoff $f_\phi(a)$.

B. Special case: $n_\phi = 1$

Next, we provide a closed-form solution to Problem 2, when the defender faces only one attacker type, i.e. $n_\phi = 1$, which is the problem we mentioned in Remark 5. For ease of notation, we will omit the superscript for the attacker type.

For $n_\phi = 1$ the matrix representation of the Bayesian Nash equilibrium definition in (6) simplifies to the definition of the Nash equilibrium

$$\begin{aligned} (p^*)^T \Omega q^* &\leq p^T \Omega q^* \quad \forall p \in \Delta_p, \\ (p^*)^T \Upsilon q^* &\geq (p^*)^T \Upsilon q \quad \forall q \in \Delta_q. \end{aligned} \quad (16)$$

Let \mathcal{Q} denote the support of the attacker's mixed strategy, i.e., if $i \in \mathcal{Q}$ then the attacker chooses \mathcal{I}_i with a nonzero probability $q_i > 0$ and if $i \notin \mathcal{Q}$ then $q_i = 0$. The support for the mixed strategy of the defender is defined in a similar way and is denoted by \mathcal{P} with probabilities p_i . In the following, we investigate one mixed strategy for the defender and one for the attacker and show how the support of the best response of the attacker, respectively defender, has to look like. We then use this to define a mixed strategy Nash equilibrium, which represents a MTD.

Lemma 4: If the attacker fixes i , $1 < i < m$, and uses the mixed strategy

$$q_j = \begin{cases} \frac{\epsilon}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right), & \text{if } j \in \{i+1, \dots, m\}, \\ 1 - \sum_{l=i+1}^m q_l, & \text{if } j = i, \\ 0, & \text{otherwise,} \end{cases} \quad (17)$$

where

$$0 \leq q_i < \max \left(1, \frac{c_F}{\mathcal{I}_i} \left(\frac{1}{\tau_{i-1}} - \frac{1}{\tau_i} \right) \right), \quad (18)$$

then $\mathcal{P} \subseteq \{i, \dots, m\}$ needs to hold for the support of the defender's best response.

Proof: First, note that $q_j > 0$ for $j \in \{i+1, \dots, m\}$, since $\tau_j > \tau_i$ if $j > i$. Further, if $q_i \in [0, 1)$ we see that the mixed strategy q given by (17) is a proper probability distribution.

Next, we look at the possible best responses of the defender to this strategy. The average cost of the defender is given by

$$\begin{aligned} p^T \Omega q &= p^T \gamma + p^T \Upsilon q \\ &= p_{1:i-1}^T \tilde{\gamma} + p_{i:m}^T \hat{\gamma} + p^T \begin{bmatrix} \Upsilon_{11} & \Upsilon_{12} \\ \Upsilon_{21} & \Upsilon_{22} \end{bmatrix} q \\ &= p_{1:i-1}^T \tilde{\gamma} + p_{i:m}^T \hat{\gamma} + [p_{1:i-1}^T \quad p_{i:m}^T] \begin{bmatrix} 0 \\ \Upsilon_{22} q_{i:m} \end{bmatrix} \\ &= p_{1:i-1}^T \tilde{\gamma} + p_{i:m}^T \hat{\gamma} + p_{i:m}^T \Upsilon_{22} q_{i:m}, \end{aligned}$$

where $\tilde{\gamma}^T = [\frac{c_F}{\tau_1}, \dots, \frac{c_F}{\tau_{i-1}}]$, $\hat{\gamma}^T = [\frac{c_F}{\tau_i}, \dots, \frac{c_F}{\tau_m}]$, and

$$\Upsilon_{22} = \begin{bmatrix} \mathcal{I}_i & 0 & \dots & 0 \\ \mathcal{I}_i & \mathcal{I}_{i+1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{I}_i & \mathcal{I}_{i+1} & \dots & \mathcal{I}_m \end{bmatrix}.$$

With that we can determine that

$$\Upsilon_{22} q_{i:m} = \begin{bmatrix} q_i \mathcal{I}_i \\ \sum_{j=i}^{i+1} q_j \mathcal{I}_j \\ \vdots \\ \sum_{j=i}^m q_j \mathcal{I}_j \end{bmatrix} = \begin{bmatrix} q_i \mathcal{I}_i + \frac{c_F}{\tau_i} - \frac{c_F}{\tau_{i+1}} \\ \vdots \\ q_i \mathcal{I}_i + \frac{c_F}{\tau_i} - \frac{c_F}{\tau_m} \end{bmatrix},$$

where we used (17) for the values of q_j for $j > i$. This leads to the following average cost of the defender

$$\begin{aligned} p^T \Omega q &= p_{1:i-1}^T \tilde{\gamma} + p_{i:m}^T \hat{\gamma} + p_{i:m}^T \Upsilon_{22} q_{i:m} \\ &= p_{1:i-1}^T \tilde{\gamma} + p_{i:m}^T \hat{\gamma} + \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \sum_{j=i}^m p_j - p_{i:m}^T \hat{\gamma} \\ &= p_{1:i-1}^T \tilde{\gamma} + \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \sum_{j=i}^m p_j. \end{aligned}$$

Now assume $\mathcal{P} \subseteq \{i, \dots, m\}$, then $p_{1:i-1} = 0$ and $\sum_{j=i}^m p_j = 1$, such that the average cost turns into $p^T \Omega q = q_i \mathcal{I}_i + \frac{c_F}{\tau_i}$, which shows us that the defender is indifferent among its action, as it obtains the same average cost no matter how the distribution $p_{i:m}$ is chosen.

