Probabilistic Current-State Opacity is Undecidable

Anooshiravan Saboori and Christoforos N. Hadjicostis

***Abstract*— Increasing concerns about security and privacy in applications of discrete event systems have led to various notions of *opacity* for systems that are modeled as (possibly non-deterministic) finite automata with partial observation on their transitions. Specifically, a system is current-state opaque if the entrance of the system state to a given set of *secret states* remains opaque (uncertain), until the system leaves this set of secret states, to an intruder who observes system activity through some projection map. While this notion has been shown useful for security requirements in many applications (including encryption using pseudo-random generators and coverage properties in sensor networks), it does not provide a quantifiable *measure* for characterizing the security of a given system. In this paper, we extend this framework to systems that can be modeled as probabilistic finite automata, obtaining in the process the probability of observing sequences of observations that violate current-state opacity. We then introduce and analyze the notion of *probabilistic current-state opacity* which can be used to provide a measure of a given system's opacity. We show that verifying probabilistic current-state opacity is undecidable in general, though it can become decidable in specific settings.**

I. Introduction

Motivated by the increased reliance of many applications on shared cyber-infrastructures (ranging from defense and banking to health care and power distribution systems), various notions of *security and privacy* have received considerable attention from researchers. A number of such notions focus on characterizing the *information flow* from the system to the intruder [1]. *Opacity* falls in this category and aims at determining whether a given system's *secret* behavior (i.e., a subset of the behavior of the system that is considered critical and is usually represented by a predicate) is kept opaque to outsiders [2], [3]. More specifically, this requires that the intruder (modeled as an observer of the system's behavior) never be able to establish the truth of the predicate.

In our earlier work [3], [4], we consider opacity with respect to predicates that are state-based. More specifically, we consider a scenario where we are given a discrete event system (DES) that can be modeled as a non-deterministic finite automaton (NFA) with partial observation on its transitions (captured via a natural projection map). Allowing the initial state of the system to be (partially) unknown, we define the secret behavior of the system as the *evolution* of the system's state (at some point in time) to a subset of the set of secret states S, which is taken to be known and fixed over the length of the observation. The intruder is assumed to have full knowledge of the system model and be able to track the observable transitions in the system via the observation of the associated labels. One particular notion of interest from [3] is *current-state opacity* which requires that the membership of the system current-state to the set of secret states S remain opaque until the system enters a state outside S.

Current-state opacity as defined in [3] does not consider the likelihood of violating the current-state opacity requirement; instead it simply reports whether a given system is opaque or not. This binary outcome (that the system is current-state opaque or not) might be inadequate in cases where different behaviors in the system (which may or may not violate current-state opacity) have unequal likelihood of occurring. As an example, consider a mobile agent moving in a 2-dimensional grid covered (partially) by a network of sensors (with some cells of the grid covered by possibly more than one sensor). Using the sensor information, the network can be used to obtain the set of possible locations of the agent at any given point in time. In this context, the notion of current-state opacity characterizes whether all trajectories that the agent can follow do not expose that she/he is currently visiting certain strategic (secret) locations (cells). However, consider the following two extreme scenarios: (i) the probability that the agent follows a trajectory that exposes its current strategic (secret) location (cell) is 10^{-6}, (ii) the probability that the agent follows a trajectory that exposes its current strategic (secret) location (cell) is $1 - 10^{-6}$. In both scenarios, the system will be classified as a system that violates current-state opacity (as defined in [3]) despite the huge discrepancy in the likelihood of witnessing a sequence of observations that reveals that the system current-state belongs to the set of secret states.

Motivated by the absence of likelihood information in existing work on current-state opacity, in our work [5], we consider a probabilistic setting and devise appropriate opacity measures to quantify opacity. More specifically, we consider a scenario where we are given a stochastic discrete event system (SDES) that can be modeled as a probabilistic finite automaton (PFA) with partial observation on its transitions; assuming that the initial-state distribution vector

This material is based upon work supported in part by the National Science Foundation (NSF), under NSF ITR Award 0426831 and NSF CNS Award 0834409. The research leading to these results has also received funding from the European Commission (EC) 7^{th} Framework Programme (FP7/2007-2013) under grant agreements INFSO-ICT-223844 and PIRG02-GA-2007-224877. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of NSF or EC.

A. Saboori is with the Coordinated Science Laboratory, and the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign. C. N. Hadjicostis is with the Department of Electrical and Computer Engineering, University of Cyprus, and with the Coordinated Science Laboratory, and the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign. Corresponding author's address: 75 Kallipoleos Avenue, P.O. Box 20537, 1678 Nicosia, Cyprus. E-mail: chadjic@ucy.ac.cy.

ISBN 978-963-311-370-7

is known, we define the notion of T-*almost current-state opacity* and analyze its verification complexity. T-almost current-state opacity considers the probability of violating current-state opacity along all sequences of events of length k, and requires that this probability lies below a threshold for all possible lengths ($k = 0, 1, 2, \ldots$).

