

Using Subobservers to Synthesize Opacity-Enforcing Supervisors

Richard Hugh Moulton¹, Behnam Behinaein Hamgini¹, Zahra Abedi Khouzani¹,
Rômulo Meira-Góes², Fei Wang³ and Karen Rudie^{1,4,5}

October 8, 2021

Abstract

In discrete-event system control, the worst-case time complexity for computing a system’s observer is exponential in the number of that system’s states. This results in practical difficulties since some problems require calculating multiple observers for a changing system, e.g., synthesizing an opacity-enforcing supervisor. Although calculating these observers in an iterative manner allows us to synthesize an opacity-enforcing supervisor and although methods have been proposed to reduce the computational demands, room exists for a practical and intuitive solution. Here we extend the subautomaton relationship to the notion of a subobserver and demonstrate its use in reducing the computations required for iterated observer calculations. We then demonstrate the subobserver relationship’s power by simplifying state-of-the-art synthesis approaches for opacity-enforcing supervisors under realistic assumptions.

1 Introduction

Discrete-event systems (DES) research has recently focused on opacity, a system property that ensures that secret information cannot be distinguished from non-secret information. The literature includes methods for verifying if a system is opaque, namely, whether a hostile agent could—based on its observations—determine whether or not the system is in a secret state, or whether or not the system has generated a secret event sequence. A natural question to ask if a system is not opaque is “how can we make this system opaque?” One approach is to alter which events are visible to an adversary, however this is rarely in the power of the system designer. An alternative approach is to use supervisory control to disable event occurrences so that the controlled system is opaque [10].

Researchers have examined opacity-enforcing supervisors under various assumptions [1, 2, 7, 22, 23, 26, 27]. The naïve method of producing such a supervisor starts by constructing a plant automaton representing the system and an observer automaton capturing the adversary’s view of the system based on which events the adversary can observe. The standard method for constructing an observer automaton relies on converting a nondeterministic finite automaton to a deterministic finite automaton (NFA-to-DFA conversion) [9], which is known to require asymptotic computing

¹Department of Electrical and Computer Engineering, Queen’s University, Kingston, Canada

²Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, USA

³College of Information Science and Engineering, Huaqiao University, Xiamen, China

⁴School of Computing, Queen’s University, Kingston, Canada

⁵Ingenuity Labs Research Institute, Queen’s University, Kingston, Canada

time that is exponential in the number of states of the plant. Although supervisory control can then restrict plant behaviour to enforce opacity, a control-aware adversary will then be able to revise its plant estimates. As a result, repeated plant evolution necessitates repeated NFA-to-DFA operations. In this paper we present a method for synthesizing an opacity-enforcing supervisor that relies on only one NFA-to-DFA operation on the states of the uncontrolled plant.

Although we motivate our approach with the problem of plant evolution due to supervisory control, it is more generally a method for tracking the joint behaviour of a plant and an observer through multiple steps of plant evolution, without computing an observer at each intermediary stage. In addition supervisory control, a plant may evolve over time for other reasons as well. For example, a discrete-event process may be a naturally time-varying dynamic discrete-event system [8], may require online control [6], or may be best modelled by a time-varying automaton because more becomes known about the process over time. Alternatively, a discrete-event process may be controlled by a number of decentralized agents with different views of the plant, [13], where controls enacted by one agent may allow another agent to update its own estimates of the plant.

The power of our approach is that it leads to efficient and intuitive computations. In the remainder of this paper we review previous approaches to the opacity control problem, we introduce the subobserver property and show that it can be used to incrementally refine the joint behaviour of a plant and an observer, and, finally, we demonstrate how this approach can be used to efficiently synthesize an opacity-enforcing supervisor.

2 Problem Definition and Related Works

To lay the theoretical foundation for our contribution, we introduce the formalisms of DES, define the opacity control problem, and review the literature for related concepts and methods.

2.1 Discrete-event systems

DES are used to model processes that are discrete, asynchronous and potentially nondeterministic [16]. The discrete nature of these systems is captured with an automaton: at any given time the system is in a specific state and the system moves between states via event-based transitions.

Mathematically, a DFA G is a 5-tuple $(Q, \Sigma, \delta_G, q_0, Q_m)$ where Q is the finite set of states in the system, Σ is the finite alphabet of events that can occur, $\delta_G : Q \times \Sigma \rightarrow Q$ is a partial function, $q_0 \in Q$ is the initial state, and $Q_m \subseteq Q$ is the set of marked states. From a control theory perspective the automaton G represents the plant, or underlying process, to be controlled. The alphabet Σ represents the events that can occur within the plant and can be divided into disjoint sets by considering whether individual events are controllable or uncontrollable (Σ_c and Σ_{uc} respectively) and whether they are observable (Σ_o and Σ_{uo} respectively). The transition function δ_G leads inductively from single events to the language $L(G)$, which is the plant's full range of behaviour. Because δ_G is usually a partial function, we use the notation $\delta_G(q, \sigma)!$ to denote that a specific transition $\delta_G(q, \sigma)$ is defined. We will also use the notation $\langle q', \sigma, q \rangle_G$ to denote that $\delta_G(q', \sigma) = q$. The set of marked states may represent the completion of a process or, in our case, a set of distinctive states that we wish to keep indistinguishable to an adversary. In this paper we call these "secret states."

A particular DFA representation of a language induces a binary relation called the equiresponse relation (Definition 1, from Lin and Wonham [14]) on Σ^* , the Kleene closure of Σ . The equiresponse relation for a DFA associates strings with the states they lead to, so that two strings are related

if they lead to the same state. Because the equireponse relation is a right congruence on Σ^* , it is also an equivalence relation and, like all equivalence relations, defines a partition. We note that different DFA representations of a language will induce different equireponse relations and there is no canonical equireponse relation.

Definition 1 (Equireponse relation). Given an automaton $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, the **equireponse relation** of G is a right congruence $eq(G)$ on Σ^* and defined by

$$s \equiv s' \pmod{eq(G)} \iff \delta_G(q_0, s) = \delta_G(q_0, s')$$

and if $\delta_G(q_0, s)$ is undefined then $\delta_G(q_0, s) = \delta_G(q_0, s')$ if and only if $\delta_G(q_0, s')$ is undefined as well.

2.1.1 Observing discrete-event systems

When an automaton contains unobservable events, an observer automaton can be used to capture an agent's beliefs about the system. This observer automaton can be computed by transforming the original automaton, first by substituting the empty string ϵ for all transitions whose events are not in Σ_o [4, p. 76], and then performing NFA-to-DFA conversion [4, pp. 87-90]. Note that we do not use the standard marking associated with NFA-to-DFA conversion (whereby a subset labelling a state in the DFA is marked if any element in the subset is marked in the original NFA); rather, we mark a state in the DFA only if *all* elements of the subset label are marked states in the original NFA. As will be discussed in Section 5.1, this way of marking states in the observer automaton will distinguish states in which an adversary can be sure that the original system is in a secret state.

Throughout this paper we denote the transformation of an automaton to an observer automaton with respect to the alphabet Σ_o by $T_{\Sigma_o}(\cdot)$; we denote the associated change in languages generated by these automata as the projection $P_{\Sigma_o}(\cdot)$ (Definition 2, from Cassandras and Lafortune [4, p. 57]).

Definition 2 (Projection of strings). For an alphabet Σ and another alphabet $\Sigma_o \subseteq \Sigma$, we define the **projection** $P_{\Sigma_o} : \Sigma^* \rightarrow \Sigma_o^*$ recursively as follows:

$$\begin{aligned} P_{\Sigma_o}(\epsilon) &= \epsilon \\ P_{\Sigma_o}(\sigma) &= \begin{cases} \sigma & \text{if } \sigma \in \Sigma_o \\ \epsilon & \text{if } \sigma \in \Sigma \setminus \Sigma_o \end{cases} \\ P_{\Sigma_o}(s\sigma) &= P_{\Sigma_o}(s)P_{\Sigma_o}(\sigma) \quad \text{for } s \in \Sigma^*, \sigma \in \Sigma \end{aligned}$$

The projection operation implicitly defines an equivalence relation over Σ^* , where two strings s and s' are related if they share a projection, $P_{\Sigma_o}(s) = P_{\Sigma_o}(s')$. Each cell in the associated partition contains those strings whose projections are equal. Different projections will define different partitions on Σ^* and we say that one partition refines another partition if it continues to distinguish between all elements that the latter partition distinguishes between (Definition 3) [17, Ch. 9].

Definition 3 (Refinement of a partition). Let ρ_1 and ρ_2 be binary relations on a set X . We say that ρ_1 **refines** ρ_2 , or that ρ_2 is coarser than ρ_1 , denoted $\rho_1 \leq \rho_2$ if

$$(\forall x, y \in X) (x, y) \in \rho_1 \implies (x, y) \in \rho_2.$$

2.1.2 Controlling discrete-event systems

In addition to observing DES processes, we can also control them. In DES the controller that modifies which events can occur in a plant G is called the *supervisor* and is formally a pair $\mathcal{S} = (S, \phi)$. Here, $S = (X, \Sigma, \xi, x_0, X_m)$ is a DFA that shares an alphabet with G but has a different state space, transition function and set of marked states [16]. The second component of the pair, $\phi : X \rightarrow \Gamma$, is a state feedback map that maps supervisor states, $x \in X$, to control patterns $\gamma \in \Gamma$,

$$\gamma := \phi(x) \in \{0, 1\}^\Sigma.$$

These control patterns indicate whether the supervisor will enable an event, denoted by 1, or disable it, denoted by 0, with the requirement that no uncontrollable event can be disabled. The controlled DES process is then represented by the supervisor composed with the plant and is denoted by

$$S/G = Ac(X \times Q, \Sigma, \xi \times \delta, (x_0, q_0), X \times Q_m)$$

where Ac is a function that restricts an automaton to its accessible portion: those states that can be reached from the initial state via a string in the automaton's language [16]. A control pattern can also be defined using the supervisor's transition function, where an event σ is enabled in $\phi(x)$ if $\xi(x, \sigma)!$ and it is disabled otherwise. A great amount of work exists in the literature for determining whether a system is controllable, whether a system is observable, and how to synthesize supervisors [4, 24].

2.2 The notion of opacity

When we talk about the opacity property of a system, we use the plain language meaning of an opaque system. That is, we say that a system is opaque if an observer of the system is unable to unambiguously determine some characteristic of that system. Mazaré formalized this notion of opacity for computer science and showed that under this definition it is decidable whether or not a particular system property is opaque [15]. Formulating this notion of opacity for transition systems, Bryans et al. used epistemic logic and possible world models to frame opacity as the inability of an observer to determine the truth of a particular predicate related to the system [3].

