Automatica 80 (2017) 162–171

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

A new approach for the verification of infinite-step and K-step opacity using two-way observers☆

Xiang Yin [a,1], Stéphane Lafortune [b]

[a] Department of Automation, Shanghai Jiao Tong University, Shanghai 200240, China

[b] Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:
Received 13 June 2016
Received in revised form
3 November 2016
Accepted 27 January 2017
Available online 28 March 2017

Keywords:
Discrete event systems
Infinite-step opacity
K-step opacity
Two-way observer

ABSTRACT

In the context of security analysis for information flow properties, where a potentially malicious observer (intruder) tracks the observed behavior of a given system, infinite-step opacity (respectively, K-step opacity) holds if the intruder can never determine for sure that the system was in a secret state for any instant within infinite steps (respectively, K steps) prior to that particular instant. We present new algorithms for the verification of the properties of infinite-step opacity and K-step opacity for partially-observed discrete event systems modeled as finite-state automata. Our new algorithms are based on a novel separation principle for state estimates that characterizes the information dependence in opacity verification problems, and they have lower computational complexity than previously-proposed ones in the literature. Specifically, we propose a new information structure, called the *two-way observer*, that is used for the verification of infinite-step and K-step opacity. Based on the two-way observer, a new upper bound for the delay in K-step opacity is derived, which also improves previously-known results.

1. Introduction

We investigate the verification of an important information-flow property called *opacity* that arises in security analysis of networked cyber and cyber–physical systems. We adopt a discrete-event framework, where the system under consideration is modeled as a partially-observed finite-state automaton and the security properties of interest for opacity are captured in terms of a set of secret states of the automaton. In this manner, the focus of the analysis is on the event-driven dynamics of the cyber or cyber–physical system of interest, captured in a discrete transition structure with unobservable events, and on the resulting observation properties during system operation. The system is said to be opaque if the secret cannot be revealed to an intruder that is potentially malicious. The intruder is modeled as an external observer that knows the transition structure of the system but can only observe part of the system's behavior.

To the best of our knowledge, the notion of opacity was initially introduced in Mazaré (2004), where it was motivated by the analysis of cryptographic protocols. It was then extended to the framework of Discrete Event Systems (DES) in Bryans, Koutny, Mazaré, and Ryan (2008) and Bryans, Koutny, and Ryan (2005). Several notions of opacity have been studied in order to capture different types of privacy requirements in the context of DES; among them we mention language-based opacity (Lin, 2011), current-state opacity (Saboori & Hadjicostis, 2007), initial-state opacity (Saboori & Hadjicostis, 2013), initial-and-final-state opacity (Wu & Lafortune, 2013), K-step opacity (Falcone & Marchand, 2014; Saboori & Hadjicostis, 2011b), and infinite-step opacity (Saboori & Hadjicostis, 2012b). If a given system is not opaque, then one is also interested in enforcing opacity. The opacity enforcement problem has been studied extensively under different enforcement mechanisms, e.g., using supervisory control (Badouel, Bednarczyk, Borzyszkowski, Caillaud, & Darondeau, 2007; Darondeau, Marchand, & Ricker, 2014; Dubreil, Darondeau, & Marchand, 2010; Saboori & Hadjicostis, 2012a; Takai & Kumar,

☆ This work was partially supported by NSF grants CCF-1138860 (Expeditions in Computing project ExCAPE: Expeditions in Computer Augmented Program Engineering) and CNS-1421122, and by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA. This work was done when the first author was at the University of Michigan. The material in this paper was presented at the 13th International Workshop on Discrete Event Systems, May 30–June 1, 2016, Xi'an, China. This paper was recommended for publication in revised form by Associate Editor Christoforos Hadjicostis under the direction of Editor Christos G. Cassandras.

E-mail addresses: xiangyin@umich.edu (X. Yin), stephane@umich.edu (S. Lafortune).

[1] Fax: +1 7347638041.

http://dx.doi.org/10.1016/j.automatica.2017.02.037
0005-1098/

2009; Takai & Oka, 2008; Yin & Lafortune, 2015b, 2016b), using dynamic observers (Cassez, Dubreil, & Marchand, 2012; Yin & Lafortune, 2015a; Zhang, Shu, & Lin, 2015), using insertion or edit functions (Wu & Lafortune, 2014; Wu, Raman, Lafortune, & Seshia, 2016), and using run-time techniques (Falcone & Marchand, 2014). Most of the above-mentioned works assume that the system is modeled as a finite-state automaton. Recently, the notion of opacity was extended to other classes of system models, including timed systems (Cassez, 2009), Petri nets (Bryans et al., 2005; Tong, Li, Seatzu, & Giua, 2016), pushdown systems (Kobayashi & Hiraishi, 2013), and stochastic systems (Bérard, Chatterjee, & Sznajder, 2015; Keroglou & Hadjicostis, 2013; Saboori & Hadjicostis, 2014). Several applications of opacity have also been investigated in the literature; see, e.g., Saboori and Hadjicostis (2011a) and Wu, Sankararaman, and Lafortune (2014). The reader is referred to the recent survey (Jacob, Lesage, & Faure, 2016) for more references on this active research area.