T-almost current-state opacity can only quantify current-state opacity when the given system is *not* current-state opaque. This can be a limitation in some scenarios. As an example, consider the sensor network in the previous example and assume that based on a particular set of readings from the sensor network, the intruder/observer is (at least) $1 - 10^{-6}$ confident that the current location (cell) of the agent is a strategic (secret) location (i.e., given the sequence of sensor readings, the probability that the current state of the system is within the set of secret states is at least $1 - 10^{-6}$ — but not one). In such case, the system will be classified as current-state opaque (because the intruder cannot be absolutely certain about the membership of the system current-state to the set of secret states) and hence T-almost current state opaque. However, this oversight raises several questions regarding the appropriateness of the notion of current-state opacity and T-almost current-state opacity for applications where the *confidence* of the intruder can serve as a measure of security. An example of areas where such confidence concerns have been considered are anonymity protocols [6], [7]. Such systems consist of a set of users whose actions generate associated outputs that are observed by intruders who then try to guess the identity of the originator of the action. The goal of anonymity protocols is to hide the origin (user) for certain actions in the system despite the observed outputs. In order to describe the level of protection offered by an anonymity protocol, three notions have been defined in [7]: (i) *beyond suspicion*: when the sender of a message appears no more likely to be the originator than any other potential sender, (ii) *probable innocence*: when the sender appears no more likely to be the originator of the message than not to be the originator, and (iii) *possible innocence*: when there is a *non-trivial* probability that the sender is someone else.

In this paper, we introduce the notion of *probabilistic current-state opacity* by considering the probability that the system current-state lies in the set of secret states (and, hence, the confidence of the intruder) given any possible sequence of observations in the system, and by requiring that this probability lies below a threshold for all possible behavior in the system. We motivate this notion using anonymity protocols for web transactions. Finally, we establish that, in general, the verification of probabilistic current-state opacity is an undecidable problem.

The work in this paper is related to existing security work in stochastic settings. In particular, [8], [9] focus on systems whose state transition functions are captured by probability distributions and define the *advantage* Adv of the intruder for a given sequence of observations ω as the increase in the conditional probability $Pr(S|\omega)$ that the system current-state resides within the set of secret states S (given ω is observed) compared to the probability $Pr(S|\epsilon)$ before any observation is available, i.e., $Adv(\omega) = Pr(S|\omega) - Pr(S|\epsilon)$. The authors of [8], [9] define three notions of opacity depending on the value of function Adv as follows: (i) *strict opacity* requires $Adv(\omega) = 0$ for all possible sequences ω of observations; (ii) *cryptographic opacity* requires $Adv(\omega)$ to be negligible for all possible observations; and (iii) *plausible opacity* requires $Adv(\omega) \neq 1 - Pr(S|\epsilon)$ for any possible sequence ω of observations (i.e., $Pr(S|\omega) \neq 1$ for any sequence ω of observations) which is equivalent to the notion of current-state opacity [3]. The authors of [8], [9] consider the verification of all three notions using exhaustive computation which is possible because the given (stochastic) system is assumed to have *finite behavior* (i.e., the length of possible state sequences in the system is finite). The notion of probabilistic current-state opacity introduced in this paper essentially extends the notion of cryptographic opacity in [8], [9] to probabilistic finite automata (with possibly infinite behavior). We show that even for this special class of probabilistic systems, verifying probabilistic current-state opacity is an undecidable problem. This implies that verifying cryptographic opacity for stochastic systems with possibly infinite behavior is also undecidable.

II. Preliminaries and Background

Let Σ be an alphabet and denote by Σ^* the set of all finite-length strings of elements of Σ, including the empty string ϵ. For a string t, $|t|$ denotes the length of t, whereas for a set X, $|X|$ denotes its cardinality. A language $L \subseteq \Sigma^*$ is a subset of finite-length strings from Σ^*. For a string ω, the *prefix-closure* of ω is defined as $\overline{\omega} = \{t \in \Sigma^* | \ \exists s \in \Sigma^* : ts = \omega\}$ where ts denotes the concatenation of strings t and s [10].

A non-deterministic finite automaton (NFA) is denoted by $G = (X, \Sigma, \delta, X_0)$, where $X = \{0, 1, \ldots, N-1\}$ is the set of states, Σ is the set of events, $\delta : X \times \Sigma \to 2^X$ (where 2^X is the power set of X) is the non-deterministic state transition function, and $X_0 \subseteq X$ is the set of possible initial states. The function δ can be extended from the domain $X \times \Sigma$ to the domain $X \times \Sigma^*$ in the routine recursive manner: $\delta(i, ts) := \bigcup_{j \in \delta(i,t)} \delta(j, s)$, for $t \in \Sigma$ and $s \in \Sigma^*$, with $\delta(i, \epsilon) := i$. The behavior of NFA G is captured by $L(G) := \{s \in \Sigma^* \mid \exists i \in X_0 \{\delta(i, s) \neq \emptyset\}\}$.

In general, only a subset Σ_{obs} of the events can be observed. Typically, one assumes that Σ can be partitioned into two sets, the set of observable events Σ_{obs} and the set of unobservable events Σ_{uo} (so that $\Sigma_{obs} \cap \Sigma_{uo} = \emptyset$ and $\Sigma_{obs} \cup \Sigma_{uo} = \Sigma$). The natural projection $P : \Sigma^* \to \Sigma_{obs}^*$ can be used to map any trace executed in the system to the sequence of observations associated with it. This projection is defined recursively as $P(ts) = P(t)P(s),\ t \in \Sigma, s \in \Sigma^*$, with $P(t) = t$ if $t \in \Sigma_{obs}$ and $P(t) = \epsilon$ if $t \in \Sigma_{uo} \cup \{\epsilon\}$ [10].