Building on this work, Saboori and Hadjicostis characterized a number of state-based opacities for DES including initial state, k -step, and infinite step opacities [19, 20, 21]. Alternatively, Lin considered opacity to be a broader concept than Bryans et al. and defined opacity in terms of languages, distinguishing between strong and weak opacity in terms of the degree to which strings in one language are confused by an observer with strings in another language [12].

More recently, Wu and Lafortune unified four common kinds of opacity: language-based, initial-state, current-state, and initial-and-final-state. They considered a problem setup where adversaries had full knowledge of the system's structure but only partial observations of its behaviour. Their key result was a set of polynomial-time algorithms for transforming each of these four kinds of opacity into the other under the authors' problem formulation. The one exception to this is that language-based opacity cannot be transformed into initial-state opacity if either the secret or non-secret language is not prefix-closed [25]. Without loss of generality, therefore, we consider in this paper the problem of enforcing current-state opacity, which we define in line with Wu and Lafortune [25].

Definition 4 (Current-state opacity). Given a plant $G = (Q, \Sigma, \delta, q_0, Q_m)$, a projection P_{Σ_o} , and a disjoint sets of secret states $Q_S \subseteq Q$ and non-secret states $Q_{NS} \subseteq Q$, $Q_S \cap Q_{NS} = \emptyset$, G is

current-state opaque if for every string that leads from an initial state to a secret state, there exists a string that leads from an initial state to a non-secret state and these two strings have identical projections.

$$\begin{aligned} &\forall s \in L(G) \text{ such that } \delta(q_0, s) \in Q_S, \\ &\exists t \in L(G) \text{ such that } \delta(q_0, t) \in Q_{NS} \text{ and } P_{\Sigma_o}(s) = P_{\Sigma_o}(t) \end{aligned}$$

2.3 Enforcing opacity through supervisory control

It is natural to want to ensure that a given system satisfies our well-defined notion of opacity. This problem has been addressed in the literature in many different forms, in this paper we name it the *opacity control problem*. This problem is defined by a set of general parameters [3, 7]:

1. the set of events that are observable and controllable for the supervisor (Σ_s and Σ_c respectively);
2. the set of events that are observable by the adversary (Σ_a);
3. the secret and type of opacity to enforce; and
4. the adversary's knowledge of any supervisory control policy.

Since the supervisor and adversary each have a set of observable events, we label these Σ_s and Σ_a respectively to avoid ambiguity. The most general formulation of the opacity control problem assumes no relationship between Σ_s , Σ_c , and Σ_a . Other reasonable relationships to assume between these three sets is that the adversary sees a subset of the events that the supervisor does, $\Sigma_a \subseteq \Sigma_s$, that the supervisor can only control events that it can see, $\Sigma_c \subseteq \Sigma_s$, and that the system contains only events that are controllable and observable by the supervisor, $\Sigma_c = \Sigma_s = \Sigma$. Regardless of the specific relationship between event sets, the most principled way to approach the opacity control problem is to consider an adversary that is aware of the supervisor's control policy and thereby avoid any appeals to security through obscurity. In this paper we use the formulation given in Problem 1.

Problem 1 (Opacity Control Problem). Given a DFA $G = (Q, \Sigma, \delta, q_0, Q_m)$ where Q_m is the set of secret states, an alphabet of events the supervisor can observe $\Sigma_s \subseteq \Sigma$ and an alphabet of events the supervisor can control $\Sigma_c \subseteq \Sigma_s$. Given an supervisor-aware adversary and an alphabet of events that it can observe $\Sigma_a \subseteq \Sigma_s$. Find a supervisor \mathcal{S} such that the closed-loop behaviour of the plant composed with the supervisor, \mathcal{S}/G , is current-state opaque with respect to the alphabet observable by the adversary, Σ_a , the set of secret states $Q_S = Q_m$, and the set of non-secret states $Q_{NS} = Q \setminus Q_m$.

The naïve solution for this formulation is to iteratively calculate the adversary's view of the plant, apply supervisory control, and then check the opacity of the controlled system until a fixed point is reached. This approach has been noted by multiple authors and dismissed as very computationally expensive due to the fact that an observer automaton must be calculated for every iteration [7, 21].

2.4 Other approaches to the Opacity Control Problem

The opacity control problem has been addressed in the literature under different assumptions. In some works the adversary is unaware of the supervisor, while in others the adversary has complete knowledge of the supervisor’s control policy. We consider the latter case, since this introduces the requirement to iteratively calculate the supervisor’s control policy and the adversary’s estimate of the plant; Table 1 summarizes the characteristics of these works.

Reference	Approach	Event Sets	Opacity Type
[1]	Sub-lattice of kernels	$\Sigma_a \subseteq \Sigma_s$ $\Sigma_c = \Sigma$	Language-based
[22]	Supremal closed, controllable, and opaque sublanguage	$\Sigma_a \subseteq \Sigma_s$ $\Sigma_c \subseteq \Sigma_s$ $\Sigma_s = \Sigma$	Language-based
[2]	Supremal controllable, observable, and opaque language	$\Sigma_a = \Sigma_s$ $\Sigma_c \subseteq \Sigma_s$	Language-based
[7]	Condensed state estimates	$\Sigma_a \subseteq \Sigma_s$ $\Sigma_c \subseteq \Sigma_s$	Current-state
[26]	All-Inclusive Controller for Opacity	$\Sigma_a \subseteq \Sigma_s$	Current-state
[27]	All-Enforcement Structure	$\Sigma_a = \Sigma_s$	Current-state
[23]	Augmented I-observer	–	Current-state
This paper	Plant behaviour composed with the adversary’s view	$\Sigma_a \subseteq \Sigma_s$ $\Sigma_c \subseteq \Sigma_s$	Current-state

Table 1: Comparing opacity control problem formulations with a supervisor-aware adversary. Σ is the set of all events, Σ_c is the supervisor’s set of controllable events, Σ_s is the supervisor’s set of observable events, and Σ_a is the adversary’s set of observable events.

Addressing language-based opacity, Badouel et al. presented an iterative method to design a supervisor that guaranteed concurrent opacity assuming there are no uncontrollable events in the plant [1]. Takai and Oka proposed a method with one-step convergence under the strong assumption that, for any pair of indistinguishable strings, any uncontrollable and observable event that can occur after one string can occur after the other [22]. Finally, Ben-Kalefa and Lin iteratively enforced opacity and controllability for the language generated by the plant [2]. This is the naïve approach applied to languages and still requires the construction of a supervisor to enforce the final language produced.

Considering current-state opacity, Yin and Lafortune presented the *All-Inclusive Controller for Opacity* [26] and the *All Enforcement Structure* [27], which both embed a game between supervisor and plant in a bipartite transition structure where the supervisor’s goal is to enforce the specified property and the plant’s is to violate the specified property. More similar to our approach to the opacity control problem are *condensed state estimates*, [7], and the *augmented I-observer*, [23].

Dubreil et al. showed that solving the opacity control problem under full observation induces

a general solution to the opacity control problem when $\Sigma_a \subseteq \Sigma_s$ [7]. Their approach to enforcing opacity through supervisory control produces an automaton whose states track the plant’s actual state as well as condensed state estimates, which contain the set of states that the adversary believes the plant could be in if the last transition it observed was the last transition that occurred in the plant. From the condensed state estimates, the supervisor then reasons about “loosing paths,” which are the traces from the current state that would lead to the disclosure of a secret [7].

Our approach to the opacity control problem is inspired by the condensed state estimates method, which we improve upon in two ways. First, the parallel composition between the plant and an observer represent the adversary’s beliefs about the plant more intuitively than condensed state estimates do. Second, our representation of the adversary’s beliefs removes the need for unrolling condensed state estimates to determine whether a particular string will disclose a system secret.

Finally, Tong et al.’s *augmented I-observer* is the parallel composition between the plant and the observer representing the adversary’s view of the plant [23]. The supremal G-opaque sublanguage – a language-based property equivalent to current-state opacity – is computed and a supervisor synthesized to enforce it as a specification [23, Definition 7]. This process iterates until the system’s generated language is G-opaque, i.e., the system is current-state opaque. The augmented I-observer method allows Σ_a and Σ_s to be incomparable, but Tong et al. restricted the majority of their discussion to non-supervisor-aware adversaries. With a supervisor-aware adversary, this generality comes with the additional computational cost of updating both the plant and augmented I-observer at every step of the iteration. With the reasonable assumption that $\Sigma_a \subseteq \Sigma_s$, our method is able to simply refine the parallel composition of the plant and adversary’s view instead.

3 Relating Observers: Subautomata versus Subobservers

The key to our method is that we are able to relate each successive observer automaton to the one that came before it. To begin, when supervisory control is applied to enforce a regular-language specification for an automaton G , the controlled system can be modelled as another automaton, G' . We begin, therefore, with an understanding that there is a relationship between G and G' . This is the subautomaton relationship, a general relationship that goes beyond supervisory control [4, p. 86].

Definition 5 (Subautomaton). We say that $G' = (Q', \Sigma, \delta_{G'}, q'_0, Q'_m)$ is a **subautomaton** of $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, denoted by $G' \sqsubseteq G$, if

$$\delta_{G'}(q'_0, s) = \delta_G(q_0, s) \quad \forall s \in L(G')$$

The subautomaton relationship is a strong form of correspondence between two automata. It allows us to reason, for example, that the states of G' are a subset of those in G , that the initial state in G' is the same as in G , and, more generally, that we can match the states in G' with a corresponding state in G [4, p. 86]. This relationship can be very helpful when reasoning about the language produced by each automaton or about termination conditions for certain algorithms.

When the specification of legal behaviour is a subautomaton of G , the resulting automaton G' will also be a subautomaton of G . Although the idea of a specification is not directly applicable to the opacity control problem, we demonstrate in Section 5.3 how our proposed method inherently ensures that supervisory control produces a subautomaton at each iteration.

Beyond the relationship between the uncontrolled and controlled plant, we are interested in how an adversary’s observations of a system evolve as supervisory control is enacted. We would like to




Figure 1: We wish to characterize the relationship that exists between H' and H .

know whether or not we can establish a relationship between the observer automata H' and H , transformations of G' and G respectively with respect to the same alphabet of observable events (Figure 1). Does the subautomaton relationship survive the transformation; is H' a subautomaton of H ? More formally, is it true, given automata $G, G', H := T_{\Sigma_o}(G)$ and $H' := T_{\Sigma_o}(G')$, that

$$G' \sqsubseteq G \implies H' \sqsubseteq H?$$

Unfortunately this statement is not true because H' and H will have their states coming from $2^{Q'}$ and 2^Q respectively, allowing for the possibility that states in H' will not properly match up with states in H (Figure 2). Although the transition graph of any subautomaton is necessarily a subgraph of the parent automaton's transition graph, the same is not true for automaton related by the subobserver property since states may be “split” from H to H' . Even if we avoid this, the state labels in H and H' are semantically meaningful, namely they encode the adversary's belief about the state of the system, and these semantics would be lost if we allowed, for example, the state labelled $\{5\}$ to be matched with the state labelled $\{5, 6\}$.

