

Hierarchical Supervisory Control of Networked and Cyber-Attacked Discrete-Event Systems

Shaowen Miao ^{a,b}, Jan Komenda ^{c,*}, Feng Lin ^d

^a*Robotics and Autonomous Systems Thrust, Hong Kong University of Science and Technology (Guangzhou), 511453, China*

^b*Department of Automation, Xiamen University, Xiamen 361102, China*

^c*Institute of Mathematics, Czech Academy of Sciences, 115 67 Prague, Czechia*

^d*Department of Electrical and Computer Engineering, Wayne State University, Detroit, MI 48202, USA*

Abstract

In standard supervisory control of discrete-event systems, partial (incomplete) observations are given by deterministic functions such as natural projections, which erase unobservable events, or masks, which can represent indistinguishable events, where two or more different events yield the same observation. However, communication channels in modern technological systems are not always reliable and can be attacked by malicious external agents. In that case, the plant observations obtained by the supervisor may not be deterministic, e.g., due to delays and losses, or external attacks. This paper considers a unified supervisory control framework with set-valued (nondeterministic) observations and proposes a simplified version of nondeterministic observability, together with a generalized normality. It shows how the results of hierarchical control can be extended to the networked and cyber-attacked discrete-event systems at the same time.

Key words: Discrete-event systems, hierarchical supervisory control, networked systems, cyber-attacks.

1 Introduction

Recent developments in various networks and connectivity features have led to nonreliable communication channels, which suffer from internal issues such as delays and losses of information [20], as well as increased vulnerability to attacks by external malicious agents [4]. Supervisory control theory [16] has been proposed as a formal approach for controlling automata to address safety and nonblockingness issues in technological systems. The emergence of networked embedded systems and sensor/actuator networks has stimulated further development of supervisory control of discrete-event systems (DESs), particularly in hierarchical [6,10,15,19] and decentralized control [9,12,17] of large-scale DESs.

In [15], observer property was introduced to guarantee

the abstracted plant has a significantly reduced state space compared to the original plant, i.e., $|Q_{hi}| \ll |Q|$, where $|Q_{hi}|$ and $|Q|$ denote the number of states in the abstracted and original plants, respectively. It is well-known that the best algorithms for computing supremal normal sublanguages are exponential in the number of states of the plant; thus, the computation of supremal normal languages based on the original plant is of order $\mathcal{O}(2^{|Q|})$. By contrast, the hierarchical control approach in this paper enables the computation based on the abstracted plant, with a complexity of $\mathcal{O}(2^{|Q_{hi}|})$.

In previous studies on hierarchical supervisory control, considering partial controllability, Schmidt and Breindl [10] introduced a local control consistency to ensure the preservation of supremal controllable sublanguages. For partially observed systems, Komenda and Masopust [6] proposed a modified observation consistency to preserve supremal normal sublanguages. However, these works only address deterministic observations. In contrast, this paper focuses on nondeterministic observations, a key distinction and central challenge in this work.

A common characteristic of networked DESs or DESs

* This work was supported by the Czech Academy of Sciences under RVO 67985840.

* Corresponding author.

Email addresses: smiao585@connect.hkust-gz.edu.cn (Shaowen Miao), komenda@ipm.cz (Jan Komenda), flin@wayne.edu (Feng Lin).

subject to cyber attacks is that observation is not only incomplete (partial) but may also be represented by a set-valued (nondeterministic) observation. The research in [8] has led to network observability and network controllability, which serve as necessary and sufficient conditions for achieving a language in the presence of control and observation delays and losses. More recently, there has been research into cyber attack (CA-) observability and CA-controllability [18], where attackers can edit the control and observation channel of the supervisor loop. Such manipulations of the observation channel known as sensor attacks can result in the supervisor receiving altered observations that may be nondeterministic.

Given that computing supervisors for DESs with incomplete observations involves exponential complexity [16], hierarchical control serves as a crucial approach for reducing the complexity of controller synthesis. This paper presents a unified framework that encompasses all systems with nondeterministic observations, including networked and cyber-attacked DESs, and addresses the hierarchical control problem concerning normality. By doing so, it effectively addresses the non-supremality issue of observability and mitigates the state-space explosion issue in [8,13,14,18]. Also, note that the hierarchical control problem considered in this paper is not comparable to the decentralized control problem in [12,17].

In this paper, starting from Section 3, our starting point is the concept of nondeterministic (ND-) observability, which encompasses both network observability and CA-observability. We first introduce a stronger yet simpler property called strong ND-observability, which coincides with standard observability in the case of (deterministic) natural projections. Furthermore, we present ND-normality, a property preserved by language unions, ensuring that supremal ND-normal sublanguages always exist. We then investigate the hierarchical control of DESs with set-valued observations and show that under a condition we call nondeterministic observation consistency (NOC), ND-normality is preserved between the original and abstracted system, cf. Section 4. Moreover, in Section 5, we show that under a slightly stronger condition, called modified NOC (MNOC), maximal permissiveness is preserved. This means that the supremal ND-normal sublanguages of the abstracted (high-level) plant can be utilized in the original system without losing maximal permissiveness, ensuring both ND-normality and supremality. Finally, in Section 6, we specialize our results to networked DESs and briefly discuss the case of set-valued observations due to sensor attacks.

2 Preliminaries

We first recall some basics of automata and supervisory control theory [3,16]. Let $\mathbb{N}$ denote the set of non-negative integers. For a set A , $|A|$ denotes its cardinality. For an alphabet (finite nonempty set of events) Σ ,

the set of all finite strings over Σ is denoted by Σ^* , including the empty string ε . Let $\Sigma^{\leq N} = \bigcup_{i=0}^N \Sigma^i$, where $N \in \mathbb{N}$. For a string s , $|s|$ denotes its length. By convention, $|\varepsilon| = 0$. A language L is a subset of Σ^* . Let st denote the concatenation of two strings s and t , and let $L_1 L_2 = \{s \in \Sigma^* \mid \exists s_1 \in L_1, \exists s_2 \in L_2 : s = s_1 s_2\}$ denote the concatenation of two languages L_1 and L_2 . The prefix closure of L is $\bar{L} = \{s \in \Sigma^* \mid \exists t \in \Sigma^* : st \in L\}$. Language L is prefix-closed if $L = \bar{L}$.