Now let $\mathcal{P} \not\subseteq \{i, \dots, m\}$, then $\sum_{j=i}^m p_j = 1 - p_{1:i-1}^T \mathbf{1}_{i-1}$, such that the average cost becomes

$$p^T \Omega q = q_i \mathcal{I}_i + \frac{c_F}{\tau_i} + p_{1:i-1}^T \left(\tilde{\gamma} - \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \mathbf{1}_{i-1} \right).$$

The defender chooses $\mathcal{P} \not\subseteq \{i, \dots, m\}$ if, and only if,

$$\begin{aligned} q_i \mathcal{I}_i + \frac{c_F}{\tau_i} + p_{1:i-1}^T \left(\tilde{\gamma} - \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \mathbf{1}_{i-1} \right) &\leq q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \\ \Leftrightarrow p_{1:i-1}^T \left(\tilde{\gamma} - \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \mathbf{1}_{i-1} \right) &\leq 0. \end{aligned}$$

Since the elements of both $p_{1:i-1}$ and $\tilde{\gamma}$ are positive and $\frac{c_F}{\tau_1} > \frac{c_F}{\tau_2} > \dots > \frac{c_F}{\tau_{i-1}}$, we obtain the following necessary condition for choosing $\mathcal{P} \not\subseteq \{i, \dots, m\}$,

$$\frac{c_F}{\tau_{i-1}} \leq q_i \mathcal{I}_i + \frac{c_F}{\tau_i}.$$

Therefore, if

$$\frac{c_F}{\tau_{i-1}} > q_i \mathcal{I}_i + \frac{c_F}{\tau_i},$$

then $p_{1:i-1}^T \left(\tilde{\gamma} - \left(q_i \mathcal{I}_i + \frac{c_F}{\tau_i} \right) \mathbf{1}_{i-1} \right) > 0$ and the defender chooses $\mathcal{P} \subseteq \{i, \dots, m\}$ to minimize its cost. Reformulating this inequality gives us the upper bound

$$q_i < \frac{c_F}{\mathcal{I}_i} \left(\frac{1}{\tau_{i-1}} - \frac{1}{\tau_i} \right).$$

Note that if τ_{i-1} is strictly dominated by τ_i , this upper bound is larger than 1 and therefore automatically fulfilled, if $q_i \in [0, 1)$. However, if τ_{i-1} is *not* strictly dominated by τ_i then both

$$\frac{c_F}{\tau_{i-1}} > \frac{c_F}{\tau_i} \quad \text{and} \quad \frac{c_F}{\tau_{i-1}} \leq \frac{c_F}{\tau_i} + \mathcal{I}_i$$

hold, which means there exists $\rho_i \in (0, 1]$ such that

$$\frac{c_F}{\tau_{i-1}} = \frac{c_F}{\tau_i} + \rho_i \mathcal{I}_i$$

holds and we can determine ρ_i as

$$\rho_i = \frac{c_F}{\mathcal{I}_i} \left(\frac{1}{\tau_{i-1}} - \frac{1}{\tau_i} \right).$$

Hence, if q is chosen according to (17) such that q_i fulfills (18), the defender chooses $\mathcal{P} \subseteq \{i, \dots, m\}$. ■

Lemma 4 shows us the support of best responses for the defender to the attack strategy (17), however, it still leaves the open question how to choose i such that (18) is fulfilled.

Lemma 5: There exists a unique index $i = i^* \in (1, m-1)$ so that (17) is a proper probability distribution and (18) holds if $1 - \sum_{j=2}^m \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) < 0$ and $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) < 1$.

Proof: From (17), we obtain that $q_i = 1 - \sum_{j=i+1}^m \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right)$, which is strictly decreasing as i decreases. Since $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) < 1$ holds, we know that $q_i \in (0, 1)$ for $i = m-1$. Further, since $1 - \sum_{j=2}^m \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) < 0$ we know that $q_i < 0$ for $i = 1$. Hence, there exists an $i = i^* > 1$ such that $q_i \geq 0$ while for $i < i^*$ we have $q_i < 0$, such that the mixed strategy in (17) is not a proper probability distribution and therefore not a valid strategy. For $i = i^*$ we can, therefore, show that

$$\begin{aligned} 1 - \sum_{l=i^*}^m \frac{c_F}{\mathcal{I}_l} \left(\frac{1}{\tau_{l-1}} - \frac{1}{\tau_l} \right) &< 0 \\ \Leftrightarrow 1 - \sum_{l=i^*+1}^m \frac{c_F}{\mathcal{I}_l} \left(\frac{1}{\tau_{l-1}} - \frac{1}{\tau_l} \right) - \frac{c_F}{\mathcal{I}_{i^*}} \left(\frac{1}{\tau_{i^*-1}} - \frac{1}{\tau_{i^*}} \right) &< 0 \\ \Leftrightarrow q_{i^*} - \frac{c_F}{\mathcal{I}_{i^*}} \left(\frac{1}{\tau_{i^*-1}} - \frac{1}{\tau_{i^*}} \right) &< 0 \\ \Leftrightarrow q_{i^*} < \frac{c_F}{\mathcal{I}_{i^*}} \left(\frac{1}{\tau_{i^*-1}} - \frac{1}{\tau_{i^*}} \right). \end{aligned}$$