All is not lost, however. Given our insistence that state labels carry semantic meaning, we can relax the definition of subautomaton to allow for this and produce a relationship that is suitable to relate two observers together. We call this the subobserver relationship, Definition 6.

Definition 6 (Subobserver). We say that $G' = (Q', \Sigma, \delta_{G'}, q'_0, Q'_m)$ is a **subobserver** of $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, denoted $G' \sqsubseteq G$, if

$$\delta_{G'}(q'_0, s) \subseteq \delta_G(q_0, s) \quad \forall s \in L(G').$$

Because we parallel the definition of a subautomaton in our definition of a subobserver it immediately follows that all subautomaton relationships are also subobserver relationships and that the converse is not true. We claim that the subobserver property does complete the relationships between automata G, G', H , and H' seen in Figure 1.

Before proving this result formally, we begin by establishing some technical notes. To begin, throughout our proofs we will define automata in terms of G . Specifically, G' is a subautomaton of G and $H' := T_{\Sigma_o}(G')$ and $H := T_{\Sigma_o}(G)$ are observer automata with our characteristic marking.

$$\begin{aligned} G &= (Q, \Sigma, \delta, q_0, Q_m) & H &= T_{\Sigma_o}(G) = (2^Q, \Sigma_o, \delta_H, q_{0h}, Q_{mh}) \\ G' &= (Q', \Sigma, \delta', q_0, Q'_m) & H' &= T_{\Sigma_o}(G') = (2^{Q'}, \Sigma_o, \delta_{H'}, q_{0h'}, Q_{mh'}) \end{aligned}$$

Next, we define a state's ϵ -reach as the set of states that can be reached from it via the empty string (Definition 7) [4, p. 71].




Figure 2: Counterexample showing that H' is not a subautomaton of H . Initial states are denoted by a small arrow; marked states are denoted with a double ring. Observable events have a solid arrow and unobservable events have a dashed arrow. Controllable events have marks across their arrows.

Definition 7 (ϵ -reach). For an automaton G and projection P_{Σ_o} , we call the set of states that can be reached from state q via a string whose projection is ϵ as the ϵ -reach of q . Formally,

$$\epsilon R_G(q) := \{q' \in Q \mid \exists s \in \Sigma^*, P_{\Sigma_o}(s) = \epsilon, \delta(q, s) = q'\}.$$

By definition a state is always in its own ϵ -reach, $q \in \epsilon R(q)$.

We also show that if $G' \sqsubseteq G$ then a state's ϵ -reach in G' is a subset of that state's ϵ -reach in G .

Lemma 1 (The ϵ -reach of any state in a subautomaton is a subset of that state's ϵ -reach in the parent automaton). *Given automata $G = (Q, \Sigma, \delta_G, q_0, Q_m)$ and $G' = (Q', \Sigma, \delta_{G'}, q_0', Q'_m)$. Then*

$$G' \sqsubseteq G \implies \epsilon R_{G'}(q) \subseteq \epsilon R_G(q) \forall q \in Q'.$$

Proof. To begin, $G' \sqsubseteq G$ implies that $Q' \subseteq Q$. We can therefore determine the ϵ -reach of state q

with respect to both G and G' for all $q \in Q'$.

$$\begin{aligned}
\epsilon R_{G'}(q) &= \{q' \in Q' \mid \exists s \in \Sigma^*, P_\Sigma(s) = \epsilon, \delta_{G'}(q, s) = q'\} \\
&\subseteq \{q' \in Q' \mid \exists s \in \Sigma^*, P_\Sigma(s) = \epsilon, \delta_{G'}(q, s) = q'\} \\
&\quad \cup \{q' \in Q' \mid \exists s \in \Sigma^*, P_\Sigma(s) = \epsilon, \delta_G(q, s) = q', \neg \delta_{G'}(q, s)\} \\
&\quad \cup \{q' \in Q \setminus Q' \mid \exists s \in \Sigma^*, P_\Sigma(s) = \epsilon, \delta_G(q, s) = q'\} && \text{by set inclusion} \\
&= \{q' \in Q \mid \exists s \in \Sigma^*, P_\Sigma(s) = \epsilon, \delta_G(q, s) = q'\} \\
&= \epsilon R_G(q)
\end{aligned}$$

□

We are now ready to prove that the relationship between the observer automata H and H' (the dashed line in Figure 1) is the subobserver relationship defined in Definition 6.

Theorem 1 (The transformation of a subautomaton is itself a subobserver). *Given automata G and G' and their respective transformations, $H := T_{\Sigma_o}(G)$ and $H' := T_{\Sigma_o}(G')$. Then*

$$G' \sqsubseteq G \implies H' \tilde{\sqsubseteq} H.$$

Proof. The proof proceeds by induction.

Base case: $s = \epsilon$, $|s| = 0$.

$$\begin{aligned}
q_{0h'} &= \epsilon R_{G'}(q_0) \\
&\subseteq \epsilon R_G(q_0) && \text{by Lemma 1} \\
&= q_{0h}
\end{aligned}$$

Inductive Hypothesis: Suppose $s \in L(H')$, $|s| = n$, and $\bar{Q}' = \delta_{H'}(q_{0h'}, s) \subseteq \delta_H(q_{0h}, s) = \bar{Q}$.

Inductive Step: Now we show that for $s\sigma \in L(H')$, $|s\sigma| = n + 1$, then $\delta_{H'}(q_{0h'}, s\sigma) \subseteq \delta_H(q_{0h}, s\sigma)$.

We proceed by translating the set of states $\delta_{H'}(q_{0h'}, s\sigma)$ into set builder notation and then adding well-chosen states to produce the set of states $\delta_H(q_{0h}, s\sigma)$.

$$\begin{aligned}
\delta_{H'}(q_{0h'}, s\sigma) &= \delta_{H'}(\bar{Q}', \sigma) && \text{by definition in our IH} \\
&= \bigcup_{\substack{q' \in \bar{Q}' \\ \delta_{G'}(q', \sigma)!}} \epsilon R_{G'}(\delta_{G'}(q', \sigma)) && \text{by definition of } H' \\
&\subseteq \bigcup_{\substack{q' \in \bar{Q}' \\ \delta_{G'}(q', \sigma)!}} \epsilon R_G(\delta_{G'}(q', \sigma)) && \text{by Lemma 1} \\
&\subseteq \bigcup_{\substack{q' \in \bar{Q}' \\ \delta_{G'}(q', \sigma)!}} \epsilon R_G(\delta_{G'}(q', \sigma)) \cup \bigcup_{\substack{q' \in \bar{Q}' \\ \neg \delta_{G'}(q', \sigma)!, \delta_G(q', \sigma)!}} \epsilon R_G(\delta_G(q', \sigma)) && \text{by set inclusion} \\
&\subseteq \bigcup_{\substack{q' \in \bar{Q}' \\ \delta_{G'}(q', \sigma)!}} \epsilon R_G(\delta_{G'}(q', \sigma)) \cup \bigcup_{\substack{q' \in \bar{Q}' \\ \neg \delta_{G'}(q', \sigma)!, \delta_G(q', \sigma)!}} \epsilon R_G(\delta_G(q', \sigma)) \cup \\
&\quad \bigcup_{\substack{q \in \bar{Q} \setminus \bar{Q}' \\ \delta_G(q, \sigma)!}} \epsilon R_G(\delta_G(q, \sigma)) && \text{by set inclusion} \\
&= \bigcup_{\substack{q' \in \bar{Q}' \\ \delta_{G'}(q', \sigma)!}} \epsilon R_G(\delta_{G'}(q', \sigma)) \cup \bigcup_{\substack{q \in \bar{Q} \setminus \bar{Q}' \\ \delta_G(q, \sigma)!}} \epsilon R_G(\delta_G(q, \sigma)) && \text{combining the last two terms} \\
&= \bigcup_{\substack{q \in \bar{Q} \\ \delta_G(q, \sigma)!}} \epsilon R_G(\delta_G(q, \sigma)) && \text{since by IH } \bar{Q}' \subseteq \bar{Q} \\
&= \delta_H(\bar{Q}, \sigma) \\
&= \delta_H(q_{0h}, s\sigma) && \text{by definition in our IH}
\end{aligned}$$

□

Having proven Theorem 1, we can relate our four automata from Figure 1: G' is a subautomaton of G , H' and H are transformations of G' and G respectively, and H' is a subobserver of H .

4 Updating the Observer Automaton

An adversary's view of a plant is tied to the plant's structure, so we explicitly link the adversary's plant estimate and the plant's actual state through the parallel composition of their respective automata (Definition 8 [4, p. 80]).

Definition 8 (Parallel Composition). Given two automata $G_1 = (Q_1, \Sigma_1, \delta_1, q_{01}, Q_{m1})$ and $G_2 = (Q_2, \Sigma_2, \delta_2, q_{02}, Q_{m2})$, their **parallel composition** is defined as:

$$G_1 \parallel G_2 = Ac(Q_1 \times Q_2, \Sigma_1 \cup \Sigma_2, \delta_{12}, (q_{01}, q_{02}), Q_{m1} \times Q_{m2})$$

$$\text{where } \delta_{12}((q_1, q_2), \sigma) = \begin{cases} (\delta_1(q_1, \sigma), \delta_2(q_2, \sigma)) & \text{if } \sigma \in \Sigma_1 \cap \Sigma_2 \\ (\delta_1(q_1, \sigma), q_2) & \text{if } \sigma \in \Sigma_1 \setminus \Sigma_2 \\ (q_1, \delta_2(q_2, \sigma)) & \text{if } \sigma \in \Sigma_2 \setminus \Sigma_1 \\ \text{undefined} & \text{otherwise} \end{cases}$$

The parallel composition $G \parallel H$, contains all of the information required to track the evolution of the plant's behaviour and the adversary's beliefs about the plant. We will denote states in $G \parallel H$ using the notation $\frac{q}{A}$ where q is a state in G and A is a state in H . It is useful to observe that a state $\frac{q}{A}$ is accessible in $G \parallel H$ if and only if there is a string $s \in L(G)$ that leads to state q in G and whose projection leads to estimate A in H (Lemma 2). As a corollary to this Lemma, if a state $\frac{q}{A}$ is accessible in $G \parallel H$ then the state q is an element of the estimate A .