A (deterministic) projection $R : \Sigma^* \rightarrow \Gamma^*$, with $\Gamma \subseteq \Sigma$, is a morphism (for concatenation) defined by $R(a) = \varepsilon$ for $a \in \Sigma \setminus \Gamma$, $R(a) = a$ for $a \in \Gamma$, and $R(a_1 a_2 \cdots a_n) = R(a_1) R(a_2) \cdots R(a_n)$ for $a_1 a_2 \cdots a_n \in \Sigma^*$. The corresponding inverse map $R^{-1} : \Gamma^* \rightarrow 2^{\Sigma^*}$ is defined as $R^{-1}(t) = \{s \in \Sigma^* \mid R(s) = t\}$. These definitions can be extended to languages in the usual way.

A deterministic finite automaton (DFA) is a quintuple $G = (Q, \Sigma, \delta, q_0, F)$, where Q is a finite set of states, Σ is an alphabet, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of marked states, and $\delta : Q \times \Sigma \rightarrow Q$ is the transition function that can be extended to the domain $Q \times \Sigma^*$ in the usual way. The language generated by G is $L(G) = \{s \in \Sigma^* \mid \delta(q_0, s) \in Q\}$, and the language marked by G is $L_m(G) = \{s \in \Sigma^* \mid \delta(q_0, s) \in F\}$. By definition, $L(G) = \bar{L_m(G)}$ and $L_m(G) \subseteq L(G)$. DFA G is nonblocking if $\bar{L_m(G)} = L(G)$.

For a partially-observed DFA G , the alphabet Σ is partitioned into the observable alphabet Σ_o and the unobservable alphabet $\Sigma_{uo} = \Sigma \setminus \Sigma_o$. The set-valued (nondeterministic) observation is denoted by $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$, where the specific mapping depends on the application scenario (e.g., network DESs or DESs under cyber-attacks) and will be specified later.

A language K is observable w.r.t. $L(G)$ and $P : \Sigma^* \rightarrow \Sigma_o^*$ if for all $s \in \bar{K}$ and $\sigma \in \Sigma$, $s\sigma \in \bar{K}$ implies $P^{-1}P(\bar{K})\sigma \cap L(G) \subseteq \bar{K}$. A language K is normal w.r.t. $L(G)$ and P if $\bar{K} = P^{-1}P(\bar{K}) \cap L(G)$.

The parallel composition of languages $L_1 \subseteq \Sigma_1^*$ and $L_2 \subseteq \Sigma_2^*$ is the language $L_1 \parallel L_2 = P_1^{-1}(L_1) \cap P_2^{-1}(L_2)$ with $P_i : (\Sigma_1 \cup \Sigma_2)^* \rightarrow \Sigma_i^*$ for $i = 1, 2$. In particular, for DFAs G_1 and G_2 , we have $L(G_1 \parallel G_2) = L(G_1) \parallel L(G_2)$ and $L_m(G_1 \parallel G_2) = L_m(G_1) \parallel L_m(G_2)$, where $G_1 \parallel G_2$ denotes the parallel composition of automata G_1 and G_2 [3]. Languages L_1 and L_2 are said to be nonconflicting if $\bar{L_1} \parallel \bar{L_2} = \bar{L_1} \parallel \bar{L_2}$.

3 Nondeterministic Normality

By the definition of set-valued function [1], for an observation $t \in \Sigma_o^*$, we have $\theta^{-1}(t) = \{s' \in \Sigma^* \mid t \in \theta(s')\}$. It

can be extended to language $L \subseteq \Sigma_o^*$ as

$$\begin{aligned}\theta^{-1}(L) &= \bigcup_{t \in L} \theta^{-1}(t) = \bigcup_{t \in L} \{s' \in \Sigma^* \mid t \in \theta(s')\} \\ &= \{s' \in \Sigma^* \mid \theta(s') \cap L \neq \emptyset\}.\end{aligned}$$

Thus, in particular, we have

$$\theta^{-1}\theta(s) = \{s' \in \Sigma^* \mid \theta(s') \cap \theta(s) \neq \emptyset\}. \quad (1)$$

We now define the notion of ND-observability, which works for general nondeterministic observation and can be made specific for a particular situation if needed. It serves as an abstraction of network observability [8] or CA-observability [18]. Each of these, together with controllability [14,18], provides necessary and sufficient conditions for a specification to be achieved by a closed-loop system in networked or cyber-attacked DESs.

Definition 1 (ND-observability) *A language $K \subseteq L(G)$ is ND-observable w.r.t. $L(G)$ and θ if*

$$\begin{aligned}(\forall sa \in \overline{K})(\exists t \in \theta(s))(\forall s' \in \theta^{-1}(t)) \\ [s' \in \overline{K} \wedge s'a \in L(G) \Rightarrow s'a \in \overline{K}].\end{aligned}$$

Note that by Eq. (1), $\exists t \in \theta(s) : s' \in \theta^{-1}(t) \Leftrightarrow s' \in \theta^{-1}\theta(s)$. However, we cannot simplify the formula $(\exists t \in \theta(s))(\forall s' \in \theta^{-1}(t))$ in the definition of ND-observability due to the inconsistency of quantifiers, making it difficult to implement hierarchical control in networked or cyber-attacked DESs. Thus, we consider a stronger condition, which will serve as a basis for ND-normality defined further below.