Upon observing some string $\omega \in \Sigma_{obs}^*$ (sequence of observations), the state of the system might not be identifiable uniquely due to the lack of knowledge of the initial state, the partial observation of events, and/or the non-deterministic behavior of the system. We denote the set of states that

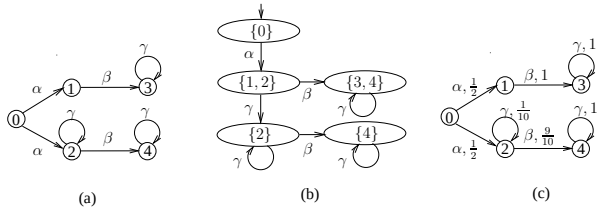

Fig. 1. (a) Non-deterministic automaton G; (b) current-state estimator $G_{0,obs}$ associated with G assuming that $X_0 = \{0\}$; (c) PFA H.

the system might reside in *given that ω was observed* as the current-state estimate. The *current-state estimator* (or observer) is a deterministic finite automaton (DFA) $G_{0,obs}$ which captures these estimates and can be constructed as follows [10]. Each state of $G_{0,obs}$ is associated with a unique subset of states of the original NFA G (so that there are at most $2^{|X|} = 2^N$ states). The initial state of $G_{0,obs}$ is associated with X_0, representing the fact that the initial state could be any state in X_0. At any state Z of the estimator ($Z \subseteq X$), the next state upon observing an event $\alpha \in \Sigma_{obs}$ is the unique state of $G_{0,obs}$ associated with the set of states that can be reached from (one or more of) the states in Z with a string of events that generates the observation α. The details of the construction of the current-state estimator are illustrated in the following example; more details can be found in [10].

Example 1: Consider the NFA G in Figure 1-a with initial state $X_0 = \{0\}$. Assuming that $\Sigma_{obs} = \Sigma = \{\alpha, \beta, \gamma\}$, the current-state estimator $G_{0,obs}$ in Figure 1-b is constructed as follows. Starting from the initial state $\{0\}$ and observing α, the current state is any of the states in $\{1, 2\}$; at this new state, the set of possible transitions is the union of all possible transitions for each of the states in $\{1, 2\}$. Following this procedure, $G_{0,obs}$ can be completed as in Figure 1-b. By convention, one does not include the state of the current-state estimator that corresponds to the empty set of state estimates and can be reached from the initial state of the current-state estimator via a sequence of observations that cannot be generated by G; we have followed this convention in Figure 1-b. ■

A *stochastic* discrete event system (SDES) is modeled in this paper as a *probabilistic finite automaton* (PFA) $H = (X, \Sigma, p, \pi_0)$ where $X = \{0, 1, \ldots, N-1\}$ is the set of states; Σ is the set of events; $p(i', \alpha|i)$ is the state transition probability defined for $i, i' \in X$, and $\alpha \in \Sigma$, as the probability that event α occurs and the system transitions to state i' given that the system is in state i. When $p(i', \alpha|i) = 0$, state i' is not reachable from state i via event α (hence, in the diagram representing the given PFA, we do not draw such transitions); finally, π_0 is the initial-state probability distribution vector [11]. In the given SDES, the probabilities of all transitions from any state i add up to at most 1, i.e., $\sum_{i' \in X} \sum_{\alpha \in \Sigma} p(i', \alpha|i) \leq 1$. When the probabilities do not add up to exactly 1, this means that the system remains in state i with the remainder probability. In this paper, we assume that for all $i \in X$, we have

$$\sum_{i' \in X} \sum_{\alpha \in \Sigma} p(i', \alpha|i) = 1. \tag{A1}$$

An example of a probabilistic finite automaton that satisfies this assumption can be seen in Figure 1-c.

We denote the current-state probability distribution vector after observing the sequence of observations $\omega = \alpha_0 \alpha_1 \ldots \alpha_n$ by π_ω. For convenience, we index the elements of the $N \times 1$ column vector π starting from 0 so that the i^{th} element of π, denoted by $\pi(i)$, $0 \leq i \leq N-1$, corresponds to the probability that the current state is state i. Also, we denote the 1-norm of the vector π by $\|.\|$, i.e., $\|\pi\| = \sum_{i=0}^{N-1} |\pi(i)|$. It can be shown [11] that for $\omega = \alpha_0 \alpha_1 \ldots \alpha_n$

$$\pi_\omega = \frac{A_{\alpha_n} A_{\alpha_{n-1}} \ldots A_{\alpha_0} \pi_0}{\|A_{\alpha_n} A_{\alpha_{n-1}} \ldots A_{\alpha_0} \pi_0\|}, \tag{1}$$

where A_{α_i}, $0 \leq i \leq n$, is the $N \times N$ state transition probability matrix associated with the observable event $\alpha_i \in \Sigma_{obs}$. Specifically, the $(j,k)^{th}$ element of matrix A_{α_i} denotes the conditional probability $p(j, \alpha_i|k)$ that the system transitions to system state j and generates the observation α_i given that it is in state k. Note that N is the number of states of the PFA H, and $\|.\|$ is the vector 1-norm. For example, for PFA in Figure 1-c we have $A_\alpha = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$, $A_\beta = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \frac{9}{10} & 0 & 0 \end{bmatrix}$, and $A_\gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{10} & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$.