Lemma 2. *Given the parallel composition $G \parallel H$ where $G = (Q, \Sigma, \delta_G, q_0, Q_m)$ is a DFA and $H := T_{\Sigma_o}(G)$ is an observer automaton. A state $\frac{q}{A}$ is accessible in $G \parallel H$ if and only if*

$$(\exists s \in L(G))(\delta_G(q_0, s) = q)((\forall s' \in L(G) \mid P_{\Sigma_o}(s) = P_{\Sigma_o}(s'))(\delta_G(q_0, s') = q' \in A)).$$

Proof. See the Appendix. □

The parallel composition has strong properties; first, it follows that $L(G \parallel H) = L(G)$ [11]. More interesting is that the accessible states of $T_{\Sigma_o}(G \parallel H)$ are pairwise disjoint sets of states of $G \parallel H$, i.e., every state in $G \parallel H$ belongs to one and only one state in $T_{\Sigma_o}(G \parallel H)$. An automaton that satisfies this property is a *state-partition automaton* (SPA). The SPA property allows for simple refinement operations, but we will not explicitly use it as Cho and Marcus used it, [5], since exploiting the SPA property would require a second transformation operation, $T_{\Sigma_o}(G \parallel H)$, and would introduce additional preprocessing overhead to our algorithms.

4.1 Using the subobserver property

We are given the parallel composition of a plant, G , and an adversary's view of that plant, H . Given that the plant's behaviour has been restricted, the **REFINE** algorithm (Algorithm 1) returns the parallel composition of the restricted plant and the adversary's updated view of the plant without explicitly calculating this adversary view. To begin, we remove states from the parallel composition that should be made inaccessible. Next, we account for the adversary's ability to reason about the restricted plant. If there is a string s in $L(G)$ that is no longer possible in G' , then the adversary cannot confuse other strings in $L(G')$ with s . We must therefore remove the state that s leads to in G from estimates where the adversary confused s with other strings.

To simplify our reasoning, we assume that the plant G and the parallel composition $G \parallel H$ have isomorphic state transition diagrams. Although this is not always the case, we will demonstrate in Section 5.3 that the parallel composition operation inherently creates an automaton \hat{G} that is language equivalent to G and whose state transition diagram is isomorphic with that of $G \parallel H$.

Data: Automaton $G \parallel H$, Δ , a list of states to make inaccessible in $G \parallel H$.
Result: An automaton M , the accessible portion of $G' \parallel H'$.

```

1  $M = G \parallel H$ ;
  /*  $A_r(A)$  is a set of states that should be removed from the adversary's
     estimate  $A$ . */
2 Initialize  $A_r(\cdot)$  as  $\emptyset$  for each estimate;
  /* Determine from  $\Delta$  if any other states should be made inaccessible */
3  $\Delta_Q = \Delta$ ;
4 while  $\Delta_Q$  is not empty do
5   | foreach transition leading into  $\frac{q}{A}$ ,  $t = \langle \frac{q'}{A'}, \sigma, \frac{q}{A} \rangle_M$  do
6   |   | if  $\sigma \in \Sigma_{uc}$  then the transition  $t$  is uncontrollable
7   |   |   | Add  $\frac{q'}{A'}$  to  $\Delta_Q$ ;
8   |   | Remove  $\frac{q}{A}$  from  $\Delta_Q$ ;
9   |   | Remove  $\frac{q}{A}$  from  $M$ ;
  /* States in  $M$  sharing the same estimate as newly inaccessible states need
     to be updated. */
10 foreach State  $\frac{q}{A}$  in  $G \parallel H$  do
11   | if  $\frac{q}{A}$  is not accessible in  $M$  then
12   |   |  $A_r(A) = A_r(A) \cup q$ ;
  /* Relabel states in  $M$  to account for the observer's updated estimates of
     the plant. */
13 foreach State  $\frac{q}{A}$  in  $M$  do
14   | Relabel  $\frac{q}{A}$  to  $\frac{q}{A \setminus A_r(A)}$ ;
  /* Ensure only the accessible portion of  $M$  is returned */
15  $M = Ac(M)$ ;
```

Algorithm 1: The REFINE algorithm

4.2 Technical developments

Before proving that the REFINE algorithm is correct, we need to develop a few ideas that will be necessary to connect our state-based formulation of opacity with the string-based confusion of the adversary. First, we know that REFINE is guaranteed to produce a subautomaton of G in the first elements of the states of M because it removes states from $G \parallel H$.

Lemma 3. *Given the parallel composition $G \parallel H$, where $G = (Q, \Sigma, \delta_G, q_0, Q_m)$ is a DFA and $H := T_{\Sigma_o}(G)$ is an observer automaton, and Δ is a list of states to remove from $G \parallel H$. Then the first elements of the states remaining in $G \parallel H$ represent a subautomaton of G .*

Proof. See the Appendix. □

Second, we note that a subautomaton G' can be defined by its parent automaton G and the set of strings that have been removed from $L(G)$. We define an equivalence relation ρ on Σ^* , which relates strings if they are both in $L(G')$, were both removed from $L(G)$, or were never in $L(G)$.

Definition 9. The binary relation $\rho_{G,G'}$ on Σ^* is defined as $(s, s') \in \rho_{G,G'}$ if and only if

$$(s, s' \in L(G')) \vee (s, s' \in L(G) \setminus L(G')) \vee (s, s' \in \Sigma^* \setminus L(G)).$$

Third, in **REFINE** we treat strings the same if they lead to the same state in G and the same estimate in H . We define an equivalence relation $match_{\Sigma_o}$ on Σ^* to capture this, where strings are related if they will receive identical treatment from **REFINE**.

Definition 10. Given a DFA G and observer automaton $H := T_{\Sigma_o}(G)$. The binary relation $match_{\Sigma_o}$ on Σ^* is defined as

$$(s, s') \in match_{\Sigma_o} \iff (s \equiv s' \pmod{eq(G)} \wedge (P_{\Sigma_o}(s) \equiv P_{\Sigma_o}(s') \pmod{eq(H)}))$$

and if $s \notin L(G)$ then $(s, s') \in \mu$ if and only if $s' \notin L(G)$ as well.

We can show that the partition defined by $match_{\Sigma_o}$ is equivalent to the partition defined by $eq(G \parallel H)$.

Lemma 4. Given a DFA $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, a set of observable events $\Sigma_o \subset \Sigma$, and an observer automaton $H := T_{\Sigma_o}(G)$. Then

$$match_{\Sigma_o} = eq(G \parallel H).$$

Proof. We will show that two strings $s, s' \in L(G)$ are related in $match_{\Sigma_o}$ if and only if they are related in $eq(G \parallel H)$.

$$\begin{aligned} (s, s') \in match_{\Sigma_o} & \\ \iff (s \equiv s' \pmod{eq(G)} \wedge (P_{\Sigma_o}(s) \equiv P_{\Sigma_o}(s') \pmod{eq(H)}) & \text{by Definition 10} \\ \iff \delta_G(q_0, s) = \delta_G(q_0, s') \wedge \delta_H(A_0, P_{\Sigma_o}(s)) = \delta_H(A_0, P_{\Sigma_o}(s')) & \text{by Definition 1} \\ \iff \delta_{G \parallel H}(\frac{q_0}{A_0}, s) = \delta_{G \parallel H}(\frac{q_0}{A_0}, s') & \text{by Definition 8} \\ \iff s \equiv s' \pmod{eq(G \parallel H)} & \text{by Definition 1} \end{aligned}$$

□

Finally, we can build on Lemma 4 by showing that if $match_{\Sigma_o}$ refines $\rho_{G, G'}$, then we can produce H' by relabelling states in H (i.e., the state transition graph of H' is guaranteed to be a subgraph of H 's state transition graph). This condition implicitly enforces that **REFINE** must remove states from $G \parallel H$ and thereby treat strings that lead to the same state and estimate identically. This is a reasonable restriction since we are considering state-based formulations of opacity and because alternate automaton representations of a language can be used if strings currently related by $match_{\Sigma_o}$ should be treated differently.

Lemma 5. Given DFAs $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, a set of observable events $\Sigma_o \subset \Sigma$, and an observer automaton $H = T_{\Sigma_o}(G)$. If G' is a subautomaton of G and $H' = T_{\Sigma_o}(G')$, then

$$(\forall s, s' \in L(H))(s \equiv s' \pmod{eq(H)}) \wedge match_{\Sigma_o} \leq \rho_{G, G'} \implies s \equiv s' \pmod{eq(H')}.$$

Proof. The proof proceeds by contradiction. Assume that $match_{\Sigma_o} \leq \rho_{G, G'}$ and that there exists strings s and s' in $L(H)$ such that $s \equiv s' \pmod{eq(H)}$ but $s \not\equiv s' \pmod{eq(H')}$.

$$s \equiv s' \pmod{eq(H)} \implies \delta_H(A_0, s) = \delta_H(A_0, s') \tag{1}$$

We can enumerate exactly those states that appear in $\delta_H(A_0, s)$, and therefore that appear in $\delta_H(A_0, s')$ as well.

$$\begin{aligned}\delta_H(A_0, s) &= \{q \in Q \mid (\exists t \in L(G))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q)\} \\ &= \{q \in Q \mid (\exists t \in L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q)\} \cup \\ &\quad \{q \in Q \mid (\exists t \in L(G) \setminus L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q)\} \\ &= \delta_{H'}(A_0, s) \cup \{q \in Q \mid (\exists t \in L(G) \setminus L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q)\}\end{aligned}$$

For notational convenience, we will denote the second term as

$$Q_\rho(s) := \{q \in Q \mid (\exists t \in L(G) \setminus L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q)\}. \quad (2)$$

These sets of states are disjoint, which can be shown by contradiction. Assume that there exists a state $q \in Q$ that belongs to the intersection of these two sets.