Definition 2 (Strong ND-Observability) *A language $K \subseteq L(G)$ is strongly ND-observable w.r.t. $L(G)$ and θ if*

$$\begin{aligned}(\forall sa \in \overline{K})(\forall s' \in \theta^{-1}\theta(s)) \\ [s' \in \overline{K} \wedge s'a \in L(G) \Rightarrow s'a \in \overline{K}].\end{aligned}$$

Intuitively, the definition of ND-observability requires the second implication only for s' stemming from a single “branch” t of $\theta(s)$. In contrast, in the definition of strong ND-observability, we need to consider all branches (all $t \in \theta(s)$), namely all $s' \in \theta^{-1}\theta(s)$. This clearly means that strong ND-observability is a stronger condition than ND-observability.

Note that if there are no communication delays, losses, and external attacks, i.e., $\theta(s) = \{P(s)\}$, both ND-observability and strong ND-observability reduce to standard observability. The main difference between

standard observability and strong ND-observability is that, while in the former $s' \in P^{-1}P(s)$ means $P(s') = P(s)$, in the latter $s' \in \theta^{-1}\theta(s)$ means $\theta(s') \cap \theta(s) \neq \emptyset$. Similar to standard observability, strong ND-observability is not preserved by language unions. Therefore, it is natural to introduce the following concept of ND-normality.

Definition 3 (ND-Normality) *A language $K \subseteq L(G)$ is ND-normal w.r.t. $L(G)$ and θ if*

$$(\forall s \in \overline{K})(\forall s' \in \theta^{-1}\theta(s))[s' \in L(G) \Rightarrow s' \in \overline{K}].$$

Note that ND-normality can be equivalently rewritten as $\theta^{-1}\theta(\overline{K}) \cap L(G) \subseteq \overline{K}$, which suggests that it can be checked similarly to normality if observers and inverse projections based on θ have regular representations. This is the case, e.g., in networked DESs [8]. Similar to normality, ND-normality is preserved by language unions. Thus, the unique supremal ND-normal sublanguage of any language $K \subseteq L(G)$, which is the union of all sublanguages of K that are ND-normal, always exists.

Proposition 1 *Let $K_i \subseteq L(G)$ for $i \in I$ be an arbitrary family of languages that are ND-normal w.r.t. $L(G)$ and θ . Then $\bigcup_{i \in I} K_i$ is ND-normal w.r.t. $L(G)$ and θ .*

Proof: Let $s \in \overline{\bigcup_{i \in I} K_i} = \bigcup_{i \in I} \overline{K_i}$ and $s' \in \theta^{-1}\theta(s)$ such that $s' \in L(G)$. We will show that $s' \in \bigcup_{i \in I} \overline{K_i}$. From the assumption, there exists some $i \in I$ such that $s \in \overline{K_i}$. Since K_i is ND-normal, we have $s' \in \overline{K_i} \subseteq \bigcup_{i \in I} \overline{K_i}$, which shows that $\bigcup_{i \in I} K_i$ is ND-normal. ■

We denote the supremal sublanguage of K that is ND-normal w.r.t. $L(G)$ and θ by $\text{supN}_{\text{ND}}(K, L(G), \theta)$. In the following proposition, it is shown that if θ is a morphism for concatenation, i.e., $\theta(sa) = \theta(s)\theta(a)$ for $s \in \Sigma^*$ and $a \in \Sigma$, then ND-normality implies strong ND-observability.

Proposition 2 *Let K be ND-normal w.r.t. $L(G)$ and θ and let $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$ be a morphism for concatenation. Then, K is strongly ND-observable w.r.t. $L(G)$ and θ .*

Proof: To show that K is strongly ND-observable, let $sa \in \overline{K}$ and $s' \in \theta^{-1}\theta(s)$ such that $s' \in \overline{K}$ and $s'a \in L(G)$. Since θ is a morphism, from $\theta(s) \cap \theta(s') \neq \emptyset$, we obtain that $\theta(sa) \cap \theta(s'a) = \theta(s)\theta(a) \cap \theta(s')\theta(a) \neq \emptyset$. Therefore, $s'a \in \theta^{-1}\theta(sa)$, and applying normality to the strings sa and $s'a$ yields $s'a \in \overline{K}$. ■

4 Hierarchical Control of Discrete-Event Systems with Nondeterministic Observations

Since the computation of supervisors for DESs with incomplete observations is of exponential complexity [16],

abstraction is one of the crucial techniques for reducing computational complexity. Furthermore, hierarchical control is a major step towards modular control [7,11] when projections to local alphabets are viewed as abstractions.

It is important to note that network controllability is not required in predictive control of networked DESs [14] and CA-controllability is standard controllability w.r.t. a larger set of uncontrollable events [18]. Given that controllability issues have been widely studied in hierarchical control, we only study the preservation of partial observation-related properties under abstraction.

The abstraction (high-level) alphabet is denoted by $\Sigma_{hi} \subseteq \Sigma$, the high-level abstraction by the (deterministic) projection $\beta: \Sigma^* \rightarrow \Sigma_{hi}^*$, and the corresponding restricted observations and abstractions by $\theta_{hi}: \Sigma_{hi}^* \rightarrow 2^{(\Sigma_{hi} \cap \Sigma_o)^*}$ and $\beta_o: 2^{\Sigma_o^*} \rightarrow 2^{(\Sigma_{hi} \cap \Sigma_o)^*}$, cf. Fig. 1. Note that β_o in Fig. 1 is a language generalization of $\beta_o: \Sigma_o^* \rightarrow (\Sigma_{hi} \cap \Sigma_o)^*$, namely $\beta_o(L) = \bigcup_{t \in L} \beta_o(t)$.

In Theorems 1 and 2 below, we assume that the diagram in Fig. 1 commutes. This assumption is general and sufficient to prove both theorems. However, in Section 6.1, this assumption is not necessary and we use special properties of observation with delays and losses to prove Theorem 1, see Proposition 4.

We provide the following sufficient conditions to ensure the commutativity of the diagram in Fig. 1. Note that, in ordinary DESs with deterministic observation, the commutativity of such a diagram holds trivially by the definition of deterministic observation.

