2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

On a Notion of Approximate Opacity for Discrete-Time Stochastic Control Systems*

Siyuan Liu, Xiang Yin, and Majid Zamani

Abstract— **This paper investigates a confidentiality property called opacity for discrete-time stochastic control systems. In order to quantitatively evaluate the security guarantee, a notion of approximate initial-state opacity for stochastic control systems is proposed. Then, we introduce a new notion of so-called opacity-preserving stochastic simulation functions to quantify the distance between two systems in a probabilistic setting, while preserving approximate initial-state opacity across them. In addition, we show that for a class of stochastic control systems satisfying incremental input-to-state stability property, one can construct their finite abstractions (a.k.a finite Markov decision processes) together with a corresponding opacity-preserving stochastic simulation function between them.**

I. Introduction

Security analysis of cyber-physical systems (CPS) has attracted considerable attention in the past years. The notions of security for CPS can be classified into two categories. The first category deals with the capabilities of the intruders while the second one focuses on the information flow from the systems to the intruders. In this paper, we investigate a security property called *opacity*, which belongs to the second category. In this case, we assume that there are outside intruders who can monitor the system by observing the external behaviors (i.e. outputs) of the system. Opacity property essentially determines whether or not any trace that reveals secret behaviors of the system is indistinguishable from other traces to an intruder [1], [2], [3].

The concept of opacity was originally proposed by [4] in the realm of computer science for the analysis of cryptographic protocols. Since then, the opacity problem has been widely studied in Discrete-Event Systems (DES) context. In order to capture various types of secret requirements, different notions of opacity have been proposed, including state-based notions in [5], [6], [7] and language-based notions in [8]. In practical situations, the state-based notions of opacity of DES are generally classified into the so-called initial-state opacity [7], current-state opacity [1], K-step opacity [5], and infinite-state opacity [6]. We refer interested readers to the recent surveys in [9], [3] describing those different notions of opacity in details.

The verification of opacity has also been widely studied in [5], [6], [10]. Unfortunately, it is known (cf. the previous references) that existing algorithms for verifying opacity have exponential time complexity. Efficient approaches to overcome the computational complexity are highly demanding [11]. More recently, several promising abstraction-based techniques have been employed for efficiently verifying and enforcing opacity [2], [12]. Particularly, in [12], several notions of opacity-preserving (bi)simulation relations are proposed. The proposed relations essentially characterize the distance between two systems in terms of preservation of opacity. By constructing abstractions of the original systems and leveraging the relations between them, the complexity of the verification algorithms can be mitigated.

Since opacity is an information-flow property, its definition depends on the information model of the system. The existing results are mostly based on the event-observation models. In this context, it is assumed that outputs of the systems are symbolic so that different outputs can be precisely distinguished by the observer. However, many real life applications are metric systems (e.g. discrete-time control systems with continuous state sets and continuous output sets) in the sense that their outputs are physical signals equipped with some metrics. In this setting, it is not practical for the intruder to unambiguously distinguish outputs due to the measurement errors. A more reasonable concept of opacity for metric systems called approximate opacity was lately proposed in [13]. These new notions provide relaxed versions of opacity to quantitatively evaluate the security guarantee level with respect to the measurement precision of the intruder.

On the other hand, in real-world applications, a small probability of violation of the opacity could be tolerable. Hence, instead of simply asking if a system is opaque or non-opaque, it is more applicable to evaluate the possibility of being not secure for stochastic systems. This direction has been recently explored in the context of stochastic DES [14], [15], [16], [17], [18], [11]. For example, in [14] three different notions of probabilistic opacity were introduced for current-state opacity; this approach has also been extended to infinite-step opacity by [11]. In [16], Jensen-Shannon divergence was adopted to quantify secrecy loss in stochastic systems. Opacity of (partially-observed) Markov decision processes has also been studied in [15], [17], [18]. Note that most of the existing works on opacity analysis of stochastic DES are based on finite systems. In discrete-time stochastic control systems, however, the state-sets are usually infinite, which makes the verification problem very challenging and even undecidable. Therefore, efficient abstraction techniques,

*This work was supported in part by the German Research Foundation (DFG) through the grant ZA 873/1-1 and by China Scholarship Council.

Siyuan Liu is with the Department of Electrical and Computer Engineering, Technical University of Munich, Germany. Xiang Yin is with the Department of Automation, Shanghai Jiao Tong University, China. Majid Zamani is with the Computer Science Department, University of Colorado Boulder, USA. Majid Zamani is with the Computer Science Department, Ludwig Maximilian University of Munich, Germany. Emails: sy.liu@tum.de, yinxiang@sjtu.edu.cn, majid.zamani@colorado.edu.

978-1-5386-8266-1/$31.00 ©2020 AACC

together with suitably defined notions of stochastic opacity, are needed in order to quantitatively evaluate the security level of large-scale infinite stochastic systems.

In this paper, we introduce a new notion of initial-state opacity for discrete-time stochastic control systems, which is called (δ,ε)-approximate initial-state opacity. Our notion can be regarded as the stochastic counterpart of the notion of approximate opacity introduced in [13]. In particular, the δ-approximate initial-state opacity proposed in [13] requires that given a threshold parameter $\delta \geq 0$ (based on the measurement precision of the intruder), for any state run starting from a secret state, there always exists another state run starting from a non-secret state, such that the largest distance between their output runs is smaller than δ. In this paper, the aim is still to ensure that discerning which of the states was the originating one is difficult for an intruder based on its observation. Particularly, starting from two initial states, the output trajectories are considered indistinguishable if the probability measures of them remaining in any set of interest are close to each other. The parameter ε is used to bound the probability distance and δ captures the measurement precision of the intruder. In the special case when $\delta = 0$, the notion boils down to ε-approximate initial-state opacity, and the parameter ε can be captured *exactly* by the well-known *total variation distance*, see, e.g. [19], [20], [21], [22].

In addition, we introduce a notion of initial-state opacity-preserving stochastic simulation function. This function is essentially used to quantitatively relate two systems in terms of opacity satisfactions in a probabilistic setting. Unfortunately, the existing notions of stochastic simulation functions do not necessarily preserve opacity. We show that if our notion of opacity-preserving stochastic simulation function exists between systems Σ and $\hat{\Sigma}$, then $\hat{\Sigma}$ being ε-approximate initial-state opaque implies that Σ is $(2\lambda, \varepsilon + 2\bar{\varepsilon}_\lambda)$-approximate initial-state opaque for any arbitrarily chosen $\lambda > 0$ and $\bar{\varepsilon}_\lambda \geq 0$ which is a function of λ (cf. equation (6)). As a result, this allows us to efficiently verify opacity of a complex system with possibly uncountable number of states by analyzing it over its simpler (potentially finite) abstraction. Finally, we propose a scheme to construct abstractions (in the form of finite Markov decision processes (MDPs)) for a class of incrementally input-to-state stable discrete-time stochastic control systems. We show that under some mild assumptions, the original and finite systems are related to each other through an opacity-preserving stochastic simulation function. Existing techniques for computing total variation distance in the case of labelled Markov chain (LMC) [20], [21] can be readily employed for the verification of ε-approximate initial-state opacity on the constructed finite MDPs.