$$\begin{aligned}q \in \delta_{H'}(A_0, s) \cap Q_\rho(s) &\implies (\exists t \in L(G'), t' \in L(G) \setminus L(G')) \\ &\quad (\delta_G(q_0, t) = \delta_G(q_0, t') = q) \\ &\quad (P_{\Sigma_o}(t) = P_{\Sigma_o}(t') = s) \quad \text{definition of } \delta_{H'}(A_0, s) \text{ and } Q_\rho(s) \\ &\implies (\exists t \in L(G'), t' \in L(G) \setminus L(G')) \\ &\quad (\delta_G(q_0, t) = \delta_G(q_0, t') = q) \\ &\quad (\delta_H(A_0, P_{\Sigma_o}(t)) = \delta_H(A_0, P_{\Sigma_o}(t'))) \quad \text{since } P_{\Sigma_o}(t) = P_{\Sigma_o}(t') \\ &\implies (\exists t \in L(G'), t' \in L(G) \setminus L(G')) \\ &\quad (t \equiv t' \pmod{eq(G)}) \\ &\quad (P_{\Sigma_o}(t) \equiv P_{\Sigma_o}(t') \pmod{eq(H)}) \quad \text{Definition 1} \\ &\implies (\exists t \in L(G'), t' \in L(G) \setminus L(G')) \\ &\quad ((t, t') \in \text{match}_{\Sigma_o}) \quad \text{Definition 10} \\ &\implies (\exists t \in L(G'), t' \in L(G) \setminus L(G')) \\ &\quad (t, t' \in L(G') \vee t, t' \in L(G) \setminus L(G')) \quad \text{since } \text{match}_{\Sigma_o} \leq \rho_{G, G'}\end{aligned}$$

Since $L(G')$ and $L(G) \setminus L(G')$ are disjoint, this is a contradiction. We can therefore say that $\delta_{H'}(A_0, s) \cap Q_\rho(s) = \emptyset$ and that

$$\delta_{H'}(A_0, s) = \delta_H(A_0, s) \setminus Q_\rho(s). \quad (3)$$

For the sake of contradiction, we assumed that $s \not\equiv s' \pmod{eq(H')}$. By Definition 1 this implies that $\delta_{H'}(A_0, s) \neq \delta_{H'}(A_0, s')$. From (3) we know for s and s' that

$$\begin{aligned}\delta_{H'}(A_0, s) &= \delta_H(A_0, s) \setminus Q_\rho(s) \\ \delta_{H'}(A_0, s') &= \delta_H(A_0, s') \setminus Q_\rho(s')\end{aligned}$$

and since we have reasoned that $\delta_H(A_0, s) = \delta_H(A_0, s')$, (1), this implies that $Q_\rho(s) \neq Q_\rho(s')$.

$$\begin{aligned}
Q_\rho(s) \neq Q_\rho(s') &\implies (\exists q^* \in Q)(q^* \in Q_\rho(s))(q^* \notin Q_\rho(s')) && \text{WLOG, since } eq(H) \text{ is symmetric} \\
&\implies (\exists t \in L(G) \setminus L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q^*) \\
&\quad (\exists t' \in L(G) \setminus L(G'))(P_{\Sigma_o}(t') = s')(\delta_G(q_0, t') = q^*) && \text{by (2)} \\
&\implies (\exists t \in L(G) \setminus L(G'))(P_{\Sigma_o}(t) = s)(\delta_G(q_0, t) = q^*) \\
&\quad (\exists t' \in L(G'))(P_{\Sigma_o}(t') = s')(\delta_G(q_0, t') = q^*) && \text{since } L(G') \cap L(G) \setminus L(G') = \emptyset \\
&\implies (\exists t \in L(G) \setminus L(G'), t' \in L(G'))(P_{\Sigma_o}(t) = s)(P_{\Sigma_o}(t') = s') \\
&\quad (t \equiv t' \pmod{eq(G)}) && \text{by Definition 1} \\
&\implies (\exists t \in L(G) \setminus L(G'), t' \in L(G')) \\
&\quad (t \equiv t' \pmod{eq(G \parallel H)}) && \text{since } s \equiv s' \pmod{eq(H)} \\
&\implies (\exists t \in L(G) \setminus L(G'), t' \in L(G')) \\
&\quad ((t, t') \in match_{\Sigma_o}) && \text{by Lemma 4} \\
&\implies (\exists t \in L(G) \setminus L(G'), t' \in L(G')) \\
&\quad (t, t' \in L(G') \vee t, t' \in L(G) \setminus L(G')) && \text{since } match_{\Sigma_o} \leq \rho_{G, G'}
\end{aligned}$$

Again, since we know that $L(G')$ and $L(G) \setminus L(G')$ are disjoint, this is a contradiction. This implies that $Q_\rho(s) = Q_\rho(s')$, which further implies that $s \equiv s' \pmod{eq(H)}$. We have therefore proved that

$$(\forall s, s' \in L(H))(s \equiv s' \pmod{eq(H)}) \wedge match_{\Sigma_o} \leq \rho_{G, G'} \implies s \equiv s' \pmod{eq(H')}.$$

□

Lemma 5 guarantees us that for any state in H —cell in the partition defined by $eq(H)$ —we can produce the corresponding state in H' by relabelling that state. It goes beyond this too, by telling us that this relabelling removes exactly those states in Q_ρ from the state label in H .

4.3 Proof of correctness for the REFINE algorithm

As a first point, the REFINE algorithm is guaranteed to terminate because it works by removing and relabelling states in the input DFA, $G \parallel H$. For the first loop, the number of states and transitions to be inspected is finite and each state and transition can only be inspected once, therefore the loop terminates. For the second loop, there are a finite number of states to assess for accessibility and no changes are made to the structure of M . For the third loop, there are a finite number of states to relabel and relabelling a states in M never requires relabelling another state in M , therefore the loop terminates. Finally, taking the accessible portion of M requires assessing the accessibility of a finite number of states, so this operation also terminates.

Secondly, we claim that the REFINE algorithm uses the subobserver relationship to calculate $G' \parallel H'$ without computing H' explicitly. This is formalized as Theorem 2.

Theorem 2. *Given the parallel composition $G \parallel H$ and list of states Δ where: $G = (Q, \Sigma, \delta_G, q_0, Q_m)$ is a DFA; $H := T_{\Sigma_o}(G)$ is an observer automaton; and Δ is a list of states in $G \parallel H$. Denote by G' the automaton that results in the first elements of $G \parallel H$ once the states in Δ are made inaccessible. Then the REFINE algorithm produces $G' \parallel H'$ where $H' := T_{\Sigma_o}(G')$.*

Proof. We show that, given $G \parallel H$, where H is the transformed automaton $T_{\Sigma_o}(G)$, and Δ , a list of states to remove from $G \parallel H$ that produces G' in the first elements of the states in G , the **REFINE** algorithm constructs $G' \parallel H'$.

We assume that Δ is not equal to the set of states in $G \parallel H$, otherwise G' is the null automaton and we simply need to return $G' \parallel H' = \text{null}$. For the case where Δ is not equal to the set of states in $G \parallel H$, we reason about the elements of the states separately, since these correspond to the states of G' and H' .

First, we argue that the first loop removes the correct states from M to make the states in Δ inaccessible. Specifically, after the first loop:

$$\text{a state appears in } \Delta_Q \iff (\text{that state appears in } \Delta) \vee \\ (\text{that state leads uncontrollably to another state in } \Delta_Q.)$$

1. \implies . States are only added to Δ_Q at Line 3 if they are in Δ and Line 7 if they lead uncontrollably to a state in Δ_Q . So every state in Δ_Q meets one of those two conditions.
2. \impliedby . First, Line 3 makes Δ_Q equal to Δ before beginning the loop. Second, for any state $\frac{q'}{A'}$ that leads uncontrollably to a state $\frac{q}{A}$ in Δ_Q , there is an uncontrollable transition that leads from $\frac{q'}{A'}$ to $\frac{q}{A}$. When the iteration for the state $\frac{q}{A}$ is reached in the loop, that transition will be identified as uncontrollable and $\frac{q'}{A'}$ will be added to Δ_Q .

Since states are never removed from Δ_Q , this guarantees that all states in Δ are in Δ_Q and that all states that lead uncontrollably to a state in Δ_Q are in Δ_Q as well.

After the first loop is complete, Δ_Q contains exactly those states that must be removed to make the states in Δ inaccessible. These are exactly the states that are removed from M , since any state in Δ_Q is removed at Line 9 and states are not removed from M anywhere else in **REFINE**. Therefore the set of states in Δ have been made inaccessible and we denote in theorem statement the first element of the states in M as an automaton G'

Second, because we created G' by removing states from $M = G \parallel H$, we know by Lemma 3 that $G' \sqsubseteq G$. Since $G' \sqsubseteq G$, $H := T_{\Sigma_o}(G)$, and $H' := T_{\Sigma_o}(G')$, we know from Theorem 1 that $H' \tilde{\sqsubseteq} H$. We argue that because $H' \tilde{\sqsubseteq} H$, the estimates related to some strings need to be updated and that this can be accomplished by relabelling the states. We also argue that because the strings removed from $L(G)$ to produce $L(G')$ were determined by a set of states in $G \parallel H$, we have that $\text{match}_{\Sigma_o} \leq \rho_{G,G'}$ by Lemma 4 and therefore relabelling the estimates from H is all that is required to produce H' (Lemma 5).

The second loop in **REFINE** produces a list of states $A_r(A)$ for each estimate A . We claim that a state will be added to $A_r(A)$ if and only if it was in the estimate A to begin with and there are no longer any other strings in G' whose projections lead to the states in A :

$$q \in A_r(A) \iff (q \in A) \wedge ((\forall s \in L(G) \mid \delta_G(q_0, s) = q \wedge \delta_H(A_0, P_{\Sigma_o}(s)) = A)(q \notin \delta_{H'}(A_0, P_{\Sigma_o}(s))))).$$

We will reason about the two directions separately.

1. \implies . If $q \in A_r(A)$ then it was added at Line 12 because the state $\frac{q}{A}$ was in $G \parallel H$ but no longer accessible in M at Line 10. Because the state $\frac{q}{A}$ was accessible in $G \parallel H$, Lemma 2 tells us that $q \in A$, satisfying the consequent's first expression. For the consequent's second

expression, we will proceed by contradiction. Assume that there exists a string $s \in L(G)$ such that

$$\delta_G(q_0, s) = q \wedge \delta_H(A_0, P_{\Sigma_o}(s)) = A \wedge q \in \delta_{H'}(A_0, P_{\Sigma_o}(s)).$$

Because $q \in \delta_{H'}(A_0, P_{\Sigma_o}(s))$, we know that there exists a string $s' \in L(G')$ such that $P_{\Sigma_o}(s) = P_{\Sigma_o}(s')$ and $\delta_{G'}(q_0, s') = q$. Since $P_{\Sigma_o}(s) = P_{\Sigma_o}(s')$, we have that $q \in \delta_{H'}(A_0, P_{\Sigma_o}(s'))$ and therefore the state $\frac{q}{A}$ is accessible in $G' \parallel H'$. But if $\frac{q}{A}$ is accessible in $G' \parallel H'$, then it would have been accessible in M at Line 11 since the first elements of the states in M reflect G' . The state q therefore not have been added to $A_r(A)$ at Line 12, which is a contradiction.