Proposition 3 *If the following conditions*

- (i) $\theta(\varepsilon) = \{\varepsilon\}$,
- (ii) $\theta_{hi}\beta(a) = \beta_o\theta(a)$ for $a \in \Sigma$, and
- (iii) θ is a morphism, i.e., $\theta(sa) = \theta(s)\theta(a)$ for $s \in \Sigma^*$ and $a \in \Sigma$

hold, then the diagram in Fig. 1 commutes, that is, $\theta_{hi}\beta(s) = \beta_o\theta(s)$ for any $s \in \Sigma^*$.

Proof: We show this by structural induction on Σ^* . First, condition (i) is used in the base step $\theta_{hi}\beta(\varepsilon) = \beta_o\theta(\varepsilon)$, where $\beta(\varepsilon) = \varepsilon$. Assuming the equality $\theta_{hi}\beta(s) = \beta_o\theta(s)$ for $s \in \Sigma^*$, we can see that if conditions (i)–(iii) hold, then $\theta_{hi}\beta(sa) = \theta_{hi}\beta(s)\theta_{hi}\beta(a) = \beta_o\theta(s)\theta_{hi}\beta(a) = \beta_o\theta(s)\beta_o\theta(a) = \beta_o\theta(sa)$, where (iii) is used in the first and last equation, and (ii) is used in the second last. ■

Proposition 3 shows that the morphism property (iii), where the mapping θ may yield a language rather than a single string due to nondeterminism, plays a critical role in ensuring the commutativity of observation with

a projection (such as abstraction). This property is also important for another reason: in automata theory, it is well-known that monoid (homo)morphisms preserve the regularity of languages [5].




Fig. 1. Commutative diagram of abstractions and mappings.

We now formulate the hierarchical supervisory control problem, which seeks to determine the conditions under which the supervisor S for the original plant G can be computed solely based on an abstracted plant G_{hi} .

Problem 1 (Hierarchical Supervisory Control)

Let G be a low-level plant over Σ , and let $K_{hi} \subseteq \Sigma_{hi}^*$ be a high-level specification. The abstracted high-level plant G_{hi} is defined by $L(G_{hi}) = \beta(L(G))$ and $L_m(G_{hi}) = \beta(L_m(G))$. The hierarchical supervisory control problem is to determine, based on G_{hi} and K_{hi} , a low-level supervisor S such that $L_m(S/G) = K_{hi} \parallel L_m(G) = \beta^{-1}(K_{hi}) \cap L_m(G)$.

In the remainder of the paper, we use simplified notations $L = L(G)$ and $L_m = L_m(G)$ to make the claims and proofs more concise. We now propose the NOC condition, which will play the role of sufficient conditions for the problem above.

Definition 4 (NOC) A prefix-closed (plant) language $L \subseteq \Sigma^*$ is nondeterministic observation consistent (NOC) w.r.t. β , θ , and θ_{hi} if for all strings $t, t' \in \beta(L)$ with $\theta_{hi}(t) \cap \theta_{hi}(t') \neq \emptyset$, there exist $s, s' \in L$ such that $\beta(s) = t$, $\beta(s') = t'$, and $\theta(s) \cap \theta(s') \neq \emptyset$.

Intuitively, NOC requires that for every two high-level strings t, t' that have a nonempty intersection of observations, their corresponding low-level strings s, s' also have a nonempty intersection of observations.

Note that we will first work with general nondeterministic observation mappings since it is more convenient to derive general forms of sufficient (consistency) conditions. The results will be specialized to networked and cyber-attacked DESs in Section 6. In Theorem 1, we provide an initial solution to Problem 1, which applies when the high-level specification K_{hi} is ND-normal w.r.t. abstracted plant $\beta(L)$ and θ_{hi} . In this case, K_{hi} can be used as a supervisor for the original system, with the resulting closed-loop language $K_{hi} \parallel L_m = \beta^{-1}(K_{hi}) \cap L_m$. A more general solution to Problem 1, when K_{hi} is not ND-normal, is provided in Theorem 2.

Theorem 1 Assume that the diagram in Fig. 1 commutes. If L is NOC w.r.t. β , θ , and θ_{hi} , and K_{hi} and L_m

are nonconflicting, then K_{hi} is ND-normal w.r.t. $\beta(L)$ and θ_{hi} iff $K_{hi}\|L_m$ is ND-normal w.r.t. L and θ .

Proof: $(\Rightarrow)$: Assume that K_{hi} is ND-normal. We need to show that $K_{hi}\|L_m$ is ND-normal. Let $s \in \overline{K_{hi}\|L_m} = \overline{\beta^{-1}(K_{hi}) \cap L_m}$, and suppose $s' \in L$ such that $s' \in \theta^{-1}\theta(s)$, i.e., $\theta(s) \cap \theta(s') \neq \emptyset$. By nonconflictingness of K_{hi} and L_m we have $s \in \overline{\beta^{-1}(K_{hi}) \cap L_m} \subseteq \overline{\beta^{-1}(K_{hi})} = \overline{\beta^{-1}(K_{hi})}$. It needs to be shown that $s' \in \overline{K_{hi}\|L_m}$. Since $\theta(s) \cap \theta(s') \neq \emptyset$, using the commutative diagram in Fig. 1, we have $\beta_o\theta(s') = \theta_{hi}\beta(s')$ and $\beta_o\theta(s) = \theta_{hi}\beta(s)$. Clearly, $\theta(s) \cap \theta(s') \neq \emptyset$ implies that $\beta_o(\theta(s) \cap \theta(s')) \neq \emptyset$ as well. Since the abstraction on observable events β_o is a (deterministic) projection, $\beta_o(\theta(s) \cap \theta(s')) \subseteq \beta_o\theta(s) \cap \beta_o\theta(s')$ always holds. Thus, we also have $\theta_{hi}\beta(s) \cap \theta_{hi}\beta(s') = \beta_o\theta(s) \cap \beta_o\theta(s') \neq \emptyset$. Recall that $s \in \overline{\beta^{-1}(K_{hi})}$, i.e., $\beta(s) \in \overline{K_{hi}}$ and $s' \in L$, i.e., $\beta(s') \in \beta(L)$. Using the ND-normality of K_{hi} , we immediately deduce from $\theta_{hi}\beta(s) \cap \theta_{hi}\beta(s') \neq \emptyset$ that $\beta(s') \in \overline{K_{hi}}$, i.e., $s' \in \beta^{-1}(\overline{K_{hi}})$. Since K_{hi} and L_m are nonconflicting and $s' \in L = \overline{L_m}$, we obtain that $s' \in \overline{\beta^{-1}(K_{hi}) \cap L_m} = \overline{K_{hi}\|L_m}$, which was to be shown.