Hence, (18) holds for $i = i^*$. Now assume we choose $i = j > i^*$ such that $q_i \in [0, 1)$ if $i = j$ and also $q_i \geq 0$ if $i = j - 1$. Then we obtain that

$$\begin{aligned} 1 - \sum_{l=j}^m \frac{c_F}{\mathcal{I}_l} \left(\frac{1}{\tau_{l-1}} - \frac{1}{\tau_l} \right) &\geq 0 \\ \Leftrightarrow 1 - \sum_{l=j+1}^m \frac{c_F}{\mathcal{I}_l} \left(\frac{1}{\tau_{l-1}} - \frac{1}{\tau_l} \right) - \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) &\geq 0 \\ \Leftrightarrow q_j - \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) &\geq 0 \\ \Leftrightarrow q_j &\geq \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right). \end{aligned}$$

Hence, (18) does not hold for any $i \neq i^*$. Therefore, $i = i^*$ is the smallest index for which $q_i^* \in [0, 1)$ and the unique index for which (18) holds. ■

Similar to Lemma 4 we can find a mixed strategy p with support $\mathcal{P} = \{i, i + 1, \dots, m\}$ such that the best response of the attacker has support $\mathcal{Q} \subseteq \{i, i + 1, \dots, m\}$.

Lemma 6: If the defender uses the mixed strategy

$$p_j = \begin{cases} 1 - \frac{\mathcal{I}_i}{\mathcal{I}_{i+1}}, & \text{if } j = i, \\ \frac{\mathcal{I}_i}{\mathcal{I}_j} - \frac{\mathcal{I}_i}{\mathcal{I}_{j+1}}, & \text{if } j \in \{i + 1, \dots, m - 1\}, \\ \frac{\mathcal{I}_i}{\mathcal{I}_m}, & \text{if } j = m, \\ 0, & \text{otherwise,} \end{cases} \quad (19)$$

then the support of the attacker's best response needs to satisfy $\mathcal{Q} \subseteq \{i, i + 1, \dots, m\}$.

Proof: First note that since $\mathcal{I}_j > \mathcal{I}_i$ if $j > i$, we see that each $p_i \in (0, 1)$. Furthermore, we can verify that $\sum_{j=1}^m p_j = 1$. Hence, the mixed strategy described by (19) is a proper probability distribution.

Next, we look at the possible best responses of the attacker to this strategy. The average cost of the attacker is

$$\begin{aligned} p^T \Upsilon q &= p^T \begin{bmatrix} \Upsilon_{11} & \Upsilon_{12} \\ \Upsilon_{21} & \Upsilon_{22} \end{bmatrix} q \\ &= p_{i:m}^T \Upsilon_{21} q_{1:i-1} + p_{i:m}^T \Upsilon_{22} q_{i:m} \\ &= \sum_{j=1}^{i-1} q_j \mathcal{I}_j + p_{i:m}^T \Upsilon_{22} q_{i:m}, \end{aligned}$$

where we used that $\Upsilon_{12} = 0$ and $\Upsilon_{21} = 1_{i-1}[\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{i-1}]$. Due to the chosen p , we obtain that $p_{i:m}^T \Upsilon_{22} = 1_{m-i+1}^T \mathcal{I}_i$, which results in the following average cost

$$\begin{aligned} p^T \Upsilon q &= \sum_{j=1}^{i-1} q_j \mathcal{I}_j + \mathcal{I}_i \sum_{l=i}^m q_l \\ &= \sum_{j=1}^{i-1} q_j \mathcal{I}_j + \mathcal{I}_i \left(1 - \sum_{j=1}^{i-1} q_j \right) \\ &= \mathcal{I}_i + \sum_{j=1}^{i-1} q_j (\mathcal{I}_j - \mathcal{I}_i) \leq \mathcal{I}_i. \end{aligned}$$

The upper bound comes from the fact that $\mathcal{I}_j < \mathcal{I}_i$ if $j \in \{1, \dots, i - 1\}$ (see Assumption 8). Hence, we see that the

best response of the attacker to the defender's strategy p is any mixed strategy q with support $\mathcal{Q} \subseteq \{i, i + 1, \dots, m\}$. ■

One notable difference between the results given in Lemma 4 and Lemma 6 is that the defender's mixed strategy p given by (19) is valid for all i , while q given by (17) has the additional constraint (18). However, Lemma 5 shows us that under a certain condition there exists a unique index i for which (18) holds. Next, we show that for a specific choice of i the strategies (17) and (19) form a mixed strategy Nash equilibrium.

Theorem 2: Let $i = i^* \in [1, m - 1]$ be the *smallest* index for which $q_i^* \in [0, 1)$ in (18) holds such that q^* is a proper probability distribution. The mixed strategies p^* and q^* given by (19) and (17), respectively, form a mixed strategy Nash equilibrium such that p^* is a MTD if, and only if, $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) \leq 1$.