2. \Leftarrow . Because $(q \in A) \wedge (\forall s \in L(G) \mid \delta_G(q_0, s) = q \wedge \delta_H(A_0, P_{\Sigma_o}(s)) = A)$ we have that the state $\frac{q}{A}$ was accessible in $G \parallel H$. The last expression in the antecedent tells us, however, that for all these strings s that led to $\frac{q}{A}$ in $G \parallel H$, $(q \notin \delta_{H'}(A_0, P_{\Sigma_o}(s)))$. This tells us that for such strings s , there does not exist a string $s' \in L(G)$ such that $\delta_{G'}(q_0, s') = q$. This means that the state $\frac{q}{A}$ must now be inaccessible in M after the first loop. This means that during the iteration for $\frac{q}{A}$ in the second loop, it is inaccessible and therefore q is added to $A_r(A)$ at Line 12. Since states are never removed from $A_r(A)$, this guarantees that $q \in A_r(A)$.

Third, the third loop correctly updates all of the estimates in M because all states in M are visited and because all states in M sharing an estimate must be updated in the same way (Lemma 5).

To conclude, after the **REFINE** algorithm has completed its work, the resulting automaton M is equal to $G' \parallel H'$ such that

$$\delta_M\left(\frac{q_0}{A_0}, s\right) = \frac{q}{A} \iff \delta_{G'}(q_0, s) = q \wedge \delta_{H'}(A_0, P_{\Sigma_o}(s)) = A.$$

□

4.4 Time complexity analysis of the **REFINE** algorithm

We tie the asymptotic time complexity of the **REFINE** algorithm to the size of the automaton $G \parallel H$ with N states, an alphabet of size $|\Sigma|$, and, because it is a DFA, an upper bound of $N \times |\Sigma|$ transitions.

Line 1 is a memory operation where we denote the input $G \parallel H$ by M . This happens in constant time. Line 2 is also a memory operation, where the set of events to remove from at most N estimates are initialized to \emptyset . The complexity class for this operation is $\mathcal{O}(N)$. Line 3 is another memory operation where we denote the input Δ by Δ_Q . This happens in constant time.

The first loop, starting at Line 4, will in the worst case examine every transition in $G \parallel H$ once. Practically speaking, transitions could be stored in a sorted list with all uncontrolled transitions at the start of the list. For each transition, at most four constant time operations are performed: verifying if the transition is uncontrolled; adding the origin state to Δ_Q ; and removing the destination state from both Δ_Q and M . The complexity class for this loop is $\mathcal{O}(N \times |\Sigma|)$.

The second loop, starting at Line 10, will examine every state in $G \parallel H$ once. For each state, at most two constant time operations are performed: checking if the state is inaccessible; and adding the first element of the state to $A_r(A)$. The complexity class for this loop is $\mathcal{O}(N)$.

The third loop, starting at Line 13 will examine every state in $G \parallel H$ once. For each state, exactly one constant time operation is performed: relabelling the state's estimate by removing the previously stored set $A_r(A)$ from the current estimate. The complexity class for this loop is $\mathcal{O}(N)$.

Line 15 requires returning the accessible portion of M . Practically, the inaccessible components of M can be removed during the second loop, and so this is a constant time operation.

Taken all together, the complexity class for **REFINE** is $\mathcal{O}((N \times 3) + (N \times \Sigma))$, which simplifies to $\mathcal{O}(N \times |\Sigma|)$. If we consider that the number of states in an automaton is generally larger than the number of events, this further simplifies to $\mathcal{O}(N)$.

4.5 Discussion

The subobserver relationship is a more general form of the subautomaton relationship. We showed that the subobserver relationship captures how the observer automaton changes to reflect evolutions in the plant. Specifically, the subobserver relationship allows the parallel composition of a plant and an observer automaton for that plant, $G \parallel H$, to be refined without recomputing the observer. Instead, our method removes transitions in the parallel composition if and only if they must be disabled by supervisory control and updates the observer's estimates. The **REFINE** algorithm implements this procedure and its asymptotic time complexity is linear in the number of states of $G \parallel H$.

Although the **REFINE** algorithm was developed within the framework of the opacity control problem, it more generally produces the joint behaviour of plant and observer as the plant evolves over time. In the remainder of this paper we consider the opacity control problem exclusively, but we note that this approach is applicable to any DES problem where evolving discrete-event processes are observed including online control, dynamic discrete-event systems, and decentralized control.

5 Synthesizing Opacity-Enforcing Supervisors

Our opacity control problem formulation is similar to that from other works (see Problem 1). We assume that all controllable events are observable, $\Sigma_c \subseteq \Sigma_s$, and that the adversary observes a subset of the events that the supervisor observes, $\Sigma_a \subseteq \Sigma_s$. These assumptions allow us to use Dubreil et al.'s reduction of the general opacity control problem to the opacity control problem with full observation, i.e. $\Sigma_s = \Sigma$ [7, Proposition 5]. We assume a supervisor-aware adversary, which implies that the adversary's view will evolve with $G' = \mathcal{S}/G$. For simplicity we enforce current-state opacity.




Figure 3: Comparing the calculations and data structures for each approach.

Recalling approaches to the opacity control problem, the naïve, computationally inefficient solution is to apply supervisory control and calculate observers iteratively (Figure 3a) [7]. Condensed state estimates (Figure 3b) require alternately constructing partial maps and automata; since the only states in a condensed state estimate are those that were reached by the last event observable to the adversary, “loosing traces” are calculated at each iteration to determine which states will reveal a secret [7]. The augmented I-observer (Figure 3c) solves a more general form of the problem but requires calculating the supremal G-opaque sublanguage, synthesizing a supervisor and refining both the plant and augmented I-observer at each step. Tong et al. avoid this complication by restricting themselves to the case of an adversary who is unaware of the supervisor [23].

Our method uses the subobserver property to calculate $G' \parallel H'$ directly from $G \parallel H$, refining the parallel composition of plant and adversary view (Figure 3d). Advantages of this method include: not producing any intermediary structures or languages; capturing the adversary’s beliefs better than condensed state estimates, thus not requiring “loosing traces” to be calculated; and avoiding unnecessary calculations compared to the augmented I-observer with a supervisor-aware adversary.

5.1 The SYNTHESIZE algorithm

As in Section 4, we use the parallel composition of G and $H := T_{\Sigma_a}(G)$ to track the plant’s evolving structure along with the adversary’s beliefs about the plant.

```

Data: Plant  $G$ ; secret states  $Q_m$  and non-secret states  $Q \setminus Q_m$ ; alphabet visible to the
          adversary  $\Sigma_a \subseteq \Sigma$  and alphabet controllable by supervisor  $\Sigma_c \subseteq \Sigma$ .
Result: Supervisor  $\mathcal{S}$  such that  $\mathcal{S}/G$  is current-state opaque with respect to  $Q_m$  and  $\Sigma_a$ .
/* Create the observer automaton and the parallel composition. */
1  $H = T_{\Sigma_a}(G)$ ;
2  $M = G \parallel H$ ;
3 while  $M$  has marked states do
    /*  $\Delta$  is the list of states to make inaccessible in  $M$ , i.e. the states
       that cause opacity to not be enforced. */
4   Set  $\Delta$  equal to the set of marked states in  $M$ ;
   /* Refining  $M$  reflects that the adversary changes its view of the plant
      based on its ability to reason about the policy enforced by the
      supervisor. */
5    $M = \text{REFINE}(M, \Delta)$ ;
   /* Mark any states whose new estimates contain only secret states. */
6   foreach  $State \frac{q}{A}$  in  $M$  do
7     | if  $A \subseteq Q_m$  then
8     | | Mark the state  $\frac{q}{A}$  in  $M$ ;
/* If  $M$  is an empty automaton, then there is no valid supervisor for  $G$ 
   that is able to enforce current-state opacity with respect to the
   adversary’s observations  $\Sigma_a$ . */
9  $\mathcal{S} = M$ ;

```

Algorithm 2: The SYNTHESIZE algorithm

Importantly, we mark states in H to capture the scenarios in which no non-secret state can be confused by the adversary with a true secret state. With a traditional observer automaton, a

state is marked in H when one of the states that it contains is itself marked in G ; in our work, by contrast, a state in H is marked when all the states it contains are themselves marked in G , as described in Section 2.1.1. Thus, in our work, if H – the observer – reaches a marked state, then a string has occurred in G – the plant – that allows the adversary to be sure that a secret state has been reached. This marking carries through to $G \parallel H$, where marked states must be made inaccessible to produce an opacity-enforcing supervisor.

The SYNTHESIZE algorithm (Algorithm 2) enforces current-state opacity by applying supervisory control as long as the controlled plant composed with the adversary’s estimations of the plant contains marked states (Lines 3-8), the adversary’s view of the plant is updated (Line 5), and the new composition of plant behaviour and adversarial belief is checked for current-state opacity (Lines 6-8).

5.2 Worked example

We will demonstrate how the SYNTHESIZE algorithm works for a concrete example, beginning with the plant and the initial parallel composition (Figure 4). This plant has two marked states, 6 and 10, an event that the supervisor cannot control, β , and events that the adversary cannot observe, α and γ .




Figure 4: The initial plant and parallel composition for the worked example.

The initial parallel composition has a single marked state, $\frac{10}{10}$, which is added to Δ and passed to REFINE. In order to prevent the state $\frac{10}{10}$ from being reached, we must also mark the state $\frac{9}{6,9}$ in M since it leads uncontrollably to our marked state. Removing these two states from M implies that when the supervisor is in state $\frac{8}{3,4,8}$ it will disable the event λ via supervisory control since the transition $\langle \frac{8}{3,4,8}, \lambda, \frac{9}{6,9} \rangle_M$ is no longer defined. Removing these states from M , however, means that our supervisor will not allow the plant G to reach state 9, which will affect the adversary’s estimates. Specifically here, when the plant is in state 6, the adversary will no longer confuse the string in G that led to state 6 with a string in G that leads to state 9 since the latter is now inaccessible in the plant. This leads to the new parallel composition in Figure 5a, which also has one marked state, $\frac{6}{6}$.

Removing the state $\frac{6}{6}$ from M produces the parallel composition in Figure 5b. This automaton has no marked states, indicating that it can be used as a supervisor to enforce current-state opacity in the plant G .




Figure 5: The parallel compositions computed during the worked example.