II. Discrete-Time Stochastic Control Systems

A. Notation

In this paper, a probability space is written as $(\Omega, \mathcal{F}_\Omega, \mathbb{P})$, where Ω is the sample space, $\mathcal{F}_\Omega$ is a sigma-algebra on Ω representing a set of events, and $\mathbb{P} : \mathcal{F}_\Omega \rightarrow [0,1]$ is a probability measure that assigns probabilities to events. We denote the sets of nonnegative and positive integers, respectively, by $\mathbb{N} := \{0,1,2,\ldots\}$ and $\mathbb{N}_{\geq 1} := \{1,2,3,\ldots\}$. The sets of real, positive and nonnegative real numbers are denoted by $\mathbb{R}$, $\mathbb{R}_{>0}$, and $\mathbb{R}_{\geq 0}$, respectively. Given $N \in \mathbb{N}_{\geq 1}$ vectors $x_i \in \mathbb{R}^{n_i}$, with $i \in \{1,\ldots,N\}$, $n_i \in \mathbb{N}_{\geq 1}$, and $n = \sum_i n_i$, we use $x = [x_1;\ldots;x_N]$ to denote the corresponding concatenated vector in $\mathbb{R}^n$. We denote by $\|\cdot\|$ the infinity norm of a vector $x \in \mathbb{R}^n$. Given functions f and g, the composition of them is denoted by $f \circ g$. A continuous function $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is said to be a class $\mathcal{K}$ function if it is strictly increasing and $\gamma(0) = 0$. A class $\mathcal{K}$ function $\gamma(\cdot)$ is said to be a class $\mathcal{K}_\infty$ function if $\gamma(r) \rightarrow \infty$ as $r \rightarrow \infty$.

B. Discrete-Time Stochastic Control Systems

First, we define discrete-time stochastic control systems (dt-SCS) as the sextuple

$$\Sigma = (X, U, \varsigma, f, Y, h), \tag{1}$$

where $X \subseteq \mathbb{R}^n$, $U \subseteq \mathbb{R}^m$, and $Y \subseteq \mathbb{R}^q$ are Borel spaces denoting the state, input and output sets of the system, respectively. We use $\mathcal{B}(X)$ to denote the Borel sigma-algebra on the state set X, thus $(X, \mathcal{B}(X))$ denotes the corresponding measurable space. In the probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P})$, we use $\varsigma = (\varsigma(1), \varsigma(2), \ldots)$ to denote a sequence of independent and identically distributed (i.i.d.) random variables from Ω to the measurable set V_ς, where $\varsigma(k) : \Omega \rightarrow V_\varsigma, k \in \mathbb{N}$.

The maps $f : X \times U \times V_\varsigma \rightarrow X$ and $h : X \rightarrow Y$ are measurable functions serving as the state transition relation and output map, respectively. Given an initial state $\xi(0) \in X$, the dt-SCS Σ satisfies

$$\Sigma : \begin{cases} \xi(k+1) = f(\xi(k), \nu(k), \varsigma(k)), \\ \zeta(k) = h(\xi(k)), \end{cases} \quad k \in \mathbb{N}, \tag{2}$$

where $\xi(\cdot) : \mathbb{N} \rightarrow X$, $\zeta(\cdot) : \mathbb{N} \rightarrow Y$, and $\nu(\cdot) : \mathbb{N} \rightarrow U$ are the state, output, and input signals, respectively. We use $\mathcal{U}$ to denote a collection of sequences $\nu : \Omega \rightarrow U$, where $\nu(k)$ is independent of $\varsigma(t)$ for any $k, t \in \mathbb{N}$ and $t \geq k$. A dt-SCS defined in (1) with (possibly) continuous state set can be equivalently represented as a general Markov decision process (gMDP). The finite abstractions of gMDPs are called finite MDPs, as constructed later in Section V-A, which are of crucial use for the analysis of opacity. We refer the interested readers to [23] for formal definitions of gMDPs and MDPs.

III. Approximate Opacity for Stochastic Control Systems

In this section, we introduce the concept of approximate opacity for the class of discrete-time stochastic control systems. Hereafter, we slightly modify the formulation in (1) to accommodate for initial states and secret states, as $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$, where $X_0 \subseteq X$ is a set of initial states and $X_S \subseteq X$ is a set of secret states. Given system $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$, $x \xrightarrow{u} x'$ is called a *transition* in the system if and only if $x' = f(x, u, \varsigma)$. The random sequence $\xi_{x_0\nu} : \Omega \times \mathbb{N} \rightarrow X$, which is in the form of $\xi_{x_0\nu} = (x_0, x_1, \ldots, x_n)$, is said to be a *solution process* of Σ under input sequence $\nu = (u_1, u_2, \ldots, u_n)$ satisfying

(2), with initial state $\xi_{x_0 \upsilon}(0) = x_0$. The random sequence $\zeta_{x_0 \nu} : \Omega \times \mathbb{N} \rightarrow Y$ is called the *output run* and defined as $\zeta_{x_0 \nu} = (y_0, y_1, \ldots, y_n)$ such that there exists a finite state run $\xi_{x_0 \nu} = (x_0, x_1, \ldots, x_n)$ with $y_i = h(x_i)$, for $i \in \{0, \ldots, n\}$. A finite state run and a finite output run can be extended to an infinite state run and an infinite output run as well.

In order to show our new notion of opacity, we define for any system $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$ with output set Y, set $B_Y = \{y = (y_0, y_1, \ldots, y_n) \in Y^n | n \in \mathbb{N}_{\geq 1}\}$ which is the set of all finite output sequences. For any measurable set $E \subseteq B_Y$, for some $\delta > 0$, we denote the δ neighborhood of set E by $\underline{E}_\delta$ and $\bar{E}^\delta$, where $\underline{E}_\delta$ is the largest measurable set contained in E satisfying:

$$\underline{E}_\delta = \{y \in E | \forall \bar{y} \in B_Y \setminus E, \|\bar{y}(i) - y(i)\| \geq \delta, \forall i \in \{0, \ldots, n\}\}, \quad (3)$$

and $\bar{E}^\delta$ is the smallest measurable set containing E satisfying:

$$\bar{E}^\delta = \{y \in B_Y \mid \exists \bar{y} \in E, \|\bar{y}(i) - y(i)\| \leq \delta, \forall i \in \{0, \ldots, n\}\}. (4)$$

Note that we can regard $\underline{E}_\delta$ and $\bar{E}^\delta$ respectively as the δ-deflated version and δ-inflated version of set E.