Given a PFA $H = (X, \Sigma, p, \pi_0)$ we can associate to it a unique NFA $G = (X, \Sigma, \delta, X_0)$ [11] where the state transition function $\delta : X \times \Sigma \to 2^X$ is defined for $i, i' \in X$, $\alpha \in \Sigma$ as $i' \in \delta(i, \alpha)$ if $p(i', \alpha|i) > 0$, and the set of possible initial states is defined as $X_0 = \{i | \pi_0(i) > 0\}$. Figure 1-a depicts the NFA G associated with the PFA H in Figure 1-c.

III. Problem Formulation and Motivation

A. Formulation

In this section, we first recall the formal definition of current-state opacity from [3].

Definition 1 (Current-State Opacity): Given an NFA $G = (X, \Sigma, \delta, X_0)$, a projection map P with respect to the set of observable events Σ_{obs} ($\Sigma_{obs} \subseteq \Sigma$), and a set of secret states $S \subseteq X$, automaton G is current-state opaque with respect to S and P (or $(S, P, 0)$ current-state opaque), if $\forall t \in \Sigma^*, \forall i \in X_0$

$$\{\delta(i,t) \neq \emptyset, \delta(i,t) \subseteq S\} \Rightarrow$$
$$\{\exists s \in \Sigma^*, \exists i' \in X_0 \{P(s) = P(t), \delta(i', s) \nsubseteq S\}\}. \quad ■$$

Current-state opacity requires that the membership of its current state to the set S remain opaque (uncertain) for all possible behavior in the system. One can check whether a

system is current-state opaque by constructing the current-state estimator and by verifying that no (nonempty) current-state estimate lies entirely within the set of secret states [3]. The following example illustrates this construction; the details can be found in [3].

Example 2: Consider the non-deterministic finite automaton G depicted in Figure 1-a with $\Sigma_{obs} = \Sigma = \{\alpha, \beta, \gamma\}$. Suppose that $S = \{4\}$ and $X_0 = \{0\}$. From Figure 1-b which depicts the current-state estimator associated with G, we see that the current-state estimate $\{4\}$, which lies entirely within the set of secret states, is reachable via sequences of the form $\alpha\gamma\gamma^*\beta\gamma^*$. This violates current-state opacity, thus NFA G is not current-state opaque with respect to $S = \{4\}$ and P. However, it is not hard to verify that NFA G is current-state opaque with respect to $S = \{3\}$ and P (because no state of the state estimator is associated with a nonempty subset of $\{3\}$). ■

Although the notion of current-state opacity characterizes sequences of observations (or state trajectories) that reveal that the system current-state is within the set of secret states, it does not quantify the degree of opacity of the given system. Motivated by such limitations, in this paper, we use probabilistic metrics to quantify the degree of opacity of the given system.

Definition 2 (Probabilistic Current-State Opacity): Given a PFA $H = (X, \Sigma, p, \pi_0)$, a projection map P with respect to the set of observable events Σ_{obs} ($\Sigma_{obs} \subseteq \Sigma$), and a set of secret states $S \subseteq X$, PFA H is probabilistically current-state opaque with respect to S, P, and θ (or $(S, P, 0, \theta)$-probabilistically current-state opaque), if

$$\forall \omega \in \Sigma^*_{obs} : \|\pi_\omega(S)\| - \|\pi_0(S)\| \leq \theta, \qquad (2)$$

where $\|.\|$ denotes the vector 1-norm and $\pi_0(S)$ ($\pi_\omega(S)$) denotes the vector of elements of the initial-state (current-state) probability distribution vector π_0 (π_ω) indexed by the elements in the set S (see the discussion before and after Equation (1)). ■

Note that $\pi_\omega(S)$ in Definition 2 denotes the confidence of the intruder that the current state of the system after observing the sequence of observations ω lies in the set of secret states. Similarly, $\pi_0(S)$ denotes the initial confidence of the intruder that the current state of the system, before any observation occurred in the system, lies in the set of secret states. Therefore, according to Definition 2, PFA H is probabilistically current-state opaque if, after the observations, the increase in the confidence of the intruder that the current state of the system lies in the set of secret states is less than the threshold θ. Note that (2) is also verified for $\omega \in \Sigma^*_{obs}$ that cannot be possibly be generated by the system (i.e., there does not exist $s \in L(G)$, where G is the NFA associated with PFA H, such that $P(s) = \omega$); this does not matter, however, because in those cases, (2) trivially holds. Also note that if H is $(S, P, 0, \theta)$-probabilistically current-state opaque, for some $\theta < 1 - \|\pi_0(S)\|$, then its associated NFA G is $(S, P, 0)$ current-state opaque; however, the converse is not necessarily true. The following example illustrates Definition 2.