5.3 The case of state splitting

As remarked in Section 3, we are not guaranteed that applying supervisory control to G will result in a subautomaton of G . For example, the observer automaton in Figure 6 illustrates how although the secret states 8 and 9 are repeatedly visited, the adversary is only able to determine that the plant is in a secret state once it sees a fourth a . This means that a supervisor can allow state 8 to be visited up to three times and must therefore enact different control patterns for state 8 depending on how many times the plant has visited state 8.

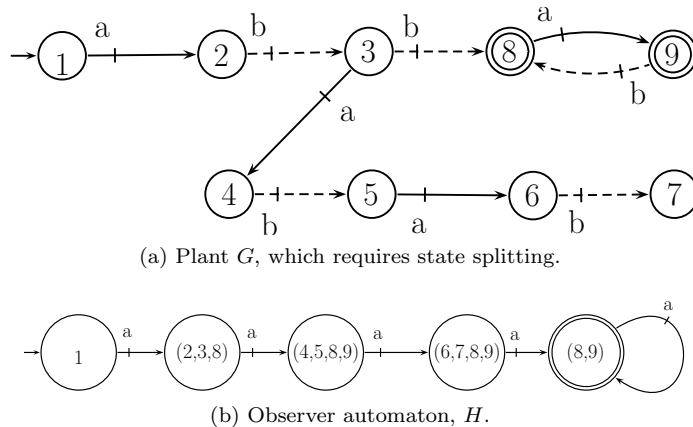


Figure 6: States 8 and 9 in G are secret and only a is observable to the adversary. The observer automaton shows that after observing a fourth a , an adversary will know the plant is in a secret state. Note that only $\{8, 9\}$ is marked in H because this is the only state whose constituent states are all marked in G .

State splitting is required for opacity enforcement when a state can be visited in the prefixes of strings leading to secret states but only a finite number of times, requiring the supervisor to know which string led to the state for a particular visit [4, p. 141]. Note that if a state can be visited an infinite number of times then the supervisor can always enact the same control pattern. Happily, this state splitting occurs implicitly in Line 2 of the SYNTHESIZE algorithm when the first parallel composition is constructed (Figure 7). Because a state that requires splitting can only be safely visited a finite number of times, it is possible to transform any DFA, G , into a language-equivalent

DFA, \hat{G} , by replicating the original state for each time that it can be safely visited.




Figure 7: Continuing from Figure 6, the first parallel composition inherently constructs the needed, language-equivalent automaton to represent the plant.

In fact, taking only the first element of the states in $G \parallel H$'s state transition diagram produces \hat{G} 's state transition diagram. Following this initial transformation, we are then guaranteed that G' will be a strict subautomaton of \hat{G} and the SYNTHESIZE algorithm can proceed from there. We note that this state splitting is not an artifact of our proposed method: any method for producing a minimally restrictive opacity-enforcing supervisor from the original plant G must be able to reason about the number of times that such states have been reached. The structure $G \parallel H$ embeds this information by splitting states the necessary number of times; it is therefore the smallest memoryless structure that can be used to solve the opacity control problem in the context of a supervisor-aware adversary.

5.4 Proof of correctness for the SYNTHESIZE algorithm

First, the SYNTHESIZE algorithm is guaranteed to terminate. As with the REFINE algorithm, termination is guaranteed because the automaton $G \parallel H$ has a finite number of states and the SYNTHESIZE algorithm makes states inaccessible until M enforces opacity or is the null automaton. Specifically, if M has any marked states then the algorithm removes states – never adding states – to make the marked states inaccessible. Since M has a finite number of states, and the algorithm terminates if M is the null automaton, the algorithm is guaranteed to terminate.

Second, we prove the correctness of SYNTHESIZE through Theorem 3.

Theorem 3. *Given a DFA $G = (Q, \Sigma, \delta_G, q_0, Q_m)$, set of secret states Q_m , set of events visible to the adversary $\Sigma_a \subseteq \Sigma$, and set of events controllable by the supervisor $\Sigma_c \subseteq \Sigma$. Then the supervisor \mathcal{S} produced by the SYNTHESIZE algorithm is a correct and maximally-permissive opacity-enforcing supervisor for G .*

Proof. We begin by proving by contradiction that the supervisor \mathcal{S} correctly enforces opacity for G . We assume that \mathcal{S} does not enforce opacity for G . In this case, at least one of the states in $\mathcal{S} = M$ must be marked, reflecting that all of the states in the adversary's estimate of the plant are secret states in G . If this was the case, however, then the while loop at Lines 3 to 8 would have executed again and removed the last controllable transition along each path to this state, rendering it inaccessible. Therefore, when the SYNTHESIZE algorithm terminates, $\mathcal{S} = M$ has no marked states and it therefore correctly enforces opacity for G .

Next, we prove by contradiction that the supervisor \mathcal{S} is maximally-permissive while enforcing opacity for G . Assume that there exists a supervisor \mathcal{S}' that correctly enforces opacity for G while

being more permissive than supervisor \mathcal{S} . This implies that there is a transition in $G \parallel H, \delta^*$, that \mathcal{S}' enables and that \mathcal{S} disables. But the only transitions that \mathcal{S} disables in $G \parallel H$ are those that lead directly to a marked state or to a state which leads uncontrollably to a marked state, so δ^* must either lead directly or uncontrollably to a marked state. Therefore, \mathcal{S}' does not correctly enforce opacity for G . \square

5.5 Time complexity analysis of the SYNTHESIZE algorithm

Although the SYNTHESIZE algorithm presents a straightforward way of seeing how we can iteratively generate a list of transitions, Δ , and then refine the parallel composition of plant and adversary's view, its asymptotic time complexity is in a higher complexity class than other methods in the literature. We therefore present Algorithm 3, which embeds the REFINE algorithm into the SYNTHESIZE algorithm in a computationally efficient manner. We tie the asymptotic behaviour of Algorithm 3 to the input of interest, the automaton G with $|Q|$ states and an alphabet of size $|\Sigma|$.

Data: Plant G ; secret states Q_m and non-secret states $Q \setminus Q_m$; alphabet visible to the adversary $\Sigma_a \subseteq \Sigma$ and alphabet controllable by supervisor $\Sigma_c \subseteq \Sigma$.
Result: Supervisor \mathcal{S} such that \mathcal{S}/G is current-state opaque with respect to Q_m and Σ_a .
/ Initialize the algorithm. */*

```

1  $H = T_{\Sigma_a}(G)$  ;
2  $M = G \parallel H$ ;
3 Produce  $\mathcal{H}$ , a hash table with adversary estimates as keys and states of  $M$  as values ;
4 foreach marked state in  $M$ ,  $m = \frac{q}{A}$  do
5   | foreach transition leading into  $m$ ,  $t = \langle m', \sigma, m \rangle_M$  do
6   |   | if  $\sigma \in \Sigma_{uc}$  then the transition  $t$  is uncontrollable
7   |   |   | Mark  $m'$  in  $M$ ;
8   |   | Remove the state  $m$  from  $M$ ;
9   |   | Go to the key  $A$  in  $\mathcal{H}$  and replace it with  $A \setminus \{q\}$ ;
10  |   | if  $A \setminus \{q\} \subseteq Q_m$  then
11  |   |   | Mark all states in the hash table slot indexed by  $\mathcal{H}(A \setminus \{q\})$  ;
12  $\mathcal{S} = M$ ;

```

Algorithm 3: An computationally efficient implementation of the SYNTHESIZE algorithm

Initializing the algorithm, Lines 1-3, requires constructing the adversary's view, $H = T_{\Sigma_a}(G)$, and the parallel composition of $G \parallel H$. The time complexity of this process is established to be $\mathcal{O}(|Q| \times 2^{|Q|})$. The production of the hash table \mathcal{H} requires an operation for each $|Q|$ state in M as well and in practice it can be constructed alongside the parallel composition.

The outer loop iterates through the marked states in M , Lines 4 to 11. This loop will run a variable number of times, but we will be able to assess the loop's complexity by establishing upper bounds for each of the lines inside this loop. First, the loop responsible for inspecting transitions in M , Lines 5-7, visits each transition in M at most once and at most one constant time operation is performed. This is guaranteed because each visited transition is visited because its destination state is marked. Because that destination state will be removed at Line 8 of the current iteration, the transition in question will not be visited again. Lines 5-7 therefore have an asymptotic time complexity over the algorithm's whole run time of $\mathcal{O}(|Q| \times 2^{|Q|} \times |\Sigma|)$. Second, the lines responsible for processing marked states in M , Lines 8-11 consists of up to four operations. Removing the state

m from M has a constant time complexity; because each state in M can be removed at most once this line has an asymptotic time complexity over the algorithm’s whole run time of $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|})$. Updating the hash table key similarly has a constant time complexity and can only occur once per state in M . This again leads to an asymptotic time complexity over the algorithm’s whole run time of $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|})$. Finally, checking if the new estimate consists of solely secret states (Line 10) takes at most $|Q_m|$ checks and then marking states in the hash table slot (Line 11) takes constant time by state. Since every state can only be marked once over the algorithm’s whole run time, this again leads to an asymptotic time complexity of $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|})$.

Taken together, the asymptotic time complexity of Algorithm 3 is $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|} \times |\Sigma|)$. Since in many problems the size of the alphabet does not grow and is much smaller than the size of the state space, we can consider $|\Sigma|$ as a constant and reduce this time complexity to $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|})$. Our result is in keeping with the exponential time required for computing an observer automaton and we note that there is no algorithm for verifying current-state opacity for a system G whose time complexity is polynomial in the number of states in G [25].

5.6 Discussion

We have demonstrated how our method for refining parallel compositions of plant and observer can be used to synthesize an opacity-enforcing supervisor. The SYNTHESIZE algorithm implements this procedure in a direct manner and Algorithm 3 trades a small increase in memory usage to reduce the asymptotic time complexity to $\mathcal{O}(|Q| \times 2^{|\mathcal{Q}|})$ where $|Q|$ is the number of states in the original plant.

We showed two examples to demonstrate how our algorithm works. The first example highlighted the interplay between enacting supervisory control and enforcing opacity when faced with a supervisor-aware adversary: multiple refinements of the automaton $G \parallel H$ may be required before an output is produced that can be used as an opacity-enforcing supervisor. The second example showed how our approach inherently addresses splits states when a secret state can be safely visited only a finite number of times without revealing the system secret. Recalling the two methods in the literature that are most similar to ours, we note that Algorithm 3 offers improvements on both approaches. Although each method has the same asymptotic time complexity, [23], our method produces only a single structure at each step: the parallel composition of plant and adversary estimate.