Now, we introduce a notion of opacity for the class of stochastic systems defined above.

Definition 3.1: Let $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$ be a dt-SCS and consider constants $\delta \geq 0$, $0 \leq \varepsilon < 1$. System Σ is (δ, ε)-approximate initial-state opaque if for any $x_0 \in X_0 \cap X_S$, there exists $x'_0 \in X_0 \setminus X_S$, so that for any input sequence ν, there exisits an input sequence ν', such that for every measurable set $E \subseteq B_Y$ and the δ neighboring sets $\underline{E}_\delta$ and $\bar{E}^\delta$ as defined in (3) and (4), the following inequalities hold:

$$\mathbb{P}(\zeta_{x'_0 \nu'} \in \underline{E}_\delta) - \varepsilon \leq \mathbb{P}(\zeta_{x_0 \nu} \in E) \leq \mathbb{P}(\zeta_{x'_0 \nu'} \in \bar{E}^\delta) + \varepsilon, \quad (5)$$

where $\zeta_{x_0 \nu}$ and $\zeta_{x'_0 \nu'}$ are the output runs of the same length, starting from x_0 under ν and x'_0 under ν', respectively.

Remark 3.2: In this definition, we use parameter ε to denote the largest allowable probability violation for the output trajectories starting from the secret and non-secret initial states x_0 and x'_0 to be δ close. Note that the value of parameter δ is chosen depending on the measurement precision of a malicious intruder. In the case that the precision of the intruder is lower than δ, the δ neighborhood of set E, i.e. $\underline{E}_\delta$ and $\bar{E}^\delta$, is indistinguishable from set E from the intruder's point of view. When $\delta = 0$, the probability inequalities in (5) boils down to $|\mathbb{P}(\zeta_{x_0 \nu} \in E) - \mathbb{P}(\zeta_{x'_0 \nu'} \in E)| \leq \varepsilon$, and we use the term ε-approximate initial-state opacity. It is worth mentioning that, in this case the parameter ε can be captured exactly by total variation distance [20], [21] for the case of finite MDPs. Thus existing techniques for computing total variation distance can be leveraged as tools to check the probability distance in (5), which would be applicable for the verification of ε-approximate initial-state opacity for finite MDPs. Although computing this distance is shown to be NP-hard [20], [21], there have been some results to approximate the distance, which are #P-hard and in PSPACE see, e.g. [19], [22]. Moreover, in the case that $\delta = 0$, $\varepsilon = 0$, and no stochasticity exists in the transition functions of the systems, this notion boils down to *exact opacity* defined in [12] for general nondeterministic transition systems.

Note that throughout the work we assume $X_0 \nsubseteq X_S$, otherwise (δ, ε)-approximate initial-state opacity is trivially violated.

IV. Stochastic Simulation Functions for Opacity

In this section, we introduce a notion of initial-state opacity-preserving stochastic simulation functions for dt-SCS. The stochastic simulation function will play an important role in analyzing opacity for dt-SCS. First, we provide the definition of initial-state opacity-preserving stochastic simulation functions.

Definition 4.1: (Initial-state opacity-preserving stochastic simulation function) Let $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$ and $\widehat{\Sigma} = (\hat{X}, \hat{X}_0, \hat{X}_S, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$ be two dt-SCS with the same output sets $Y = \hat{Y}$. A function $V : X \times \hat{X} \rightarrow \mathbb{R}_{\geq 0}$ is called an initial-state opacity-preserving stochastic simulation function (InitSOP-SSF) from $\hat{\Sigma}$ to Σ, if there exist constants $\psi \geq 0$, $\omega \geq 0$, a function $\alpha \in \mathcal{K}_\infty$, and a function $\kappa \in \mathcal{K}$ which satisfies $\kappa(s) \geq \hat{\kappa} s$, $\forall s \in \mathbb{R}_{\geq 0}$, where $0 < \hat{\kappa} < 1$, such that

1) a) $\forall x_0 \in X_0 \cap X_S$, $\exists \hat{x}_0 \in \hat{X}_0 \cap \hat{X}_S$: $V(x_0, \hat{x}_0) \leq \omega$;
   b) $\forall \hat{x}_0 \in \hat{X}_0 \setminus \hat{X}_S$, $\exists x_0 \in X_0 \setminus X_S$: $V(x_0, \hat{x}_0) \leq \omega$;
2) $\forall x \in X, \forall \hat{x} \in \hat{X}$, $\alpha(\|h(x) - \hat{h}(\hat{x})\|) \leq V(x, \hat{x})$;
3) $\forall x \in X, \forall \hat{x} \in \hat{X}$, the following conditions hold:
   a) $\forall u$, $\exists \hat{u}$, s.t. $\mathbb{E}\Big[V(f(x, u, \varsigma), \hat{f}(\hat{x}, \hat{u}, \varsigma)) \,\big|\, x, \hat{x}, u, \hat{u}\Big] - V(x, \hat{x}) \leq -\kappa(V(x, \hat{x})) + \psi$;
   b) $\forall \hat{u}$, $\exists u$, s.t. $\mathbb{E}\Big[V(f(x, u, \varsigma), \hat{f}(\hat{x}, \hat{u}, \varsigma)) \,\big|\, x, \hat{x}, u, \hat{u}\Big] - V(x, \hat{x}) \leq -\kappa(V(x, \hat{x})) + \psi$.

Now, before stating the main theorem in this section, we provide the following technical proposition which is inspired by Theorem 3.3 in [24]. This proposition shows us the usefulness of the InitSOP-SSF in the sense that it can be employed to show indistinguishability of output trajectories of two dt-SCS in a probabilistic setting.