Proof: We begin by noting that if $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) > 1$, then $q_m > 1$ such that (17) is not a proper probability distribution. Next, we show that if $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) \leq 1$ then $i^* \in [1, m - 1]$ such that $q_{i^*} \in [0, 1)$ exists. For this, we need to consider three different cases. In the first case, we assume that $1 - \sum_{j=2}^m \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) < 0$ and $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) < 1$. Then Lemma 5 shows us that a unique $i^* \in (1, m - 1)$ exists, for which (18) holds and i^* is also the smallest index for which $q_i \in [0, 1)$. In the second case, we assume that $1 - \sum_{j=2}^m \frac{c_F}{\mathcal{I}_j} \left(\frac{1}{\tau_{j-1}} - \frac{1}{\tau_j} \right) \geq 0$ and $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) < 1$. Then $i^* = 1$ guarantees that $q_i \in [0, 1)$. Furthermore, $i^* = 1$ is also the smallest index in this case for which $q_i \in [0, 1)$. In the third case, we assume that $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) = 1$, which shows us that $i^* = m - 1$ is the smallest index in this case for which $q_i \in [0, 1)$. Hence, there exists a unique index $i^* \in [1, m - 1]$, which is the smallest index such that $q_i \in [0, 1)$ if, and only, if $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) \leq 1$.

Next, if we use q^* with $i = i^*$, we see that the support of the defender's best response needs to fulfill $\mathcal{P} \subseteq \{i^*, i^* + 1, \dots, m\}$, which is fulfilled when p^* is used. Therefore, p^* is a best response to q^* .

Finally, if we use p^* with $i = i^*$, we see that the support of the defender's best response needs to fulfill $\mathcal{Q} \subseteq \{i^*, i^* + 1, \dots, m\}$, which is fulfilled when q^* is used. Therefore, q^* is a best response to p^* . Hence, p^* and q^* form a mixed strategy Nash equilibrium and p^* is a MTD according to Definition 1. ■

Theorem 2 presents one optimal solution to the optimization problem (15) when $n_\phi = 1$. Our numerical experiments in Section VII-B show that the optimal solution obtained by solving (15) coincides with the equilibrium proposed in Theorem 2.

Remark 8: Note that if $\frac{c_F}{\mathcal{I}_m} \left(\frac{1}{\tau_{m-1}} - \frac{1}{\tau_m} \right) = 1$, then we have $i^* = m - 1$ and $q_{i^*} = 0$ such that the attacker plays a pure strategy. This means there are at least two Nash equilibria, one given by p^* and \mathcal{I}_m , and one given by τ_m and \mathcal{I}_m . In this case, our matrix game is degenerate but there still exists an MTD according to Definition 1.

Remark 9: As already mentioned in Remark 3, it is often reasonable to include an attacker type that has zero payoff for

all trajectories a such that the defender does also consider the case without an attacker when choosing the threshold. For the closed-form solution presented in this section, we are able to replace \mathcal{I}_j by $\pi_1 \mathcal{I}_j$ to take the attacker type with zero payoff into account, where $\pi_1 \in (0, 1]$ is the probability that the attack is happening. Interestingly, this modification does not change the defender's equilibrium strategy (19). However, it does lead to a larger i^* in Lemma 5 (see Section VII-B). This is consistent with our intuition since the attack will be less likely to happen and the defender focuses on increasing the mean time between false alarms by choosing larger thresholds.

VII. NUMERICAL EVALUATION

For the numerical evaluation, we look at a four-tank system [31], which we linearize around the input voltage of 6 V and discretize with a sampling time of 0.5 s. We further assume that $w(k) \sim \mathcal{N}(0, 0.1I_4)$ and $v(k) \sim \mathcal{N}(0, 0.01I_2)$ and an LQG controller is used, where the LQR cost matrix for the states is I_4 and the cost matrix for the controller input is I_2 , i.e., the controller input $u(k) = -Kx(k)$ minimizes the cost function $\sum_{k=0}^{\infty} x(k)^T x(k) + u(k)^T u(k)$. The attack length is chosen to be $N = 1000$ time steps. For anomaly detection, a χ^2 detector is used such that

$$y_D(k+1) = r(k)^T r(k).$$

A. Bayesian Nash equilibrium

In this part, we solve the optimization problem (15) to find the equilibrium moving target defense for the defender. Note that we choose the defender's set of actions, the attacker type payoff functions and the factor c_F in the defender's objective function for illustrative purposes of the presented MTD framework. In practice, the defender needs to choose c_F according to its cost for false alarm and the payoff functions of the attacker types could be the result of a risk assessment as discussed in Section II-E.

The defender considers six thresholds, which correspond to the following average times between false alarms,

$$\tau \in \{10, 10^2, 10^3, 10^4, 10^5, 10^6\}. \quad (20)$$

These values are chosen to cover a wide range of average times between false alarms. Further, we use $c_F = 43200$ as the cost factor for false alarms.

The defender faces $n_\phi = 5$ attacker types. We further assume that the attack starts at $k = 0$. The first attacker type is an attacker with zero payoff, i.e., $f_1(a) = 0$ for all a . This type represents the case where there is actually no attacker present in the system and the defender only has to consider the cost induced by the false alarms. For the other attacker types, we use the average value of the plant's state at the end of the attack, i.e., $\bar{x} = \mathbb{E}\{x(N)\}$, to define the payoff. For attacker type $\phi \in \{2, \dots, 5\}$, we use $f_\phi(a) = |\bar{x}_{\phi-1}|^2$ as the payoff function. Thus, attacker type 2 attacks the water level in tank 1, attacker type 3 attacks the water level in tank 2 and so on. Since a χ^2 detector is used, we can use the results of Proposition 3 in [15] to determine the attack impact for each attacker type for a given τ .