Example 3: In this example we show that the PFA H depicted in Figure 1-c is not probabilistically current-state opaque with respect to $S = \{3\}$, P, and $\theta = 0.5$ assuming that $\Sigma_{obs} = \Sigma = \{\alpha, \beta, \gamma\}$ and that $\pi_0 = [\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}]'$. For this, we evaluate the current-state probability distribution vector $\pi_{\alpha\beta\gamma}$ after observing the sequence of observations $\alpha\beta\gamma$. It can be shown that $\pi_{\alpha\beta\gamma} = [0, 0, 0, \frac{100}{109}, \frac{9}{109}]' = [0, 0, 0, 0.92, 0.08]'$. Using the notation of Definition 2, $\pi_{\alpha\beta\gamma}(S) = [0.92]$ and $\pi_0(S) = [0.2]$, and therefore $\|\pi_{\alpha\beta\gamma}(S)\| - \|\pi_0(S)\| = 0.72$ which is larger than 0.5 and hence PFA H is not $(\{3\}, P, 0, 0.5)$-probabilistically current-state opaque. Note that, as discussed in Example 2, the NFA G associated with PFA H is $(\{3\}, P, 0)$ current-state opaque. ■

B. Motivational Example

There are many application areas where state-based notions of opacity can be used to characterize security or other requirements of interest. In our earlier work [4], we studied conditions under which the seed (or more generally the past state or the current state) of a pseudo-random generator in a cryptographic protocol can be compromised, and showed that this problem can be formulated and analyzed more precisely using the notions of current and initial-state opacity [4]; when probabilities are available, this problem can be analyzed using the framework of this paper. In the sequel, we discuss another example in the context of anonymity for web transactions which was initially studied in [7] and can be analyzed using the probabilistic current-state opacity introduced in this paper.

Example 4: "Crowds" is a system which aims to protect anonymity on the world-wide-web. Given a set of users $\mathcal{I} = \{0, 1, \ldots, I\}$, crowds partitions them into geographically diverse groups (crowds) $\mathcal{C} = \{C_1, C_2, \ldots, C_C\}$. Each crowd issues requests to the server on behalf of its members (users) [7]. A user $i \in \mathcal{I}$ is represented in a crowd by a process $j_i \in \mathcal{J} = \{j_0, j_1, \ldots, j_I\}$ called a "jondo". Any request coming from the user's browser is sent directly to the jondo.

Upon receiving a request from the browser of user i for server S, jondo j_i initiates the establishment of a random path of jondos (possibly in more than one crowd) $j_{i_1} \rightarrow j_{i_2} \rightarrow \ldots \rightarrow j_{i_m} \rightarrow S$, $i_k \in \mathcal{I}$, $1 \leq k \leq m$, that carries its user's transactions to and from the intended web server S. More precisely, when jondo j_{i_1} receives the request, it flips a biased coin to determine whether or not to forward the request to another jondo. If the result is to forward, then the jondo selects a random[1] jondo j_{i_2} (possibly in a different crowd) and forwards the request to it; otherwise the jondo submits the request to the end server for which the request was destined. In this way, each request travels from the user's browser, through a number of jondos (in possibly more than one crowd), and finally to the end server. Due to physical limitations, a jondo can only send messages to neighboring jondos (possibly in different crowds). We model this via an NFA. The non-deterministic structure of the automaton models the randomness in the selection of

[1]We describe precisely the process of selecting jondos next.

jondos. An example of such automaton is shown in Figure 2-a assuming that the set of jondos is $\mathcal{J} = \{j_0, j_1, j_2, j_3\}$. We say $j_i \rightarrow j_{i'}$ when there is a transition in the NFA that starts from j_i and ends in $j_{i'}$.

Communication between any two jondos is encrypted using a key known only to the two of them. The intruder is modeled as a set of collaborating (corrupted) jondos $\mathcal{J}^{\mathcal{C}} \subseteq \mathcal{J}$ which can belong to more than one crowd. Figure 2-b depicts a Crowds system where the set of jondos (users) $\mathcal{J} = \{j_0, j_1, j_2, j_3\}$ is partitioned into crowds $\{C_1, C_2\}$ with $C_1 = \{j_0, j_2\}$ and $C_2 = \{j_1, j_3\}$, and jondo j_0 is corrupted, i.e., $\mathcal{J}^{\mathcal{C}} = \{j_0\}$.

We assume that the messages sent from each crowd to a *different* crowd do not reveal the identity of the originator of that message unless this message is sent from a collaborating jondo to another collaborating jondo. This implies that the collaborators can only identify the originator of a message when the originator either belongs in their *own crowd* or is a collaborating jondo. Note that, due to the encryption of messages, the collaborators can *only* identify the jondo (within their own crowd) from which they *immediately* received the message and not the preceding jondos in the path (unless the preceding jondo is collaborating). We model this by assigning label pc to a transition that starts from a non-collaborating jondo j_p in crowd C and ends in a collaborating jondo j_c in the same crowd C. For example, in Figure 2-b, the transition between non-collaborating jondo j_2 in crowd C_1 and collaborating jondo j_0 in crowd C_1 has label 20. Also we assign label cc' to a transition that starts from a collaborating jondo j_c and ends in another collaborating jondo $j_{c'}$ (independent from the crowd that they belong to). Since messages received in crowd C' from a non-collaborating jondo j_p in a different crowd C do not reveal the originator, we assign the label $\{pc | j_p \in C, j_p \in \mathcal{J} - \mathcal{J}^{\mathcal{C}}, j_c \in C', j_c \in \mathcal{J}^{\mathcal{C}}, j_p \rightarrow j_c, C, C' \in \mathcal{C}, C \neq C'\}$ to any transition that starts from any non-collaborating jondo j_p in crowd C and ends in a collaborating jondo j_c in crowd C'. For example, in Figure 2-b, transitions that start from the non-collaborating jondos j_1 and j_3 in crowd C_2 and end in collaborating jondo j_0 in crowd C_1 have the label $\{10, 30\}$.