Proposition 4.2: Let $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$ and $\widehat{\Sigma} = (\hat{X}, \hat{X}_0, \hat{X}_S, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$ be two dt-SCS with the same output sets $Y = \hat{Y}$. Suppose V is an InitSOP-SSF from $\widehat{\Sigma}$ to Σ. Then, for any $a \in X_0 \cap X_S$ in Σ, there exists $\hat{a} \in \hat{X}_0 \cap \hat{X}_S$ in $\hat{\Sigma}$ (respectively, for any $\hat{a} \in \hat{X}_0 \setminus \hat{X}_S$ in $\hat{\Sigma}$, there exists $a \in X_0 \setminus X_S$ in Σ) so that for any $\hat{\nu} \in \hat{\mathcal{U}}$ in $\hat{\Sigma}$, there exists $\nu \in \mathcal{U}$ in Σ and vice versa such that the following inequality holds

$$\mathbb{P}\left\{\sup_{0 \leq k \leq n} \|\zeta_{a\nu}(k) - \hat{\zeta}_{\hat{a}\hat{\nu}}(k)\| \leq \lambda \,|\, [a; \hat{a}]\right\} \geq 1 - \bar{\varepsilon}_\lambda,$$

$$\bar{\varepsilon}_\lambda := \begin{cases} 1 - (1 - \frac{\omega}{\alpha(\lambda)})(1 - \frac{\psi}{\alpha(\lambda)})^n & \text{if } \alpha(\lambda) \geq \frac{\psi}{\hat{\kappa}}, \\ (\frac{\omega}{\alpha(\lambda)})(1 - \hat{\kappa})^n + (\frac{\psi}{\hat{\kappa}\alpha(\lambda)})(1 - (1 - \hat{\kappa})^n) & \text{if } \alpha(\lambda) < \frac{\psi}{\hat{\kappa}}, \end{cases} \quad (6)$$

for any $\lambda > 0$.

Proof: It can be readily seen that by conditions 2) and 3) in Definition 4.1, the InitSOP-SSF is a stochastic simulation function (SSF) (as defined in [24, Definition 3.2]) both from $\widehat{\Sigma}$ to Σ and from Σ to $\widehat{\Sigma}$. Since by condition 1) in Definition 4.1, $V(a, \hat{a}) \leq \omega$, the rest of the proof is concluded by applying Theorem 3.3 in [24]. ■

This proposition will be used for the proof of the following main theorem, where we show preservation of approximate initial-state opacity across related systems as in Definition 4.1. The next lemmas will be used to prove the main result.

Lemma 4.3: Suppose for two dt-SCS Σ and $\hat{\Sigma}$, the output trajectories $\zeta_{a\nu}$ and $\hat{\zeta}_{\hat{a}\hat{\nu}}$ satisfy the inequality

$$\sup_{0\leq k\leq n} \|\zeta_{a\nu}(k) - \hat{\zeta}_{\hat{a}\hat{\nu}}(k)\| \leq \lambda,$$

for some time bound n and $\lambda > 0$. Then we have:

$$\hat{\zeta}_{\hat{a}\hat{\nu}} \in \underline{E}_\lambda \Longrightarrow \zeta_{a\nu} \in E; \zeta_{a\nu} \in E \Longrightarrow \hat{\zeta}_{\hat{a}\hat{\nu}} \in \bar{E}^\lambda,$$

over time interval $[0, n]$, for any measurable set $E \subseteq B_Y$ and the modified sets $\underline{E}_\lambda$ and $\bar{E}^\lambda$ as defined in (3) and (4).

Proof: As can be seen from the definition of $\underline{E}_\lambda$ and $\bar{E}^\lambda$ in (3) and (4), given any set of output sequences $E \subseteq B_Y$, $\underline{E}_\lambda$ and $\bar{E}^\lambda$ are the λ-deflated version and λ-inflated version of set E, respectively. Since we have

$$\sup_{0\leq k\leq n} \|\zeta_{a\nu}(k) - \hat{\zeta}_{\hat{a}\hat{\nu}}(k)\| \leq \lambda,$$

then it can be readily seen that according to the structure of $\underline{E}_\lambda$ and $\bar{E}^\lambda$, $\hat{\zeta}_{\hat{a}\hat{\nu}} \in \underline{E}_\lambda$ guarantees $\zeta_{a\nu} \in E$. Similarly, $\zeta_{a\nu} \in E$ inplies $\hat{\zeta}_{\hat{a}\hat{\nu}} \in \bar{E}^\lambda$ as well. ■

This lemma essentially provides us the relation between the property satisfactions of two dt-SCS, given that the output trajectories of these two dt-SCS are close to each other. Based on this lemma, the following lemma presents another technical result of this paper.

Lemma 4.4: Suppose Σ and $\widehat{\Sigma}$ are two dt-SCS for which inequality (6) holds with initial states a and $\hat{a}$, input sequences ν and $\hat{\nu}$, a constant pair $(\lambda,\bar{\varepsilon}_\lambda)$ and any time bound n. The following inequality holds for any set $E \subseteq B_Y$ and the modified sets $\underline{E}_\lambda$ and $\bar{E}^\lambda$ as defined in (3) and (4):

$$\mathbb{P}(\hat{\zeta}_{\hat{a}\hat{\nu}} \in \underline{E}_\lambda) - \bar{\varepsilon}_\lambda \leq \mathbb{P}(\zeta_{a\nu} \in E) \leq \mathbb{P}(\hat{\zeta}_{\hat{a}\hat{\nu}} \in \bar{E}^\lambda) + \bar{\varepsilon}_\lambda, \quad (7)$$

where the satisfaction is over time interval $\{0, \ldots, n\}$.

Proof: Let us consider the events: $\mathcal{E}_1 {:=} \{\hat{\zeta}_{\hat{a}\hat{\nu}} \in \underline{E}_\lambda\}$, $\mathcal{E}_2 {:=} \{\zeta_{a\nu} \in E\}$ and $\mathcal{E}_3 {:=} \{\sup_{0\leq k\leq n} \|\zeta_{a\nu}(k) - \hat{\zeta}_{\hat{a}\hat{\nu}}(k)\| \leq \lambda\}$. Since we have from Lemma (4.3), $\mathcal{E}_1 \cap \mathcal{E}_3 \Longrightarrow \mathcal{E}_2$, thus, $\mathbb{P}(\bar{\mathcal{E}}_2) \leq \mathbb{P}(\bar{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_3) \leq \mathbb{P}(\bar{\mathcal{E}}_1) + \mathbb{P}(\bar{\mathcal{E}}_3)$, where $\bar{\mathcal{E}}_1$, $\bar{\mathcal{E}}_2$ and $\bar{\mathcal{E}}_3$ are the complements of $\mathcal{E}_1$, $\mathcal{E}_2$ and $\mathcal{E}_3$, respectively. As we have by (6), $\mathbb{P}(\bar{\mathcal{E}}_3) \leq \bar{\varepsilon}_\lambda$, now we readily get:

$$\begin{aligned}\mathbb{P}(\bar{\mathcal{E}}_2) \leq \mathbb{P}(\bar{\mathcal{E}}_1) + \bar{\varepsilon}_\lambda &\Longrightarrow 1 - \mathbb{P}(\mathcal{E}_2) \leq 1 - \mathbb{P}(\mathcal{E}_1) + \bar{\varepsilon}_\lambda \\ &\Longrightarrow \mathbb{P}(\mathcal{E}_1) \leq \mathbb{P}(\mathcal{E}_2) + \bar{\varepsilon}_\lambda,\end{aligned}$$

which gives us $\mathbb{P}(\hat{\zeta}_{\hat{a}\hat{\nu}} \in \underline{E}_\lambda) - \bar{\varepsilon}_\lambda \leq \mathbb{P}(\zeta_{a\nu} \in E)$. The proof of $\mathbb{P}(\zeta_{a\nu} \in E) \leq \mathbb{P}(\hat{\zeta}_{\hat{a}\hat{\nu}} \in \bar{E}^\lambda) + \bar{\varepsilon}_\lambda$ is similar and is omitted here due to lack of space. ■

Now, we present the main result of this paper on the preservation of opacity across related dt-SCS systems.