$(\Leftarrow)$: Assume that $K_{hi}\|L_m = \beta^{-1}(K_{hi}) \cap L_m$ is ND-normal. We need to show that K_{hi} is ND-normal. Let $t \in \overline{K_{hi}} \subseteq \beta(L)$, $t' \in \beta(L)$, and suppose $t' \in \theta_{hi}^{-1}\theta_{hi}(t)$, i.e., $\theta_{hi}(t) \cap \theta_{hi}(t') \neq \emptyset$. We need to show that $t' \in \overline{K_{hi}}$. Since $t, t' \in \beta(L)$ and $\theta_{hi}(t) \cap \theta_{hi}(t') \neq \emptyset$, the NOC (cf. Def. 4) condition implies that there exist $s, s' \in L$ such that $\beta(s) = t$, $\beta(s') = t'$, and $\theta(s) \cap \theta(s') \neq \emptyset$. Since $s' \in L$ and G is nonblocking, we have $s' \in \overline{L_m} = L$. Given $t \in \overline{K_{hi}}$, we also have $s \in \beta^{-1}(\overline{K_{hi}})$, i.e., $s \in \overline{K_{hi}\|L_m} = \overline{K_{hi}\|L_m}$. Recall that $s \in L$. By the ND-normality of $K_{hi}\|L_m$, it follows that $s' \in \overline{K_{hi}\|L_m} \subseteq \overline{\beta^{-1}(K_{hi})}$. Thus, $t' = \beta(s') \in \beta(\overline{\beta^{-1}(K_{hi})}) = \overline{\beta\beta^{-1}(K_{hi})} = \overline{K_{hi}}$. ■

Remark 1 If the specification language K_{hi} is prefix-closed, the nonconflictingness assumption is not needed. If the abstraction β is an L_m -observer [10,15], then K_{hi} and L_m are always nonconflicting. Moreover, it is well known that if the abstraction β satisfies the L_m -observer property, then the abstraction can not have more states than the original plant, which is the key motivation for hierarchical supervisory control.

5 Preservation of Supremal ND-normal Sublanguages

From hierarchical supervisory control with deterministic partial observations [6], it is shown that the observation consistency condition is insufficient to preserve supremality of normal supervisors between the abstracted and the original systems. Moreover, ND-normality is an extension of normality to nondeterministic observations,

thus NOC is also insufficient to preserve supremality of ND-normal supervisors.

In the following, we propose a stronger version of NOC for DESs with nondeterministic observations, namely MNOC, which proves beneficial for the preservation of supremal ND-normal sublanguages.

Definition 5 (MNOC) A prefix-closed (plant) language $L \subseteq \Sigma^*$ is modified nondeterministic observation consistent (MNOC) w.r.t. β , θ , and θ_{hi} if for all strings $s \in L$ and $t' \in \beta(L)$ with $\theta_{hi}\beta(s) \cap \theta_{hi}(t') \neq \emptyset$, there exists $s' \in L$ such that $\beta(s') = t'$ and $\theta(s) \cap \theta(s') \neq \emptyset$.

Refer to Fig. 2 for an illustration of MNOC. It should be clear that MNOC implies NOC. Indeed, If L is MNOC, then for any $t, t' \in \beta(L)$, we can choose an arbitrary $s \in L$ such that $t = \beta(s)$. By MNOC, we get $s' \in L$ with the required properties showing that L is NOC.




Fig. 2. Illustration of MNOC condition.

It is shown below that abstractions satisfying MNOC preserve ND-normality, i.e., abstractions of ND-normal languages remain ND-normal.

Lemma 1 If L is MNOC w.r.t. β , θ , and θ_{hi} , then ND-normality of $M \subseteq L_m$ w.r.t. L and θ implies ND-normality of $\beta(M)$ w.r.t. $\beta(L)$ and θ_{hi} .

Proof: We only need to show that $\theta_{hi}^{-1}\theta_{hi}\beta(\overline{M}) \cap \beta(L) \subseteq \overline{\beta(M)}$, as the converse inclusion always holds. Let $t' \in \theta_{hi}^{-1}\theta_{hi}(\beta(\overline{M}) \cap \beta(L)) \subseteq \beta(L)$. Then, there exists $s \in \overline{M} \subseteq L_m \subseteq L$ such that $t' \in \theta_{hi}^{-1}\theta_{hi}\beta(s)$, meaning $\theta_{hi}(t') \cap \theta_{hi}\beta(s) \neq \emptyset$. By MNOC (cf. Def. 5), there is an $s' \in L$ such that $\beta(s') = t'$ and $\theta(s) \cap \theta(s') \neq \emptyset$. This implies $s' \in \theta^{-1}\theta(s) \cap L \subseteq \theta^{-1}\theta(\overline{M}) \cap L = \overline{M}$. Consequently, $t' = \beta(s') \in \beta(\overline{M})$, thus proving that $\beta(M)$ is ND-normal. ■

We now show that MNOC guarantees that supremal ND-normal supervisors, which are computed based on abstraction, can be applied to the original system while maintaining their supremal ND-normal properties. We first formulate this result for the prefix-closed case to simplify the proof and to ensure the nonconflictingness of the high-level supervisor with the marked language. For the non-prefix-closed specifications, see the appendix.