Each collaborating jondo j_c can identify the jondo j' it forwards the message to if either jondo j' belongs to the same crowd as jondo j or jondo j' is also a collaborating jondo. This implies that a collaborating jondo j_c in crowd C who sends a message to crowd C' cannot determine the (non-collaborating) jondo who receives this message[2]. To model this, we assign label $\{cp | j_c \in C, j_c \in \mathcal{J}^{\mathcal{C}}, j_p \in C', j_p \in \mathcal{J} - \mathcal{J}^{\mathcal{C}}, j_c \rightarrow j_p, C, C' \in \mathcal{C}, C \neq C'\}$ to any transition that starts from the collaborating jondo j_c in crowd C and ends in any non-collaborating jondo j_p in crowd C'. In Figure 2-b, the labeling of the transition between the collaborating jondo j_0 in C_1 and non-collaborating jondo j_1 in C_2 follows this rule. Finally, transitions between non-collaborating jondos are not observable to intruders, and in Figure 2-b we denote these transitions with dashed lines.

As mentioned before, each jondo flips a biased coin to decide whether or not to forward the request to another jondo (possibly in a different crowd). To keep the discussion simple, we assume that each jondo decides to forward the message to each of its neighbors and the server with equal probability. For example, based on the automaton in Figure 2-a, jondo j_0 can either forward the message to the server S or to one of its neighbors j_1 and j_2; therefore, it forwards the message to the server S with probability $\frac{1}{3}$ and to each of its neighbors with probability $\frac{1}{3}$. In Figure 2-c we depict the PFA which models this. Note that for brevity, in Figure 2-c we have renamed the labels of the events in Figure 2-b as follows: $02 \leftrightarrow \alpha$, $20 \leftrightarrow \beta$, $01 \leftrightarrow \sigma$, and $\{10, 30\} \leftrightarrow \gamma$. Also, we assigned label δ_{uo} to all transitions associated with dashed lines to indicate that these transitions are unobservable.

The security of Crowds has been studied in [7] by obtaining the probability that the identity of the originating jondo is revealed to a set of collaborating jondos. This calculation is carried assuming that every jondo can directly send messages to any other jondo in Crowds (in this case the graph in Figure 2-a is *fully connected*) and that the probability that a jondo forwards a message to another jondo is the same for all jondos. Moreover, the authors in [7] only analyze the steady-state behavior of the system and as a result, the performance of the system becomes dependent only on the number of jondos, the number of collaborating jondos, and the probability of forwarding the message.

The framework we introduce in this paper can be used to analyze the security properties of Crowds (and similar protocols) in a more systematic fashion; more specifically, general jondo communication graphs (Figure 2-a) are allowed and the probability of forwarding can be different for various jondos. In this way, the notions of opacity that are introduced in this paper can be used to study both the transient and steady-state behavior of the system. For example, one of the security questions that arises in this context is whether in a given topology the collaborating jondos can determine the (non-collaborating) jondos that contributed to a path (on which at least one collaborator jondo occupies a position) with certain confidence (described via a threshold on the probability that the non-collaborating jondo contributed to the path). By defining the set of secret states as certain subsets of the set of non-collaborating jondos and by taking the set of observable events as the labels of the transitions associated with solid lines in Figure 2-b, the notion of probabilistic current-state opacity introduced in this paper can be used to answer such questions. ■

IV. Verification of Probabilistic Current-State Opacity

In the sequel, we show that the problem of verifying probabilistic current-state opacity (P-CURRENT) is undecidable using a reduction from the *emptiness problem* (EMPTY) for

[2]This can be achieved with the existence of a representative jondo in each crowd which receives the messages sent to that crowd and randomly distributes them among the members of that crowd (subject to interconnection constraints between jondos).