Theorem 4.5: Let $\Sigma = (X,X_0,X_S,U,\varsigma,f,Y,h)$ and $\widehat{\Sigma} = (\hat{X},\hat{X}_0,\hat{X}_S,\hat{U},\varsigma,\hat{f},\hat{Y},\hat{h})$ be two dt-SCS with the same output sets $Y = \hat{Y}$. Consider constants $\varepsilon \in \mathbb{R}_{\geq 0}$ and $\lambda \in \mathbb{R}_{>0}$. Assume V is an InitSOP-SSF from $\hat{\Sigma}$ to Σ as in Definition 4.1 with the corresponding constants ψ, ω, $\hat{\kappa}$ and $\mathcal{K}_\infty$ function α. Then the following implication holds:

$$\begin{aligned}&\widehat{\Sigma} \text{ is } \varepsilon\text{-approximate initial-state opaque} \\ &\Rightarrow \Sigma \text{ is } (2\lambda, \varepsilon + 2\bar{\varepsilon}_\lambda)\text{-approximate initial-state opaque}, \quad (8)\end{aligned}$$

where $\bar{\varepsilon}_\lambda \in \mathbb{R}_{\geq 0}$ is computed as in (6).

Proof: Consider an arbitrary secret initial state $x_0 \in X_0 \cap X_S$, input sequence $\nu = \{u_1, u_2, \ldots, u_n\}$ and the corresponding state run $\xi_{x_0\nu} = (x_0, x_1, \ldots, x_n)$ in Σ. Since V is an InitSOP-SSF from $\hat{\Sigma}$ to Σ, by conditions 1)a), 2) and 3)a) in Definition 4.1, there exist a secret initial state $\hat{x}_0 \in \hat{X}_0 \cap \hat{X}_S$, input sequence $\hat{\nu} = \{\hat{u}_1, \hat{u}_2, \ldots, \hat{u}_n\}$ and state run $\hat{\xi}_{\hat{x}_0\hat{\nu}} = (\hat{x}_0, \hat{x}_1, \ldots, \hat{x}_n)$ in $\hat{\Sigma}$ such that $V(x_0, \hat{x}_0) \leq \omega$, and $\forall i \in \{0, 1, \ldots, n\}$:

$$\begin{aligned}&\alpha(\|h(x_i) - \hat{h}(\hat{x}_i)\|) \leq V(x_i, \hat{x}_i), \\ &\mathbb{E}\Big[V(f(x_i, u_i, \varsigma_i)), \hat{f}(\hat{x}_i, \hat{u}_i, \varsigma_i)) \,\Big|\, x_i, \hat{x}_i, u_i, \hat{u}_i\Big] \\ &- V(x_i, \hat{x}_i) \leq -\kappa(V(x_i, \hat{x}_i)) + \psi.\end{aligned}$$

By applying Proposition 4.2, for the given λ, we have:

$$\mathbb{P}\left\{\max_{0\leq i\leq n} \|\zeta_{x_0\nu}(i) - \zeta_{\hat{x}_0\hat{\nu}}(i)\| \leq \lambda \,|\, [x_0; \hat{x}_0]\right\} \geq 1 - \bar{\varepsilon}_\lambda,$$

where $\bar{\varepsilon}_\lambda$ is computed using inequality (6). By applying (7) in Lemma 4.4, we get for any set $E \subseteq B_Y$ and the modified sets $\underline{E}_\lambda$ and $\bar{E}^\lambda$:

$$\mathbb{P}(\hat{\zeta}_{\hat{x}_0\hat{\nu}} \in \underline{E}_\lambda) - \mathbb{P}(\zeta_{x_0\nu} \in E) \leq \bar{\varepsilon}_\lambda, \quad (9)$$

$$\mathbb{P}(\zeta_{x_0\nu} \in E) - \mathbb{P}(\hat{\zeta}_{\hat{x}_0\hat{\nu}} \in \bar{E}^\lambda) \leq \bar{\varepsilon}_\lambda. \quad (10)$$

Since $\hat{\Sigma}$ is ε-approximate initial-state opaque, by (5) in Definition 3.1, there exist a non-secret initial state $\hat{x}_0' \in \hat{X}_0 \setminus \hat{X}_S$, input sequence $\hat{\nu}' = \{\hat{u}_1', \hat{u}_2', \ldots, \hat{u}_n'\}$ and state run $\xi_{\hat{x}_0'\hat{\nu}'} = (\hat{x}_0', \hat{x}_1', \ldots, \hat{x}_n')$ in $\hat{\Sigma}$, such that $\|\mathbb{P}(\hat{\zeta}_{\hat{x}_0'\hat{\nu}'} \in E) - \mathbb{P}(\hat{\zeta}_{\hat{x}_0\hat{\nu}} \in E)\| \leq \varepsilon$ holds for any set E, so we have:

$$\mathbb{P}(\hat{\zeta}_{\hat{x}_0'\hat{\nu}'} \in \underline{E}_\lambda) - \mathbb{P}(\hat{\zeta}_{\hat{x}_0\hat{\nu}} \in \underline{E}_\lambda) \leq \varepsilon, \quad (11)$$

$$\mathbb{P}(\hat{\zeta}_{\hat{x}_0\hat{\nu}} \in \bar{E}^\lambda) - \mathbb{P}(\hat{\zeta}_{\hat{x}_0'\hat{\nu}'} \in \bar{E}^\lambda) \leq \varepsilon. \quad (12)$$