Theorem 2 Assume that the diagram in Fig. 1 commutes. Let G be a DFA and $K_{hi} \subseteq \beta(L_m)$ be prefix-closed.

If L is MNOC w.r.t. β , θ , and θ_{hi} , then

$$\sup N_{ND}(K_{hi} \| L, L, \theta) = \sup N_{ND}(K_{hi}, \beta(L), \theta_{hi}) \| L.$$

Proof: Let $\sup N_{ND} = \sup N_{ND}(K_{hi}, \beta(L), \theta_{hi})$.

($\subseteq$) : Note that $\sup N_{ND}(K_{hi} \| L_m, L, \theta) \subseteq \beta^{-1}(\sup N_{ND})$ is equivalent to $\beta(\sup N_{ND}(K_{hi} \| L_m, L, \theta)) \subseteq \sup N_{ND}$. The inclusion follows directly from Lemma 1, as $\sup N_{ND}(K_{hi} \| L_m, L, \theta)$ is ND-normal by definition. Thus, leveraging the MNOC assumption, we conclude that $\beta(\sup N_{ND}(K_{hi} \| L_m, L, \theta))$ is ND-normal. This implies that $\beta(\sup N_{ND}(K_{hi} \| L_m, L, \theta)) \subseteq \sup N_{ND}$.

($\supseteq$) : We show that $\sup N_{ND} \| L$ is ND-normal. Let $s \in \sup N_{ND} \| L$ and $s' \in L$ such that $\theta(s) \cap \theta(s') \neq \emptyset$. Then, $\beta(s) \in \sup N_{ND}$ and $\beta(s') \in \beta(L)$. We can show that $\theta_{hi}\beta(s) \cap \theta_{hi}\beta(s') \neq \emptyset$ in exactly the same way as in the proof of Theorem 1, using the commutative diagram in Fig. 1. Given the normality of $\sup N_{ND}$, we obtain $\beta(s) \in \sup N_{ND}$. Thus, $s' \in \beta^{-1}(\sup N_{ND}) \cap L = \sup N_{ND} \| L$. ■

~~We provide the following formula for computing the supremal ND-normal sublanguage of the prefix-closed specification K . The proof is omitted, as it can be easily adapted from the proof of Theorem 3 in [2] by replacing $P(s) \in P(H)$ for $H \subseteq \Sigma^*$ with $\theta(s) \cap \theta(H) \neq \emptyset$ and $P(s) = P(s')$ with $\theta(s) \cap \theta(s') \neq \emptyset$.~~

Theorem 3 $\sup N_{ND}(K, L, \theta) = K \setminus \theta^{-1}\theta(L \setminus K)\Sigma^*$.

A direct way to compute $\sup N_{ND}(K, L, \theta)$ is to continuously remove illegal strings until a fixed point is reached, although this computation may be infinite: Let $H_0 = K$ and $H_{n+1} = H_n \setminus \theta^{-1}\theta(L \setminus H_n)\Sigma^*$. Then $\sup N_{ND}(K, L, \theta) = \bigcap_{n=0}^{\infty} H_n$.

Finally, we discuss the consistency conditions studied in this paper. Similar to the hierarchical supervisory control of ordinary DESs [6], it is unclear if algorithms exist to check these consistency conditions. However, there is a simple event-set condition that implies MNOC. Specifically, if $\Sigma_o \subseteq \Sigma_{hi}$, meaning all observable events are high-level events, then L is always MNOC.

Theorem 4 If $\Sigma_o \subseteq \Sigma_{hi}$, then L is (M)NOC.

Proof: Assume that $\Sigma_o \subseteq \Sigma_{hi}$. Then abstraction β_o from the diagram in Fig. 1 is an identity map. Let $s \in L$ and $t' \in \beta(L)$ be such that $\theta_{hi}\beta(s) \cap \theta_{hi}(t') \neq \emptyset$; equivalently, $t' \in \theta_{hi}^{-1}\theta_{hi}\beta(s)$. Now, take an arbitrary $s' \in L$ such that $\beta(s') = t'$. Since $t' \in \beta(L)$, we know that at least one such s' exists. We will show that any such s' satisfies $s' \in \theta^{-1}\theta(s)$, i.e., $\theta(s') \cap \theta(s) \neq \emptyset$. Given that $\beta(s') = t' \in \theta_{hi}^{-1}\theta_{hi}\beta(s)$, it follows that

$s' \in \beta^{-1}\theta_{hi}^{-1}\theta_{hi}\beta(s) = \beta^{-1}\theta_{hi}^{-1}\theta(s)$ since $\theta_{hi} \circ \beta = \theta$ as β_o is identity. Since $(\theta_{hi} \circ \beta)^{-1} = \beta^{-1} \circ \theta_{hi}^{-1}$, we finally get $s' \in \theta^{-1}\theta(s)$, which was to be shown. ■

From a practical point of view, assuming $\Sigma_o \subseteq \Sigma_{hi}$ means that no observable event should be abstracted away, but only unobservable events, to ensure that the MNOC, and thus also NOC conditions hold. Notice that if we abstract all unobservable events, the high-level supervisor will have complete, but in general still non-deterministic observations.

6 Application to Networked and Cyber-Attacked Discrete-Event Systems

The aim of this section is to specialize the general notions presented in the previous sections to common examples of supervisory control under set-valued observation, specifically to networked [8] and cyber-attacked [18] DESs.