In this paper, we study the *verification* problem for the two notions of *infinite-step* opacity and *K-step* opacity. Current-state opacity requires that the secret not be revealed to the intruder based on the *current* state estimate. In contrast, infinite-step opacity requires that the secret not be revealed for any instant along the entire observation trajectory up to the present time, *based on the observations up to the current time*. Similarly, K-step requires that the secret not be revealed within K steps prior to the current instant, *based on the observations up to the current time*. It was shown in Wu and Lafortune (2013) that language-based opacity, initial-state opacity, and current-state opacity are "equivalent" in the sense that they can be mapped to one another in polynomial time. However, infinite-step and K-step opacity appear to be incomparable with the above notions, for the following reason. Whereas current-state opacity only depends on the current state estimate of the system, infinite-step and K-step opacity allow to do *smoothing*, i.e., to improve state estimation for *earlier* time instants, using observations up to the *present* time. Therefore, infinite-step and K-step opacity are fundamentally different from current-state opacity, language-based opacity, and initial-state opacity.

One of the motivations for studying infinite-step opacity and K-step opacity is that these two notions are very useful in privacy applications. For example, privacy is an important issue in Location-Based Services (LBS); see, e.g., Gruteser and Grunwald (2003). In LBS applications, the user may want to hide some of her crucial location information (e.g., visiting a bank or a hospital). However, this information may be revealed to an intruder located at the LBS server that keeps tracking the user's queries. Therefore, a formal methodology is needed in order to verify this privacy issue in LBS. It was shown in Wu et al. (2014) that verifying whether or not the user can always hide her *current* crucial location can be formulated as a current-state opacity verification problem. However, in some cases, the user may also want that the intruder never be able to infer that she was at a crucial place at some particular instant in the past (e.g., visited bank two days ago). Clearly, current-state opacity is not sufficient to capture this requirement, since the intruder may be able to use future observations to improve its knowledge about the user's location at some particular instant. However, this requirement can be captured using the notions of infinite-step or K-step opacity.

The notions of infinite-step opacity and K-step opacity were initially studied in Saboori and Hadjicostis (2011b, 2012b), respectively. More specifically, in Saboori and Hadjicostis (2011b), two different approaches for the verification of K-step opacity were proposed; both of these approaches have the same computational complexity of $O((|E_o|+1)^K \times |E_o| \times 2^{|X|})$, where X and E_o are the set of states and the set of observable events of the system, respectively. For infinite-step opacity, a verification algorithm of complexity of $O(|E_o| \times 2^{|X|} \times 2^{|X|^2})$ was provided in Saboori and Hadjicostis (2012b).

In this paper, we propose new approaches for the verification of infinite-step opacity and K-step opacity. Specifically, our contributions are summarized as follows.

- We provide a new characterization for the delayed state estimate, which is referred to as the *separation principle*. This result reveals that the information needed in the infinite-step (K-step) opacity verification problem can be decomposed into two mutually independent parts where each of them can be computed individually and effectively.
- We propose a novel information structure called the Two-Way Observer (TW-observer) in order to capture and represent in a single structure the two parts of independent information described by the separation principle.
- Based on the TW-observer, we present a new approach for the verification of infinite-step opacity. This approach results in a new algorithm that has complexity of $O(|E_o| \times 2^{|X|} \times 2^{|X|})$, compared with $O(|E_o| \times 2^{|X|} \times 2^{|X|^2})$ for the previous approach (Saboori & Hadjicostis, 2012b).
- We show that our proposed approach can also be used to verify the notion of K-step opacity, resulting in an algorithm of complexity of $O(\min\{2^{|X|}, |E_o|^K\} \times |E_o| \times 2^{|X|})$. This approach is based on the notion of K-reduced TW-observer that we introduce. The previous algorithm for verifying K-step opacity had a complexity of $O((|E_o|+1)^K \times |E_o| \times 2^{|X|})$ (Saboori & Hadjicostis, 2011b). Therefore, our new algorithm leads to considerable improvement in verification complexity when K is relatively large.
- Using the TW-observer, we provide a new upper bound in the K-step opacity problem. We show that a system is infinite-step opaque if and only if it is $(2^{|X|}-2)$-step opaque. This also improves upon the previous upper bound of $2^{|X|^2}-2$ derived in Saboori and Hadjicostis (2011b).
- Overall, the TW-observer provides a unified and more efficient framework for the verification of infinite-step and K-step opacity, as previous approaches require different techniques for verifying these properties.

In the definitions of infinite-step opacity and K-step opacity, it is required that the intruder cannot infer that the system was at a secret state for any *specific instant* in the past. However, in some cases, it is possible that the intruder knows that the system has visited a secret state in the past, although it cannot tell the specific instant (in terms of the number of steps) the secret state was visited. We call a system *trajectory-based* infinite-step (respectively, K-step) opaque if this scenario does *not* occur; examples for trajectory-based opacity can be found in Saboori and Hadjicostis (2011b, 2012b). Therefore, infinite-step (K-step) opacity is also referred to as *non-trajectory-based* infinite-step (K-step) opacity. Trajectory-based K-step opacity is referred to as K-step *strong* opacity in Falcone and Marchand (2014), where a verification algorithm is provided. Whether one needs to use the trajectory-based notions or the non-trajectory-based notions is application dependent. In this paper, we will focus on the non-trajectory-based notions.

The remaining sections of this paper are organized as follows. Sections 2 and 3 present the system model and the definitions of the opacity properties considered in this paper, respectively. In Section 4, the above-mentioned separation principle is investigated. Section 5 describes the structure of the proposed two-way observer and discusses its properties. Section 6 presents the new approach for the verification of infinite-step opacity. In Section 7, we show how to use the two-way observer to verify K-step opacity. Finally, we conclude the paper in Section 8.

Preliminary and partial versions of some of the results in this paper are presented, without proofs, in Yin and Lafortune (2016a).