Again, since V is an InitSOP-SSF from $\hat{\Sigma}$ to Σ, by conditions 1)b), 2) and 3)b) in Definition 4.1, Proposition 4.2 and (7) in Lemma 4.4, there exist an initial state $x_0' \in X_0 \setminus X_S$, input sequence $\nu' = \{u_1', u_2', \ldots, u_n'\}$ and the corresponding state run $\xi_{x_0'\nu'} = (x_0', x_1', \ldots, x_n')$ in Σ such that

$$\mathbb{P}(\zeta_{x_0'\nu'} \in \underline{E}_{2\lambda}) - \mathbb{P}(\hat{\zeta}_{\hat{x}_0'\hat{\nu}'} \in \underline{E}_\lambda) \leq \bar{\varepsilon}_\lambda, \quad (13)$$

$$\mathbb{P}(\hat{\zeta}_{\hat{x}_0'\hat{\nu}'} \in \bar{E}^\lambda) - \mathbb{P}(\zeta_{x_0'\nu'} \in \bar{E}^{2\lambda}) \leq \bar{\varepsilon}_\lambda. \quad (14)$$

Hence, by combining inequalities (9), (11), (13), we have the following result

$$\mathbb{P}(\zeta_{x_0'\nu'} \in \underline{E}_{2\lambda}) - \mathbb{P}(\zeta_{x_0\nu} \in E) \leq \varepsilon + 2\bar{\varepsilon}_\lambda. \quad (15)$$

Additionally, combining inequalities (10), (12), (14), we get

$$\mathbb{P}(\zeta_{x_0\nu} \in E) - \mathbb{P}(\zeta_{x_0'\nu'} \in \bar{E}^{2\lambda}) \leq \varepsilon + 2\bar{\varepsilon}_\lambda. \quad (16)$$

Since $x_0 \in X_0 \cap X_S$ and input sequence ν in Σ are arbitrary, we conclude that Σ is $(2\lambda, \varepsilon + 2\bar{\varepsilon}_\lambda)$-approximate initial-state opaque. ■

Remark 4.6: The theorem provides a sufficient condition for approximate initial-state opacity based on the relation between two stochastic systems. It bridges the gap between the verification of opacity and abstraction-based techniques on stochastic systems. To be specific, when analyzing opacity for a large system Σ, the verification procedure can be highly time-consuming. By constructing an abstraction of system Σ, which appears as system $\hat{\Sigma}$ in the theorem, and leveraging the simulation relation between them, one can efficiently verify opacity of the complex system Σ. The abstraction is contructed as a finite Markov decision process (MDP) in the following subsection V-A. In addition, as mentioned before, ε-approximate initial-state opacity for the MDP can be verified easily using existing computation algorithms for total variation distance.

V. Opacity of Stochastic Control Systems based on Finite Abstractions

In this section, we show how to analyze approximate opacity for the class of dt-SCS based on their finite abstractions (finite MDPs). First, we provide the construction of finite abstractions of the concrete systems.

A. Finite Abstractions of dt-SCS

Given a dt-SCS $\Sigma = (X, X_0, X_S, U, \varsigma, f, Y, h)$, we construct a finite MDP as its finite abstraction, represented by the tuple $\widehat{\Sigma} = (\hat{X}, \hat{X}_0, \hat{X}_S, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$. The construction of finite MDPs follows a similar procedure as in [24, Algorithm 1] with some modifications to incorporate the role of secret sets. First, for the given state set X and input set U, we select finite partitions of them as $X = \cup_i \mathsf{X}_i, U = \cup_i \mathsf{U}_i$, and single representative points $\bar{x}_i \in \mathsf{X}_i,\ \bar{u}_i \in \mathsf{U}_i$ as abstract states and inputs. Then, we define the state and input sets of $\widehat{\Sigma}$ as $\hat{X} = \{\bar{x}_i, i = 1, \dots, n_x\}$, and $\hat{U} = \{\bar{u}_i, i = 1, \dots, n_u\}$, which simply consist of the selected representative points. The transition function $\hat{f} : \hat{X} \times \hat{U} \times V_\varsigma \to \hat{X}$ is defined as

$$\hat{f}(\hat{x}, \hat{u}, \varsigma) = \Pi_x(f(\hat{x}, \hat{u}, \varsigma)), \tag{17}$$

where $\Pi_x : X \to \hat{X}$ represents the map that assigns to any $x \in X$, the representative point $\hat{x} \in \hat{X}$ of the corresponding partition set containing x. The output set $\hat{Y}$ is the image of $\hat{X}$ under h, with $\hat{h} : \hat{X} \to \hat{Y}$ being the same as h except for having a restricted domain $\hat{X}$. Similarly, we use $\Pi_u : U \to \hat{U}$ to denote the map that assigns to any $u \in U$, the representative input point $\hat{u} \in \hat{U}$ of the corresponding partition set containing u.

Remark 5.1: Note that we have not defined the map $\Pi_x : X \to \hat{X}$ yet. For example, one can choose center points (if applicable) of the partitions as representative points or apply other specialized mapping rule to it. In this work, we enforce two conditions as the rule of choosing representative points for initial states and secret states as follows:

1) If $X_0 \cap \mathsf{X}_i \neq \emptyset$, we constrain the representative point of X_i to be chosen as $\bar{x}_i \in X_0 \cap \mathsf{X}_i$;
2) If $X_S \cap \mathsf{X}_i \neq \emptyset$, we constrain the representative point of X_i to be chosen as $\bar{x}_i \in X_S \cap \mathsf{X}_i$.

By the above conditions, one can observe that the initial state set $\hat{X}_0$ and secret state set $\hat{X}_S$ satisfy $\hat{X}_0 \subseteq X_0$ and $\hat{X}_S \subseteq X_S$.

Remark 5.2: In this paper, it is assumed that the abstraction maps Π_x and Π_u satisfy the inequalities

$$\|\Pi_x(x) - x\| \leq \mu_x, \forall x \in X, \|\Pi_u(u) - u\| \leq \mu_u, \forall u \in U, \tag{18}$$

where μ_x and μ_u are the *state* and *input* discretization parameter defined as

$$\mu_x := \sup\{\|x - x'\|,\ x, x' \in \mathsf{X}_i,\ i = 1, 2, \dots, n_x\}, \tag{19}$$

$$\mu_u := \sup\{\|u - u'\|,\ u, u' \in \mathsf{U}_i,\ i = 1, 2, \dots, n_u\}. \tag{20}$$

Next, we construct the InitSOP-SSF for a class of nonlinear stochastic systems.