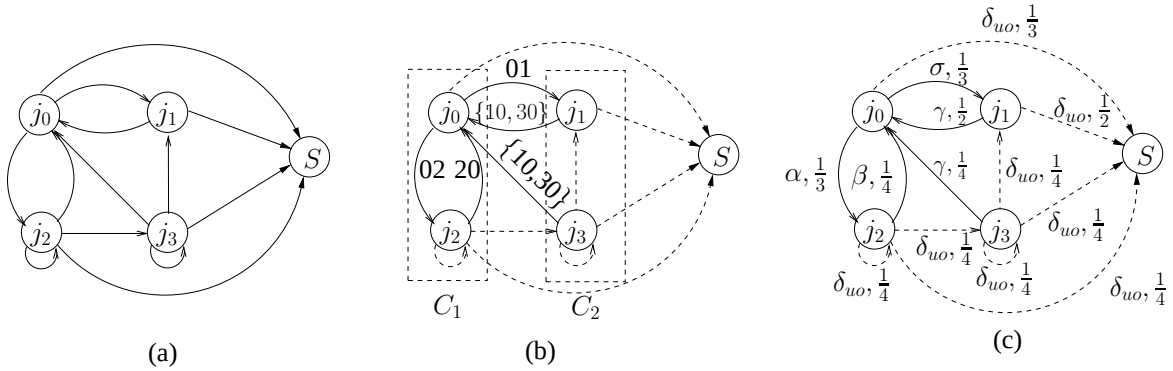

Fig. 2. (a) Example of an automaton describing the possible paths between a set of jondos $\mathcal{J} = \{j_0, j_1, j_2, j_3\}$ and a server S; (b) Automaton that models the Crowds system assuming that users are partitioned into crowds $\{C_1, C_2\}$, with $C_1 = \{j_0, j_2\}$ and $C_2 = \{j_1, j_3\}$, and where jondo j_0 is assumed corrupted, i.e., $\mathcal{J}^{\mathcal{C}} = \{j_0\}$; (c) Probabilistic automaton that models the Crowds system in part (b) assuming that each jondo decides to forward the message to its neighbors and the server with equal probability. For brevity, we rename the events in part (b) as follows $02 \leftrightarrow \alpha$, $20 \leftrightarrow \beta$, $01 \leftrightarrow \sigma$, and $\{10, 30\} \leftrightarrow \gamma$.

probabilistic automata, which is known to be undecidable [11], [12].

Definition 3 (Emptiness Problem (EMPTY)): Consider a PFA $H = (X, \Sigma, p, \pi_0)$, where for each $i \in X$, $\alpha \in \Sigma$

$$\sum_{i' \in X} p(i', \alpha | i) = 1, \tag{A2}$$

and π_0 is a column vector of size $|X|$ whose entries are all zeros except for single entry with value 1 (corresponding to the *initial state*). Given a constant θ, $0 < \theta < 1$, and a row vector y of size $|X|$ whose entries are all zeros except for single entry with value 1 (corresponding to the *final state*), does there exist $\omega = \alpha_0 \alpha_1 \dots \alpha_n \in \Sigma^*$ (for some $n \geq 0$) such that $y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0 > \theta$? ■

Definition 4 (P-CURRENT problem): Given a PFA $H = (X, \Sigma, p, \pi_0)$, a projection map P with respect to the set of observable events Σ_{obs} ($\Sigma_{obs} \subseteq \Sigma$), and a set of secret states $S \subseteq X$, is PFA H probabilistically current-state opaque with respect to S, P, and θ? ■

Theorem 1: The P-CURRENT problem is undecidable. ■

Proof: We introduce a polynomial-time algorithm which reduces every instance of the EMPTY problem to an instance of the P-CURRENT problem, and since the EMPTY problem is undecidable, this proves that the P-CURRENT problem is also undecidable. Consider an instance of the EMPTY problem: Given a PFA $H = (X, \Sigma, p, \pi_0)$ where π_0 is a column vector of size $|X|$ whose entries are all zeros except for entry i that has value 1 (i.e., $\pi_0(i) = 1$), and a row vector y of size $|X|$ whose entries are all zeros except for entry j that has value 1 (i.e., $y(j) = 1$). In the sequel, using PFA H, we construct a PFA $\hat{H}$ and show that PFA is probabilistic current-state opaque if and only there exists $\omega = \alpha_0 \alpha_1 \dots \alpha_n \in \Sigma^*$ (for some $n \geq 0$) such that $y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0 > \theta$. In this way, we reduce each instance of the EMPTY problem to an instance of the P-CURRENT problem.

We construct $\hat{H} = (X, \Sigma, \hat{p}, \pi_0)$ as follows: for $i, i' \in X$ and $\alpha \in \Sigma$, we define

$$\hat{p}(i', \alpha | i) = \frac{1}{|\Sigma|} p(i', \alpha | i).$$

Using Assumption (A2), this implies that

$$\sum_{i' \in X} \sum_{\alpha \in \Sigma} \hat{p}(i', \alpha | i) = 1,$$

i.e., PFA $\hat{H}$ satisfies Assumption (A1). Denote by $\hat{A}_\alpha$ the transition matrix that is constructed using $\hat{p}(i', \alpha | i)$. It is not hard to see that for all $\alpha \in \Sigma$,

$$\hat{A}_\alpha = \frac{1}{|\Sigma|} A_\alpha. \tag{3}$$

Now, define $S = \{j\}$ where state j is the index of the nonzero entry of the vector y and define $\Sigma_{obs} := \Sigma$. In order to avoid trivial cases, assume that $i \neq j$ (but this assumption can be easily removed by introducing dummy state if $i = j$). The initial state i is unique (characterized by the only nonzero entry of the vector π_0), which implies that