6.1 Networked Discrete-Event Systems

In networked DESs, we define $\Sigma_L \subseteq \Sigma_o$ as the set of events that may be lost in the observation channels. Let $\theta_L : \Sigma^* \rightarrow 2^{\Sigma_o^*}$ denote the observation map with communication losses, which is defined as

$$\begin{aligned} \theta_L(\varepsilon) &= \{\varepsilon\}, \\ \theta_L(a) &= \begin{cases} \{\varepsilon, a\}, & \text{if } a \in \Sigma_L \\ \{a\}, & \text{if } a \in \Sigma_o \setminus \Sigma_L, \\ \{\varepsilon\}, & \text{if } a \in \Sigma_{uo} \end{cases} \quad (2) \\ \theta_L(sa) &= \theta_L(s)\theta_L(a) \text{ for } w \in \Sigma^* \text{ and } a \in \Sigma. \end{aligned}$$

Let $\theta_D : \Sigma^* \rightarrow 2^{\Sigma_o^*}$ denote the observation map with communication delays upper-bounded by $N_o \in \mathbb{N}$ events, which is defined as $\theta_D(s) = \bigcup_{i=0}^{N_o} \{s_{-i}\}$, where s_{-i} is the prefix of s with length of $|s| - i$. Recall from [8] that the observation map with both communication delays and losses $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$, is defined as the composition of θ_D and θ_L : $\theta = \theta_L \circ \theta_D = \bigcup_{i=0}^{N_o} \theta_L(s_{-i})$.

We formulate network normality by specifying $\theta^{-1}\theta(s)$ from the definition of ND-normality. A similar definition can be derived for strong network observability¹.

Note that for $s \in \Sigma^*$, we have $\theta_L^{-1}\theta_L(s) = P_L^{-1}P_L\theta_L(s)$, where $P_L : \Sigma^* \rightarrow (\Sigma_o \setminus \Sigma_L)^*$ is a deterministic projection and $P_L^{-1} : (\Sigma_o \setminus \Sigma_L)^* \rightarrow 2^{\Sigma^*}$ is the corresponding inverse

¹ A special case of strong ND-observability, termed delay observability, was studied in [13,17].

projection. For $t \in \Sigma^*$, we have $\theta_D^{-1}(t) = \{s \Sigma^{\leq N_o}\}$. Note that we have $\theta^{-1} = \theta_D^{-1} \circ \theta_L^{-1}$. Therefore,

$$\theta^{-1}\theta(s) = \bigcup_{m=0}^{\min(N_o, |s|)} P_L^{-1}(\{P_L(s_{-m})\})\Sigma^{\leq N_o}. \quad (3)$$

~~This means that $\theta^{-1}\theta(s)$ can be replaced by $P_L^{-1}P_L\theta_L(s_{-m})w$ for all $m \leq N_o$ and $w \in \Sigma^{\leq N_o}$.~~

Definition 6 (Network Normality) A language K is network normal w.r.t. $L(G)$, θ , and N_o if

$$\forall s \in \bar{K}, \forall s' \in \Sigma^*, \forall m \in \{0, \dots, \min(N_o, |s|)\}, \\ \forall w \in \Sigma^{\leq N_o}: [P_L(s') = P_L(s_{-m}) \wedge s'w \in L] \Rightarrow s'w \in \bar{K}.$$

Note that in the case of no delays ($N_o = 0$, meaning θ_D is an identity map) and no losses ($\theta_L = P_L = P$), the partial observation is given by $\theta = P$, and network normality coincides with standard normality. In this scenario, we also have $m = 0$, and it is sufficient to consider only $P_L(s') = P(s') = P(s)$ as in standard normality. However, the delay operator $\theta_D : \Sigma^* \rightarrow 2^{\Sigma^*}$ given by $\theta_D(s) = \bigcup_{i=0}^{N_o} \{s_{-i}\}$, is not a morphism, i.e., $\theta_D(sa) \neq \theta_D(s)\theta_D(a)$. This means that network normality does not imply strong network observability (by Proposition 2) ~~unless only observation losses are considered (θ_L is a morphism); see the following example.~~

For $N_o \geq 1 \in \mathbb{N}$ and all $s\sigma \in L$, we have $\emptyset \neq \theta_L(s) \subseteq \theta(s) \cap \theta(s\sigma)$. This is because observing s entails no delay, whereas observing $s\sigma$ allows for a delay of one event; both can produce the observations in $\theta_L(s)$. If $K \neq \emptyset$, then $\varepsilon \in \bar{K}$. Thus, by induction and definition of network normality, we have $\bar{K} = L$. Consequently, $N_o \geq 1, K \neq \emptyset \Rightarrow [K \text{ is network normal} \Leftrightarrow \bar{K} = L]$ holds, meaning that network normality implies strong ND-observability in this setting.

Example 1 ~~Consider $\Sigma = \{a, \tau\}$ with $\Sigma_o = \{a\}$ and observation delays with $N_o = 1$, i.e., $\theta(s) = \{P(s), P(s_{-1})\}$. The plant is $L = \{\tau, a, ab, \tau b\}$ and $K = \{\tau, a, ab\}$. Clearly, K is not strongly network observable, because $s = \tau$ and $s' = a$ have overlapping observations ($\varepsilon \in \theta(s) \cap \theta(s')\}$, $s'b \in K$, but $sb \in L \setminus K$. However, K is network normal, because $\theta(sb) \cap \theta(s'b) = \emptyset$, which does not violate network normality.~~

We then discuss the specialization of Theorem 1 to networked DESs. This theorem, when interpreted in the context of networked DESs, states that under the NOC condition, a high-level specification language K_{hi} is network normal w.r.t. abstracted (high-level) plant iff its low-level counterpart $K_{hi}|L$ is network normal w.r.t. the original plant L , if the diagram in Fig. 1 commutes.

Remark 2 It should be noted that the diagram may not commute due to observation delays even if $N_o = 1$ and all events are observable without any losses. Indeed, the commutativity of the diagram means that $\theta_{hi}\beta(s) = \beta_o\theta(s)$. Since $\beta_o = \beta$ and $\theta(s) = \{s, s_{-1}\}$, the commutativity of the diagram requires that $\{\beta(s), [\beta(s)]_{-1}\} = \{\beta(s), \beta(s_{-1})\}$, which does not generally hold. This discrepancy arises because, in the presence of low-level events at the end of s , it may be that $[\beta(s)]_{-1} = \beta(s_{-k})$ for some $k > 1$.