B. Establishing InitSOP-SSF for a Class of Nonlinear Stochastic Systems

In this subsection, we focus on a general class of nonlinear stochastic systems Σ. We provide an InitSOP-SSF candidate for the concrete systems Σ and their finite MDPs as constructed in the previous subsection. The existence of such an InitSOP-SSF enables us to verify opacity of a continuous-space stochastic system by leveraging its finite abstraction. The establishment of InitSOP-SSF is under the following two assumptions. First, we assume that the output map h satisfies the following general Lipschitz assumption: there exists an $\tilde{\alpha} \in \mathcal{K}_\infty$ such that $\|h(x) - h(x')\| \leq \tilde{\alpha}(\|x - x'\|)$ for all $x, x' \in X$. Second, we assume that the concrete system is *incrementally input-to-state stable* as in the following definition.

Definition 5.3: A dt-SCS Σ is *incrementally input-to-state stable* if there exists function $V : X \times X \to \mathbb{R}_{\geq 0}$ such that $\forall x, x' \in X$, $\forall u, u' \in U$ the following two inequalities hold:

$$\underline{\alpha}(\|x - x'\|) \leq V(x, x') \leq \overline{\alpha}(\|x - x'\|), \tag{21}$$

$$\mathbb{E}\Big[V(f(x, u, \varsigma), f(x', u', \varsigma)) \big| x, x', u, u'\Big] - V(x, x') \leq -\bar{\kappa}(V(x, x')) + \bar{\rho}(\|u - u'\|), \tag{22}$$

for some $\underline{\alpha}, \overline{\alpha} \in \mathcal{K}_\infty$, $\bar{\kappa} \in \mathcal{K}$, and $\bar{\rho} \in \mathcal{K}_\infty \cup \{0\}$.

Now, we provide the main theorem in this subsection. We show that by adding a mild condition, the function V described in Definition 5.3 is indeed an InitSOP-SSF from the finite abstraction $\widehat{\Sigma}$ (as constructed in Subsection V-A) to the concrete system Σ.

Theorem 5.4: Let Σ be an incrementally input-to-state stable dt-SCS via a function V as in Definition 5.3 and $\widehat{\Sigma}$ be its *finite* MDP constructed as in Subsection V-A. Suppose there exists a constant $0 < \hat{\kappa} < 1$ such that the function $\bar{\kappa} \in \mathcal{K}$ satisfies $\bar{\kappa}(s) \geq \hat{\kappa} s$, $\forall s \in \mathbb{R}_{\geq 0}$. Assume that there exists a function $\gamma \in \mathcal{K}_\infty$ such that V satisfies

$$V(x, x') - V(x, x'') \leq \gamma(\|x' - x''\|), \forall x, x', x'' \in X. \tag{23}$$

Then V is an InitSOP-SSF from $\widehat{\Sigma}$ to Σ.

Proof: We start by proving condition 1) in Definition 4.1. For every initial state $x_0 \in X_0 \cap X_S$ in Σ, there always exists a representative point $\hat{x}_0 = \Pi_x(x_0)$ in $\widehat{\Sigma}$ which is inside the set $\hat{X}_0 \cap \hat{X}_S$ by the construction of $\hat{X}_0$ and $\hat{X}_S$, and $\|\hat{x}_0 - x_0\| \leq \mu_x$ holds by (18). Hence, we have $V(x_0, \hat{x}_0) \leq \overline{\alpha}(\|x_0 - \hat{x}_0\|)$ by (21), and condition 1)a) in Definition 4.1 is satisfied with $\omega = \overline{\alpha}(\mu_x)$. For every $\hat{x}_0 \in \hat{X}_0 \setminus \hat{X}_S$, by choosing $x_0 = \hat{x}_0$ which is also inside $X_0 \setminus X_S$, we get $V(x_0, \hat{x}_0) = 0 \leq \omega$. Hence, condition 1)b) in Definition 4.1 holds as well.

Let us show condition 2) in Definition 4.1 holds. Since Σ is incrementally input-to-state stable and using (21), and given the Lipschitz assumption on $h, \forall x \in X$ and $\forall \hat{x} \in \hat{X}$, one gets

$$\|h(x) - \hat{h}(\hat{x})\| \leq \tilde{\alpha}(\|x - \hat{x}\|) \leq \tilde{\alpha} \circ \underline{\alpha}^{-1}(V(x, \hat{x})),$$

which results in

$$\alpha(\|h(x) - \hat{h}(\hat{x})\|) \leq V(x, \hat{x}),$$

$\forall x \in X$ and $\forall \hat{x} \in \hat{X}$, where $\alpha(s) := (\tilde{\alpha} \circ \underline{\alpha}^{-1})^{-1}(s)$, $\forall s \in \mathbb{R}_{\geq 0}$. Hence, condition 2) in Definition 4.1 is satisfied. Let us now show that condition 3) in Definition 4.1 holds as well. Now, $\forall x \in X, \forall \hat{x} \in \hat{X}$, $\forall u \in U$ and $\forall \hat{u} \in \hat{U}$, by taking the conditional expectation from (23), we have

$$\begin{aligned}&\mathbb{E}\Big[V(f(x,u,\varsigma),\hat{f}(\hat{x},\hat{u},\varsigma)) \,\big|\, x,\hat{x},u,\hat{u}\Big]\\ &- \mathbb{E}\Big[V(f(x,u,\varsigma),f(\hat{x},\hat{u},\varsigma)) \,\big|\, x,\hat{x},u,\hat{u}\Big]\\ &\leq \mathbb{E}\Big[\gamma(\|\hat{f}(\hat{x}, \hat{u}, \varsigma) - f(\hat{x}, \hat{u}, \varsigma)\|) \,\big|\, x,\hat{x},u,\hat{u}\Big].\end{aligned}$$

Employing (22), one gets

$$\begin{aligned}&\mathbb{E}\Big[V(f(x, u, \varsigma), f(\hat{x}, \hat{u}, \varsigma)) \,\big|\, x,\hat{x},u,\hat{u}\Big]\\ &\leq V(x, \hat{x}) - \bar{\kappa}(V(x, \hat{x})) + \bar{\rho}(\|u - \hat{u}\|).\end{aligned} \tag{24}$$

Since $\hat{f}(\hat{x}, \hat{u}, \varsigma) = \Pi_x(f(\hat{x}, \hat{u}, \varsigma))$, by using (18), we get

$$\mathbb{E}\Big[\gamma(\|\hat{f}(\hat{x}, \hat{u}, \varsigma) - f(\hat{x}, \hat{u}, \varsigma)\|) \,\big|\, x,\hat{x},u,\hat{u}\Big] \leq \gamma(\mu_x).$$

Now, consider any $u \in U$. By choosing the representative input $\hat{u} = \Pi_u(u)$, which satisfies $\|u - \hat{u}\| \leq \mu_u$, we obtain

$$\begin{aligned}&\mathbb{E}\Big[V(f(x, u, \varsigma), \hat{f}(\hat{x}, \hat{u}, \varsigma)) \,\big|\, x,\hat{x},u,\hat{u}\Big] - V(x, \hat{x})\\ &\leq - \bar{\kappa}(V(x, \hat{x})) + \bar{\rho}(\mu_u) + \gamma(\mu_x).\end{aligned} \tag{25}$$