$$\|\pi_0(S)\| = 0 \tag{4}$$

(because the initial state i is different from the secret state j). Also, $\|\pi_\omega(S)\|$ denotes the cumulative probability of being in the set of secret states after observing $\omega = \alpha_0 \alpha_1 \dots \alpha_n$ and since state j is the only secret state, it is equal to the probability that system $\hat{H}$ resides in state j after observing ω. Putting all these together, we have

$$\begin{aligned}
\|\pi_\omega(S)\| - \|\pi_0(S)\| &= \|\pi_\omega(S)\| \text{ (By 4)} \\
&= \frac{y \hat{A}_{\alpha_n} \hat{A}_{\alpha_{n-1}} \dots \hat{A}_{\alpha_0} \pi_0}{\|\hat{A}_{\alpha_n} \hat{A}_{\alpha_{n-1}} \dots \hat{A}_{\alpha_0} \pi_0\|} \text{ (By (1))} \\
&= \frac{y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0}{\|A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0\|} \text{ (By (3))} \\
&= y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0,
\end{aligned}$$

where the last equation follows from the fact that

$$\|A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0\| = 1 \tag{5}$$

for any $\omega = \alpha_n \alpha_{n-1} \dots \alpha_0$. Note that (5) follows from Assumption (A2) [11]. Now, the EMPTY problem requires deciding whether

$$\exists \omega \in \Sigma^* : y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0 > \theta,$$

or, equivalently,

$$\forall \omega \in \Sigma^* : y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0 \leq \theta. \quad (6)$$

By previous discussion we have that $\|\pi_\omega(S)\| = y A_{\alpha_n} A_{\alpha_{n-1}} \dots A_{\alpha_0} \pi_0$. Therefore, deciding (6) is equivalent to deciding whether

$$\forall \omega \in \Sigma^* : \|\pi_\omega(S)\| - \|\pi_0(S)\| \leq \theta,$$

which is equivalent to deciding whether PFA $\hat{H}$ is probabilistically current-state opaque. This completes the proof. ■

V. Conclusion

In this paper, we define, analyze, and characterize the notion of probabilistic current-state opacity as an extension of the notion of current-state opacity [3] to stochastic systems. The notion of current-state opacity requires that the membership of the current state of a given NFA G to a set of secret states S remains opaque to intruders. We define probabilistic current-state opacity by limiting the maximum increase in the probability that the system current-state lies in the set of secret states given the observations (conditional probability) compared to the case when no observation is available (prior probability). We establish that the verification of probabilistic current-state opacity is an undecidable problem.

One important extension for future work is to develop restrictive condition under which the verification of probabilistic current-state opacity becomes decidable. One trivial case is when the number of strings in the automaton G associated with PFA H that reach a state in the set of secret states S from a state in the set of initial states X_0, is finite. In this case, one can enumerate all possible sequences of observations and obtain the maximum probability of residing in the set of secret states, hence, verifying probabilistic current-state opacity.

We are also interested in developing similar frameworks for other state-based notions of opacity, such as initial-state opacity [4]. Initial-state opacity requires that the membership of the initial state of the system to the set of secret states be kept opaque for the duration of the operation of the system and can potentially be handled in a similar fashion.

References

[1] R. Focardi and R. Gorrieri, "A taxonomy of trace–based security properties for CCS," in *Proc. of the 7th Workshop on Computer Security Foundations*, June 1994, pp. 126–136.

[2] J. Bryans, M. Koutny, L. Mazare, and P. Ryan, "Opacity generalised to transition systems," *International Journal of Information Security*, vol. 7, no. 6, pp. 421–435, November 2008.

[3] A. Saboori and C. N. Hadjicostis, "Notions of security and opacity in discrete event systems," in *Proc. of the 46th IEEE Conference on Decision and Control*, December 2007, pp. 5056–5061.

[4] ——, "Verification of initial–state opacity in security applications of DES," in *Proc. of the 9th International Workshop on Discrete Event Systems*, May 2008, pp. 328–333.

[5] ——, "Opacity verification in stochastic discrete event systems," submitted to the 49th IEEE Conference on Decision and Control.

[6] D. Chaum, "The dining cryptographers problem: unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, no. 1, pp. 65–75, March 1988.

[7] M. K. Reiter and A. D. Rubin., "Crowds: Anonymity for web transactions," *ACM Transactions on Information and System Security*, vol. 1, no. 1, pp. 66–92, November 1998.

[8] Y. Lakhnech and L. Mazaré, "Probabilistic opacity for a passive adversary and its application to Chaum's voting scheme," Verimag, Tech. Rep. TR-2005-4, 2005. [Online]. Available: http://www-verimag.imag.fr/index.php?page=techrep-list

[9] R. Janvier, Y. Lakhnech, and L. Mazaré, "Completing the picture: soundness of formal encryption in the presence of active adversaries," in *Proceedings of 14th European Symposium on Programming*, ser. Lecture Notes in Computer Science, vol. 3444, April 2005, pp. 172–185.

[10] C. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, 2008.

[11] A. Paz, *Introduction to Probabilistic Automata*. Academic Press, Inc., 1971.

[12] V. D. Blondel and V. Canterini, "Undecidable problems for probabilistic automata of fixed dimension," *Theory of Computing Systems*, vol. 36, no. 2, pp. 231–245, March 2008.