Fig. 2. The plot shows the MTD of the defender for three different priors of an attack happening, where the horizontal axis are the defender actions and the vertical axis shows the probability of choosing the respective action.

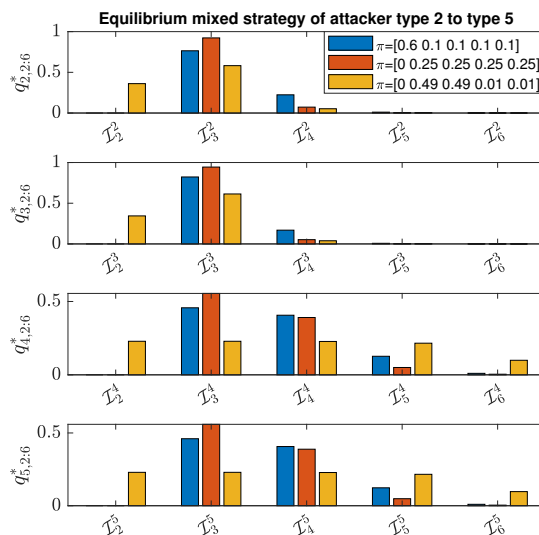


Fig. 3. For three different priors, the equilibrium mixed strategies of attacker type 2 to type 5 are shown in each of the subfigures. The horizontal axis shows the attacker actions while the vertical axis shows the probability of choosing the respective action.

We consider three different scenarios that differ in terms of their priors π_ϕ . In the first scenario, the operator assumes that it is more likely that there is no attack than that there is an attack, and thus $\pi_1 = 0.6$ and $\pi_\phi = 0.1$ for $\phi \in \{2, \dots, 5\}$. In the second scenario, the operator assumes that there is always an attack but we want to investigate how the defender's strategy changes when the attacks are still equally likely, i.e., $\pi_1 = 0$ and $\pi_\phi = 0.25$ for $\phi \in \{2, \dots, 5\}$. In the third scenario, there is also always an attack but this time the attacker is assumed to most likely attack the first and second state of the plant, i.e., $\pi_1 = 0$, $\pi_2 = \pi_3 = 0.49$, and $\pi_4 = \pi_5 = 0.01$.

Figure 2 shows the equilibrium MTD of the defender for the three different scenarios. Figure 3 shows the equilibrium mixed strategies for attacker type 2 to attacker type 5. Using Lemma 3, we can determine that for the first two scenarios the defender's actions τ_1 and τ_2 are strictly dominated and, therefore, will not be used in the Bayesian Nash equilibrium. In the third scenario, only τ_1 is strictly dominated and τ_2 is used in the Bayesian Nash equilibrium. This means that the cost for an attack that is stealthy for τ_1 is negligible compared to the cost for false alarms. Further, this also means none of the attacker types will use the attack action corresponding to these thresholds in the respective scenarios, because the attacker

wants to maximize its payoff. Hence, $p_1^* = q_{\phi,1}^* = 0$ for all ϕ and these values are not depicted in Figure 2 and Figure 3 for the sake of simplicity. Furthermore, since the attacker type 1 obtains always zero payoff it does not influence the objective value of the optimization problem in (15). Therefore, we can arbitrarily choose $q_{1,1:6}^* = \frac{1}{6}1_6$ as the equilibrium mixed strategy of attacker type 1 for all three scenarios. This is not shown in the figures for the sake of simplicity.

From Figure 2, we can observe that the defender chooses τ_6 with the highest probability in the MTD, while the attacker puts most of the probability weight on the attacks with a lower payoff than the payoff corresponding to τ_6 . Since the attack will not receive any payoff if it is detected, this observation is reasonable and also shows that the proposed moving target defense is effectively limiting the attacker payoff. This can be observed especially in the third scenario, where attacker type 2 and type 3 are more likely than attacker type 4 and type 5. In addition to that, the payoffs of attacker type 2 and type 3 are larger than for type 4 and type 5. Hence, a larger attacker payoff is more likely than in the second scenario. Therefore, in the third scenario the defender chooses τ_2 with a non-zero probability to force the attacker to remain undetected and receive a lower payoff. So we see that in this case having more false alarms outweighs the cost of having an undetectable attack.

We can also observe from Figure 2 that the defender has the same MTD for the first two scenarios. A reason for this is that in (15) the defender's constraints will not be affected by the attacker type with zero payoff. Hence, the constraint set for choosing p is the same in the first and second scenario due to the uniform prior across attacker types 2 to 5. For the different attacker types, it is interesting to see that although the defender has the same MTD in the first and second scenario, the attacker types' mixed strategies do change. We, further, observe that the mixed strategies q_2^* and q_3^* for attacker type 2 and type 3, respectively, are very close and the mixed strategies q_4^* and q_5^* are close as well, in all three investigated scenarios. For example, $\|q_2^* - q_3^*\|_\infty$ for the first, second, and third scenario is 0.0672, 0.0199, and 0.0131, respectively, while $\|q_4^* - q_5^*\|_\infty$ for the first, second, and third scenario is 0.0039, 0.0024, and 0.0053, respectively.