Hence, condition 3)a) in Definition 4.1 holds with $\psi = \bar{\rho}(\mu_u) + \gamma(\mu_x)$. Similarly, $\forall x \in X, \forall \hat{x} \in \hat{X}$, and $\forall \hat{u} \in \hat{U}$, by choosing $u = \hat{u}$, we have

$$\begin{aligned}&\mathbb{E}\Big[V(f(x, u, \varsigma), \hat{f}(\hat{x}, \hat{u}, \varsigma)) \,\big|\, x,\hat{x},u,\hat{u}\Big] - V(x, \hat{x})\\ &\leq - \bar{\kappa}(V(x, \hat{x})) + \gamma(\mu_x) \leq - \bar{\kappa}(V(x, \hat{x})) + \psi.\end{aligned}$$

Therefore, condition 3)b) in Definition 4.1 holds as well, and we conclude that V is an InitSOP-SSF from $\widehat{\Sigma}$ to Σ. ■

VI. Discussion

In this paper, we extended the notion of approximate opacity to discrete-time stochastic systems. The new notion called (δ, ε)-approximate initial-state opacity is proposed to evaluate the security level of a system to hide initial-state secret information. A stronger version of stochastic simulation function that preserves opacity was also proposed. By leveraging this simulation function, we discussed an effective approach to construct finite MDPs, which are in opacity-preserving relations with the original systems. Our result provides a promising way for verifying opacity of complex discrete-time stochastic control systems by the verification of opacity on the related finite MDPs, which can be done using computational algorithms for total variation distance. For the future work, we plan to extend our framework to cover more notions of opacity, e.g., K-step opacity, current-state opacity and infinite-state opacity.

References

[1] A. Saboori and C. N. Hadjicostis, "Notions of security and opacity in discrete event systems," in *IEEE Conference on Decision and Control*. IEEE, 2007, pp. 5056–5061.

[2] M. Noori-Hosseini, B. Lennartson, and C. Hadjicostis, "Compositional visible bisimulation abstraction applied to opacity verification," in *14th IFAC Workshop on Discrete Event Systems*, 2018, pp. 434–441.

[3] S. Lafortune, F. Lin, and C. N. Hadjicostis, "On the history of diagnosability and opacity in discrete event systems," *Annual Reviews in Control*, vol. 45, pp. 257–266, 2018.

[4] L. Mazaré, "Using unification for opacity properties," *Proceedings of the 4th Workshop on Issues in the Theory of Security*, vol. 7, pp. 165–176, 2004.

[5] A. Saboori and C. N. Hadjicostis, "Verification of k-step opacity and analysis of its complexity," *IEEE Transactions on Automation Science and Engineering*, vol. 8, no. 3, pp. 549–559, 2011.

[6] ——, "Verification of infinite-step opacity and complexity considerations," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1265–1269, 2012.

[7] ——, "Verification of initial-state opacity in security applications of discrete event systems," *Information Sciences*, vol. 246, pp. 115–132, 2013.

[8] F. Lin, "Opacity of discrete event systems and its applications," *Automatica*, vol. 47, no. 3, pp. 496–503, 2011.

[9] R. Jacob, J.-J. Lesage, and J.-M. Faure, "Overview of discrete event systems opacity: Models, validation, and quantification," *Annual Reviews in Control*, vol. 41, pp. 135–146, 2016.

[10] X. Yin and S. Lafortune, "A new approach for the verification of infinite-step and k-step opacity using two-way observers," *Automatica*, vol. 80, pp. 162–171, 2017.

[11] X. Yin, Z. Li, W. Wang, and S. Li, "Infinite-step opacity and k-step opacity of stochastic discrete-event systems," *Automatica*, vol. 99, pp. 266–274, 2019.

[12] K. Zhang, X. Yin, and M. Zamani, "Opacity of nondeterministic transition systems: A (bi) simulation relation approach," *IEEE Transactions on Automatic Control*, vol. 64, no. 12, pp. 5116–5123, 2019.

[13] X. Yin, M. Zamani, and S. Liu, "On approximate opacity of cyber-physical systems," *IEEE Transactions on Automatic Control, conditionally accepted, arXiv:1902.09411*, 2019.

[14] A. Saboori and C. N. Hadjicostis, "Current-state opacity formulations in probabilistic finite automata," *IEEE Transactions on automatic control*, vol. 59, no. 1, pp. 120–133, 2014.

[15] B. Bérard, K. Chatterjee, and N. Sznajder, "Probabilistic opacity for Markov decision processes," *Information Processing Letters*, vol. 115, no. 1, pp. 52–59, 2015.

[16] J. Chen, M. Ibrahim, and R. Kumar, "Quantification of secrecy in partially observed stochastic discrete event systems," *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 1, pp. 185–195, 2017.

[17] M. Ahmadi, B. Wu, H. Lin, and U. Topcu, "Privacy verification in POMDPs via barrier certificates," in *IEEE Conference on Decision and Control*. IEEE, 2018, pp. 5610–5615.

[18] B. Wu and H. Lin, "Privacy verification and enforcement via belief abstraction," *IEEE Control Systems Letters*, vol. 2, no. 4, pp. 815–820, 2018.

[19] D. Chistikov, A. S. Murawski, and D. Purser, "Asymmetric distances for approximate differential privacy," in *30th International Conference on Concurrency Theory (CONCUR 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[20] C. Cortes, M. Mohri, and A. Rastogi, "Lp distance and equivalence of probabilistic automata," *International Journal of Foundations of Computer Science*, vol. 18, no. 04, pp. 761–779, 2007.

[21] R. B. Lyngsø and C. N. Pedersen, "The consensus string problem and the complexity of comparing hidden markov models," *Journal of Computer and System Sciences*, vol. 65, no. 3, pp. 545–569, 2002.

[22] S. Kiefer, "On computing the total variation distance of hidden markov models," *arXiv preprint arXiv:1804.06170*, 2018.

[23] S. Haesaert, S. Soudjani, and A. Abate, "Verification of general markov decision processes by approximate similarity relations and policy refinement," *SIAM Journal on Control and Optimization*, vol. 55, no. 4, pp. 2333–2367, 2017.

[24] A. Lavaei, S. Soudjani, and M. Zamani, "From dissipativity theory to compositional construction of finite Markov decision processes," in *Proceedings of the 21st ACM International Conference on Hybrid Systems: Computation and Control*, 2018, pp. 21–30.