The method of condensed state estimates, by contrast, produces both a partial map and an automaton at every step [7]. Additionally this method requires “loosing paths” to be assessed in order to verify the opacity property because condensed state estimates account only for states that the plant might have reached if the last observed event is the last event that occurred in the plant. Our method does away with the need to reason about these “loosing paths,” which produces two advantages. First, opacity is immediately verified by the lack of marked states in the output which permits an opacity-enforcing supervisor to be easily verified by an independent party. Second, the output’s state labels are semantically meaningful: if a trace occurs in the system then this trace leads to a state in the parallel composition whose label is the current state of the system and the states that the observer believes the system could be in. This information can be used by an engineer or a computer program to make decisions about the system.

The augmented I-observer method, solves a more general problem by allowing Σ_s and Σ_a to be incomparable [23]. Balanced against this, it focuses on the case of a single iteration to enforce opacity against a non-supervisor-aware adversary. Because this method calculates the supremal

G-opaque sublanguage of the plant’s language and a supervisor to enforce this language as a specification, [23], this results in unnecessary calculations when a a supervisor-aware adversary makes multiple iterations necessary. Our method makes the realistic assumption that $\Sigma_a \subseteq \Sigma_s$ and addresses the case of a supervisor-aware adversary with only a single structure computed at each step.

6 Conclusion

The opacity control problem has been of interest in the DES literature for many years and has been addressed a number of times. It is a complex problem whose straightforward solution is computationally expensive. Compared with our method, the solutions to the opacity control problem in the literature variously require additional computations to deal with the problem in its general form or use structures that do not intuitively align with the opacity control problem.

This paper introduced and described the subobserver relationship. This relationship is analogous to the subautomaton relationship and links the observers of a plant whose structure evolves. We demonstrated the usefulness of the subobserver relationship by using it as the basis of an algorithm to solve the opacity control problem and we believe that this relationship has broader application in DES research including online control, dynamic discrete-event systems, and decentralized control.

Time complexity analysis shows that our algorithms are computationally efficient, with Algorithm 3 matching the asymptotic time complexity of previous methods in the literature. We also demonstrated that our algorithms intuitively solve the opacity control problem under reasonable assumptions. Future research applying these ideas may include relaxing the requirement that $\Sigma_c \subseteq \Sigma_s$ for the opacity control problem and applying the subobserver relationship to the iterative refinements necessary when a single supervisor enacts online control or several agents enact decentralized control over a plant.

Acknowledgements

The authors acknowledge that Queen’s University is situated on traditional Anishinaabe and Haudenosaunee Territory.

The authors would like to thank Richard Ean and Bryony Schonewille of the Queen’s Discrete-Event Systems Lab as well as the reviewers for their thoughtful and constructive comments. This research was supported by the Natural Sciences and Engineering Research Council of Canada as well as the Queen’s University Faculty of Engineering and Applied Science. RHM held a Walter C. Sumner Memorial Fellowship during the conduct of this work.

All figures with automata were produced using the Integrated Discrete-Event Systems Tool, [18], which is freely available online at <https://github.com/krudie/IDES>.

A Omitted Proofs

Proof of Lemma 2. We begin by noting that $L(G \parallel H) = L(G)$ and that because $\Sigma_o \subseteq \Sigma$, the alphabet for $G \parallel H$ is Σ as well.

\implies . If a state $\frac{q}{A}$ is accessible in $G \parallel H$, then there exists a string $s \in L(G \parallel H)$ such that $\delta_{G \parallel H}(\frac{q_0}{A_0}, s) = \frac{q}{A}$. By Definition 8, this implies that that $\delta_G(q_0, s) = q$ and that $\delta_H(A_0, P_{\Sigma_o}(s)) =$

A. Because $H := T_{\Sigma_o}(G)$, this implies that

$$((\forall s' \in L(G) \mid P_{\Sigma_o}(s) = P_{\Sigma_o}(s'))(\delta_G(q_0, s') = q' \in A)).$$

\Leftarrow . We are given that

$$(\exists s \in L(G))(\delta_G(q_0, s) = q)((\forall s' \in L(G) \mid P_{\Sigma_o}(s) = P_{\Sigma_o}(s'))(\delta_G(q_0, s') = q' \in A)).$$

We therefore know that $\delta_G(q_0, s) = q$ and that $\delta_H(A_0, P_{\Sigma_o}(s)) = A$. By Definition 8, this implies that $\delta_{G \parallel H}(\frac{q_0}{A_0}, s) = \frac{q}{A}$ and therefore $\frac{q}{A}$ is accessible in $G \parallel H$. \square

Proof of Lemma 3. According to the definition, G' is a subautomaton of G , denoted by $G' \sqsubseteq G$, if $\delta_{G'}(q'_0, s) = \delta_G(q_0, s) \quad \forall s \in L(G')$.

Denote the first elements of states remaining in $G \parallel H$ as G' . Then we know that

$$(\forall s \in L(G'))(\delta_{G'}(q_0, s) = q \implies \delta_G(q_0, s) = q)$$

since any first element of a state that survives in $G \parallel H$ was originally in G . This means that any string that occurs in G' leads to the same state that it did in G and therefore that $\delta_{G'}(q'_0, s) = \delta_G(q_0, s) \quad \forall s \in L(G')$.

We conclude that G' is a subautomaton of G . \square

References

- [1] BADOUEL, E., BEDNARCZYK, M., BORZYSZKOWSKI, A., CAILLAUD, B., AND DARONDEAU, P. Concurrent Secrets. *Discrete Event Dynamic Systems* 17, 4 (Nov 2007), 425–446.
- [2] BEN-KALEFA, M., AND LIN, F. Supervisory control for opacity of discrete event systems. In *49th Annual Allerton Conference on Communication, Control, and Computing* (Sep 2011), IEEE, pp. 1113–1119.
- [3] BRYANS, J. W., KOUTNY, M., MAZARÉ, L., AND RYAN, P. Y. Opacity generalised to transition systems. *International Journal of Information Security* 7, 6 (2008), 421–435.
- [4] CASSANDRAS, C. G., AND LAFORTUNE, S. *Introduction to Discrete Event Systems*, second ed. Springer US, Boston, MA, 2008.
- [5] CHO, H., AND MARCUS, S. I. Supremal and maximal sublanguages arising in supervisor synthesis problems with partial observations. *Mathematical Systems Theory* 22, 1 (1989), 177–211.
- [6] CHUNG, S.-L., LAFORTUNE, S., AND LIN, F. Limited lookahead policies in supervisory control of discrete event systems. *IEEE Transactions on Automatic Control* 37, 12 (1992), 1921–1935.
- [7] DUBREIL, J., DARONDEAU, P., AND MARCHAND, H. Supervisory control for opacity. *IEEE Transactions on Automatic Control* 55, 5 (2010), 1089–1100.
- [8] GRIGOROV, L., AND RUDIE, K. Near-Optimal Online Control of Dynamic Discrete-Event Systems. *Discrete Event Dynamic Systems* 16, 4 (Dec 2006), 419–449.
- [9] HOPCROFT, J. E., AND ULLMAN, J. D. *Introduction to Automata Theory, Languages and Computation*.
- [10] JACOB, R., LESAGE, J. J., AND FAURE, J. M. Overview of discrete event systems opacity: Models, validation, and quantification. *Annual Reviews in Control* 41 (2016), 135–146.
- [11] JIRÁSKOVÁ, G., AND MASOPUST, T. On properties and state complexity of deterministic state-partition automata. *Lecture Notes in Computer Science* 7604 (2012), 164–178.

- [12] LIN, F. Opacity of discrete event systems and its applications. *Automatica* 47, 3 (Mar 2011), 496–503.
- [13] LIN, F., AND WONHAM, W. M. Decentralized supervisory control of discrete-event systems. *Information Sciences* 44, 3 (1988), 199–224.
- [14] LIN, F., AND WONHAM, W. M. On observability of discrete-event systems. *Information Sciences* 44, 3 (1988), 173–198.
- [15] MAZARÉ, L. Using Unification for Opacity Properties. In *Proceedings of the Workshop on Issues in the Theory of Security (WITS'04)* (2004), vol. 7, pp. 165–176.
- [16] RAMADGE, P. J., AND WONHAM, W. M. Supervisory Control of a Class of Discrete Event Processes. *SIAM Journal on Control and Optimization* 25, 1 (Jan 1987), 206–230.
- [17] ROSEN, K. H. *Discrete Mathematics and its Applications*, eighth (in ed. McGraw-Hill Education, New York, USA, 2019).
- [18] RUDIE, K. The Integrated Discrete-Event Systems Tool. *8th International Workshop on Discrete Event Systems* (2006), 394–395.
- [19] SABOORI, A., AND HADJICOSTIS, C. N. Verification of initial-state opacity in security applications of DES. *9th International Workshop on Discrete Event Systems (WODES 08)* (2008), 328–333.
- [20] SABOORI, A., AND HADJICOSTIS, C. N. Verification of K-Step Opacity and Analysis of Its Complexity. *IEEE Transactions on Automation Science and Engineering* 8, 3 (Jul 2011), 549–559.
- [21] SABOORI, A., AND HADJICOSTIS, C. N. Opacity-enforcing supervisory strategies via state estimator constructions. *IEEE Transactions on Automatic Control* 57, 5 (2012), 1155–1165.
- [22] TAKAI, S., AND OKA, Y. A Formula for the Supremal Controllable and Opaque Sublanguage Arising in Supervisory Control. *SICE Journal of Control, Measurement, and System Integration* 1, 4 (2008), 307–311.
- [23] TONG, Y., LI, Z., SEATZU, C., AND GIUA, A. Current-state opacity enforcement in discrete event systems under incomparable observations. *Discrete Event Dynamic Systems: Theory and Applications* 28, 2 (2018), 161–182.
- [24] WONHAM, W. M., AND CAI, K. *Supervisory Control of Discrete-Event Systems*. Communications and Control Engineering. Springer International Publishing, Cham, 2019.
- [25] WU, Y. C., AND LAFORTUNE, S. Comparative analysis of related notions of opacity in centralized and coordinated architectures. *Discrete Event Dynamic Systems: Theory and Applications* 23, 3 (2013), 307–339.
- [26] YIN, X., AND LAFORTUNE, S. A new approach for synthesizing opacity-enforcing supervisors for partially-observed discrete-event systems. In *2015 American Control Conference (ACC)* (Chicago, USA, Jul 2015), IEEE, pp. 377–383.
- [27] YIN, X., AND LAFORTUNE, S. A Uniform Approach for Synthesizing Property-Enforcing Supervisors for Partially-Observed Discrete-Event Systems. *IEEE Transactions on Automatic Control* 61, 8 (2016), 2140–2154.