Finally, we look at an interesting property of the MTD obtained by solving (15). The MTD strategy in the first and second scenario, where all attacker types with a nonzero impact have the same prior, is

$$p^* = [0 \ 0 \ 0.25 \ 0.15 \ 0.1 \ 0.5]^T \quad (21)$$

while in the third scenario the MTD is

$$p^* = [0 \ 0.3333 \ 0.1667 \ 0.1 \ 0.0667 \ 0.3333]^T. \quad (22)$$

Interestingly, both (21) and (22) have the structure of the proposed MTD in Theorem 2, although Theorem 2 is only for the case where the defender faces one specific attacker type. This can be explained with the structure of the payoff for each attacker type. We determine that $\mathcal{I}_j^\phi = \alpha_\phi J_D(\tau_j)$ for $\phi \in \{2, \dots, 5\}$ (Proposition 3 of [15]). Hence, the payoff is the detector threshold times an attacker type specific constant

α_ϕ . Therefore, we have that

$$\frac{\mathcal{I}_i^\phi}{\mathcal{I}_j^\phi} = \frac{J_D(\tau_i)}{J_D(\tau_j)}.$$

Since this ratio is independent of the attacker type ϕ , the closed-form solution of the defender's equilibrium MTD for each of the attacker types is the same.

B. Closed-form solution

We finally evaluate the closed-form solution, where we only consider one attacker type. The attacker that we consider uses $f(a) = \|\bar{x}\|_\infty^2$ as its attack objective, which can be seen as the attacker type with the largest payoff among the four attacker types with a non-zero payoff in the previous section.

We consider a defender that can choose from the six different mean time between false alarms in (20). Lemma 3 with $n_\phi = 1$ shows us that $\tau_1 = 10$ is strictly dominated by the other strategies and is therefore not used in the Nash equilibrium. Setting $i = 2$, i.e., using the set of all strictly dominating strategies, we determine that $q_2 = 0.4748 < 1$ such that the condition on i in Theorem 2 holds with $i^* = 2$. Using Theorem 2, we determine that

$$p^* = [0 \ 0.3333 \ 0.1667 \ 0.1 \ 0.0667 \ 0.3333]^T \quad (23)$$

and

$$q^* = [0 \ 0.4748 \ 0.4857 \ 0.0364 \ 0.0029 \ 0.0002]^T. \quad (24)$$

These strategies are also obtained with the optimization problem (15). So we see that our closed-form solution coincides with the solution of the linear program when $n_\phi = 1$. Furthermore, the MTD (23) coincides with (22), while (24) does not coincide with any of the attacker type distributions obtained in the previous section.

Next, we evaluate the effect of the size m of the set of average times between false alarms the defender can choose from on the smallest index i in the support set of the defender and its probability q_i^* . We do so by considering values of m from 1 to 100 and set $\tau_j = 10 + (j - 1)50$, where $j \in [1, m]$. The goal here is to show how the index i changes as the size of the set of τ varies. First, for a given m the Nash equilibrium obtained from (15) coincides with the closed-form solution in Theorem 2 for all investigated m . Now let us analyse how the index i and with it q_i^* changes as m increases. In Figure 4, we see the index i of the mixed Nash equilibrium in the upper plot and q_i^* in the lower plot over m . Furthermore, we analyse two different cases, one where the attacker is always present and one where the prior of the attacker being present is 0.2 and the prior of an attacker with zero payoff is 0.8 (see Remark 9). First, note the evolution of q_i^* over m . Every time i changes, i.e., the smallest τ used in the Nash equilibrium changes for a given m , q_i^* jumps from a value close to zero back up to a larger value just to decrease to zero again as m increases until the next jump. This is consistent with $q_i^* = 1 - \sum_{l=i+1}^m \frac{c_F}{\mathcal{I}_l} \left(\frac{1}{\tau_{l-1}} - \frac{1}{\tau_l} \right)$, which is decreasing as m increases. Take, for example, the interval $m \in [2, 6]$




Fig. 4. The upper plot shows how the smallest index $i = i^*$, for which (18) holds, changes when m increases. The lower plot shows the trajectory of q_i^* in (17) over m . The solid lines represent the case with an attacker always being present, while the dashed lines represent the case where the chance of the attacker being present is 20 %.

for the case where there is always an attacker (solid line in Figure 4). In this interval $i = 2$ and we see that q_i^* has its maximum 1 at $m = 2$ and it decays to $q_i^* = 0.02456$ at $m = 6$. However, if we choose $i = 2$ for $m = 7$, then $q_i^* = -0.0148$ is smaller than zero and, therefore, using τ_2 in the Nash equilibrium does not lead to a proper probability distribution when $m = 6$. Hence, i needs to jump from 2 to 3 to guarantee that q^* is a probability distribution with $q_i^* \in [0, 1)$. Furthermore, if $m = 100$ then only the action τ_1 is strictly dominated. However, in the Nash equilibrium the smallest τ used has index $i = 3$. Therefore, we see that even if an action is not strictly dominated the action is not necessarily used in the Nash equilibrium. With our results in Theorem 2 we understand the reasons behind that in the game presented here. Similar observations are made for the case when the prior of the attacker with a non-zero payoff is 0.2 (dashed lines in Figure 4). Furthermore, we see since the attacker is only present with a chance of 20% more strategies of the defender are strictly dominated. This observation is in line with our intuition, since the attack is less likely to happen and the defender can choose larger thresholds to avoid false alarms without fearing a larger impact.

VIII. CONCLUSIONS

In this paper, we presented a moving target defense strategy against stealthy sensor attacks. To find the moving target defense, we formulated a game where the defender periodically switches the detector threshold at random and the attacker has access to all sensor measurements. While the attacker wants to maximize its payoff the defender wants to minimize its cost consisting of the cost for false alarms and the cost induced by the attacker's payoff. However, the defender is not certain about which attacker it faces and only knows the prior of the attacker's possible type. We analyzed one period of this periodic switching game and showed that for one period we can find a strategically equivalent matrix game. For this matrix game, we use the Bayesian Nash equilibrium to determine the equilibrium MTD strategy and presented a

necessary and sufficient condition for when a MTD for the defender exists. Furthermore, we showed that the MTD can be found by solving a linear program. For the case, where the defender only faces one attacker/knows the attacker type exactly, we presented a closed-form solution for the moving target defense. In the numerical evaluation, we saw how the thresholds used by the defender depend on the prior of an attack happening. The mere threat by the defender of randomly choosing a lower threshold with a low probability forces each attacker type to choose attacks with a lower impact which are stealthy for even small thresholds. The reason for that is that the attacker will not be able to get a payoff when it is detected.

If we believe that the attacker might observe the defender's switching pattern before attacking, the attacker could have a larger payoff. Therefore, one direction of future work is not to investigate the Bayesian Nash equilibrium, but the Bayesian Stackelberg equilibrium, where it is assumed that the attacker observes the defender first. Another direction of future work would be to investigate the repeated game setting. Furthermore, an in-depth analysis of the optimal choice of the set $\{\tau_1, \dots, \tau_m\}$ is also an avenue of future work. Finally, applying the proposed MTD to more complex processes than the four-tank system is another important step in validating the applicability of the MTD.

REFERENCES

- [1] R. M. Lee, M. J. Assante, and T. Conway, "Analysis of the cyber attack on the Ukrainian power grid. defense use case," *E-ISAC*, 2016.
- [2] —, "German steel mill cyber attack," *E-ISAC*, 2014.
- [3] S. M. Dibaji, M. Pirani, D. B. Flamholz, A. M. Annaswamy, K. H. Johansson, and A. Chakraborty, "A systems and control perspective of CPS security," *Annual Reviews in Control*, vol. 47, pp. 394–411, 2019.
- [4] M. S. Chong, H. Sandberg, and A. M. H. Teixeira, "A tutorial introduction to security and privacy for cyber-physical systems," in *18th European Control Conference (ECC)*, June 2019, pp. 968–978.
- [5] S. R. Etesami and T. Başar, "Dynamic games in cyber-physical security: An overview," *Dynamic Games and Applications*, Jan. 2019.
- [6] S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang, *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*, 1st ed. Springer Publishing Company, Incorporated, 2011.
- [7] J. Giraldo, A. Cardenas, and R. G. Sanfelice, "A moving target defense to detect stealthy attacks in cyber-physical systems," in *American Control Conference (ACC)*, July 2019, pp. 391–396.
- [8] P. Hespanhol, M. Porter, R. Vasudevan, and A. Aswani, "Statistical watermarking for networked control systems," in *2018 Annual American Control Conference (ACC)*, 2018, pp. 5467–5472.
- [9] Z. Zhang, R. Deng, D. K. Y. Yau, P. Cheng, and J. Chen, "Analysis of moving target defense against false data injection attacks on power grid," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 2320–2335, 2020.
- [10] A. Kanellopoulos and K. G. Vamvoudakis, "A moving target defense control framework for cyber-physical systems," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 1029–1043, 2020.
- [11] P. Griffioen, S. Weerakkody, and B. Sinopoli, "A moving target defense for securing cyber-physical systems," *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 2016–2031, 2021.
- [12] C. Murguia and J. Ruths, "CUSUM and chi-squared attack detection of compromised sensors," in *IEEE Conference on Control Applications (CCA)*, Sep. 2016, pp. 474–480.
- [13] D. I. Urbina, J. A. Giraldo, A. A. Cardenas, N. O. Tippenhauer, J. Valente, M. Faisal, J. Ruths, R. Candell, and H. Sandberg, "Limiting the impact of stealthy attacks on industrial control systems," in *ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16. New York, NY, USA: ACM, 2016, pp. 1092–1105.
- [14] A. Ghafouri, W. Abbas, A. Laszka, Y. Vorobeychik, and X. Koutsoukos, "Optimal thresholds for anomaly-based intrusion detection in dynamical environments," in *Decision and Game Theory for Security*. Springer International Publishing, 2016, pp. 415–434.

- [15] D. Umsonst and H. Sandberg, "A game-theoretic approach for choosing a detector tuning under stealthy sensor data attacks," in *IEEE Conference on Decision and Control (CDC)*, Dec. 2018, pp. 5975–5981.
- [16] L. Niu and A. Clark, "A framework for joint attack detection and control under false data injection," in *Decision and Game Theory for Security*, T. Alpcan, Y. Vorobeychik, J. S. Baras, and G. Dán, Eds. Cham: Springer International Publishing, 2019, pp. 352–363.
- [17] D. Umsonst, S. Saritaş, and H. Sandberg, "A Nash equilibrium-based moving target defense against stealthy sensor attacks," in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 3772–3778.
- [18] M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and Fault-Tolerant Control*, 3rd ed. Springer Berlin, Heidelberg, 2016.
- [19] S. X. Ding, *Model-Based Fault Diagnosis Techniques*, 2nd ed. Springer London, 2013.
- [20] H. P. Benson, *Concave Minimization: Theory, Applications and Algorithms*. Boston, MA: Springer US, 1995, pp. 43–148.
- [21] C. E. Shannon, "Communication theory of secrecy systems," *The Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, 1949.
- [22] D. Umsonst and H. Sandberg, "Experimental evaluation of sensor attacks and defense mechanisms in feedback systems," *Control Engineering Practice*, vol. 124, p. 105178, 2022.
- [23] M. O. Sayin and T. Başar, "Secure sensor design for cyber-physical systems against advanced persistent threats," in *Decision and Game Theory for Security*. Cham: Springer International Publishing, 2017, pp. 91–111.
- [24] R. M. Blank and P. D. Gallagher, "Guide for conducting risk assessments," *National Institute of Standards and Technology*, 2012.
- [25] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, "Deployed armor protection: The application of a game theoretic model for security at the los angeles international airport," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track*, ser. AAMAS '08. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2008, p. 125–132.
- [26] S. Sengupta, A. Chowdhary, D. Huang, and S. Kambhampati, "Moving target defense for the placement of intrusion detection systems in the cloud," in *Decision and Game Theory for Security*, L. Bushnell, R. Poovendran, and T. Başar, Eds. Cham: Springer International Publishing, 2018, pp. 326–345.
- [27] D. Umsonst, "Tuning of anomaly detectors in the presence of sensor attacks," KTH, Automatic Control, Tech. Rep., 2019.
- [28] K. Leyton-Brown and Y. Shoham, "Essentials of game theory: A concise multidisciplinary introduction," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 2, no. 1, pp. 1–88, 2008.
- [29] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory, 2nd Edition*. Society for Industrial and Applied Mathematics, 1998.
- [30] L. Huang and Q. Zhu, "Analysis and computation of adaptive defense strategies against advanced persistent threats for cyber-physical systems," in *Decision and Game Theory for Security*, L. Bushnell, R. Poovendran, and T. Başar, Eds. Springer International Publishing, 2018, pp. 205–226.
- [31] K. H. Johansson, "The quadruple-tank process: a multivariable laboratory process with an adjustable zero," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 3, pp. 456–465, May 2000.



David Umsonst is a doctoral student at KTH Royal Institute of Technology, Stockholm, Sweden. He received his M.Sc. degree in Electrical Engineering and Information Technology from the Technical University Darmstadt, Germany, in 2016. David received a M.Sc. degree in Engineering from KTH Royal Institute of Technology in 2016 as well due to the participation in a Double Degree program between the Technical University Darmstadt and KTH Royal Institute of Technology. In 2021, he received his Ph.D.

degree in automatic control at the Division of Decision and Control Systems at KTH Royal Institute of Technology, Stockholm, Sweden. His current research interests include security of cyber-physical control systems, anomaly detectors for control systems, optimization, robotics, and collaborative SLAM.



Serkan Saritaş is an Assistant Professor at the Electronics Engineering Department at Middle East Technical University (METU). He received his B.Sc. degree in Electrical and Electronics Engineering, M.S. degree in Computer Engineering, and Ph.D. degree in Electrical and Electronics Engineering from Bilkent University, Turkey, in 2010, 2013, and 2018, respectively. He was a visiting research student at the Department of Mathematics and Statistics at Queen's University in Kingston, Canada, for five months in 2015.

Between 2018 and 2021, he was a Postdoctoral Researcher jointly at the Division of Decision and Control Systems, the Division of Network and Systems Engineering, and the Division of Information Science and Engineering at KTH Royal Institute of Technology in Sweden. His main research interests include information security, cyber-physical systems, networked control systems, game theory, communication systems, and information theory.



György Dán (M'07-SM'17) is a Professor at KTH Royal Institute of Technology, Stockholm, Sweden. He received the M.Sc. in computer engineering from the Budapest University of Technology and Economics, Hungary in 1999, the M.Sc. in business administration from the Corvinus University of Budapest, Hungary in 2003, and the Ph.D. in Telecommunications from KTH in 2006. He worked as a consultant in the field of access networks, streaming media and videoconferencing 1999-2001. He was a visiting researcher at the Swedish Institute of Computer Science in 2008, a Fulbright research scholar at University of Illinois at Urbana-Champaign in 2012-2013, and an invited professor at EPFL in 2014-2015. He has been an area editor of Computer Communications 2014-2021 and of IEEE Trans. on Mobile Computing 2019-2023. His research interests include the design and analysis of content management and computing systems, game theoretical models of networked systems, and cyber-physical system security and resilience.



Henrik Sandberg is a Professor at the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. He received an M.Sc. degree in engineering physics and a Ph.D. in automatic control from Lund University, Lund, Sweden, in 1999 and 2004, respectively. From 2005 to 2007, he was a Post-Doctoral Scholar at the California Institute of Technology, Pasadena, USA. In 2013, he was a Visiting Scholar at the Laboratory for Information and Decision Systems (LIDS) at MIT, Cambridge, USA. He has also held visiting appointments at the Australian National University and the University of Melbourne, Australia. His current research interests include security of cyber-physical systems, power systems, model reduction, and fundamental limitations in control. Dr. Sandberg received the Best Student Paper Award from the IEEE Conference on Decision and Control in 2004, an Ingvar Carlsson Award from the Swedish Foundation for Strategic Research in 2007, and a Consolidator Grant from the Swedish Research Council in 2016. He has served on the editorial boards of IEEE Transactions on Automatic Control and the IFAC Journal Automatica. He is a Fellow of the IEEE.