Let us denote $\theta_{hi} = \theta_{L,hi} \circ \theta_{D,hi}$ as the restriction of $\theta = \theta_L \circ \theta_D$ to Σ_{hi} , i.e., $\theta_{L,hi} : \Sigma_{hi}^* \rightarrow 2^{(\Sigma_{hi} \cap \Sigma_o)^*}$ and $\theta_{D,hi} : \Sigma_{hi}^* \rightarrow 2^{\Sigma_{hi}^*}$. First, we show that the diagram in Fig. 1 commutes for the observation losses operator θ_L . This result extends the well-known commutativity of natural projections to the case of observation losses.

Lemma 2 $\theta_{L,hi}\beta(s) = \beta_o\theta_L(s)$.

Proof: This can be shown using structural induction on Σ^* . The base step is straightforward: $\theta_{L,hi}\beta(\varepsilon) = \{\varepsilon\} = \beta_o\theta_L(\varepsilon)$, since $\theta_L(\varepsilon) = \beta(\{\varepsilon\}) = \{\varepsilon\}$. Assume the induction hypothesis: $\theta_{L,hi}\beta(s) = \beta_o\theta_L(s)$. For the induction step, we need to show that $\theta_{L,hi}\beta(sa) = \beta_o\theta_L(sa)$. Since all operators involved are concatenative (unlike θ_D , they are all morphisms for concatenation), it suffices to show that $\theta_{L,hi}\beta(a) = \beta_o\theta_L(a)$ for any $a \in \Sigma$. It is easy to verify: If $a \in \Sigma \setminus \Sigma_{hi}$, we have $\theta_{L,hi}\beta(a) = \theta_{L,hi}(\varepsilon) = \{\varepsilon\} = \beta_o\theta_L(a)$, where $\beta_o\theta_L(a) = \{\varepsilon\}$ no matter of which of the three possible values $\theta_L(a)$ takes, cf. Eq. (2); If $a \in \Sigma_{hi}$, then $\theta_{L,hi}\beta(a) = \theta_{L,hi}(a) = \beta_o\theta_L(a)$, because β_o is the identity on all possible values of $\theta_L(a) = \theta_{L,hi}(a)$. ■

Finally, we need the following simple claim: if two strings have identical projections on a given alphabet, they will also have identical projections on any smaller alphabet.

Lemma 3 If $\theta_L(x) \cap \theta_L(y) \neq \emptyset$ for some $x, y \subseteq \Sigma^*$, then $\theta_{L,hi}\beta(x) \cap \theta_{L,hi}\beta(y) \neq \emptyset$ as well.

Proof: Let $x, y \subseteq \Sigma^*$ be such that $\theta_L(x) \cap \theta_L(y) \neq \emptyset$. Then there exists $s \in \theta_L(x) \cap \theta_L(y)$. We claim that $\beta(s) \in \theta_{L,hi}\beta(x) \cap \theta_{L,hi}\beta(y) \neq \emptyset$. First, $\beta(s) \in \theta_{L,hi}(\theta_L(x) \cap \theta_L(y)) \subseteq \beta_o(\theta_L(x)) \cap \beta_o(\theta_L(y))$ by using the “half-distributivity” property of the natural projection $\beta_o : \Sigma_o^* \rightarrow (\Sigma_o \cap \Sigma_{hi})^*$, which always holds. Now, by Lemma 2, we obtain $\beta(s) \in \theta_{L,hi}\beta(x) \cap \theta_{L,hi}\beta(y)$, making this set nonempty, which was to be shown. ■

We are now ready to prove the following result.

Proposition 4 Theorem 1 holds for $\theta = \theta_L \circ \theta_D = \bigcup_{i=0}^{N_o} \theta_L(s_{-i})$ by assuming only the NOC condition, without requiring the commutativity of the diagram in Fig. 1.

Proof: The commutativity of the diagram in Fig. 1 is used to show that overlapping observations in the original system imply overlapping observations in the abstraction. Specifically, in the first part ($\Rightarrow$) of the proof of Theorem 1, it needs to be established that $\theta(s) \cap \theta(s') \neq \emptyset$, i.e., $\theta_L \theta_D(s) \cap \theta_L \theta_D(s') \neq \emptyset$, implies $\theta_{hi} \beta(s) \cap \theta_{hi} \beta(s') \neq \emptyset$. It is shown below that this holds without the commutativity of the diagram.

By taking $x = \theta_D(s)$ and $y = \theta_D(s')$ in Lemma 3, we obtain that $\theta_L \theta_D(s) \cap \theta_L \theta_D(s') \neq \emptyset$ implies $\theta_{L,hi} \beta(x) \cap \theta_{L,hi} \beta(y) \neq \emptyset$ as well. To prove that $\theta_{L,hi} \theta_D \beta(s) \cap \theta_{L,hi} \theta_D \beta(s') \neq \emptyset$, it suffices to show that for all $w \in \Sigma^*$, we have $\beta \theta_D(w) \subseteq \theta_{D,hi} \beta(w)$. By definition, $\beta \theta_D(w) = \{\beta(w), \beta(w_{-1}), \dots, \beta(w_{-N_o})\}$ and $\theta_{D,hi} \beta(w) = \{\beta(w), [\beta(w)]_{-1}, \dots, [\beta(w)]_{-N_o}\}$. The inclusion is then implied by $(\forall w \in L)(\forall i \in \mathbb{N} : i \leq N_o)(\exists j \in \mathbb{N} : j \leq N_o) : \beta(w_{-i}) = [\beta(w)]_{-j}$. There is always such a $j \leq i \leq N_o$ to ensure the inclusion. Therefore, Theorem 1 holds for networked DESs even without the commutativity of the diagram in Fig. 1. ■

Note that Theorem 2 enables the computation of the supremal network normal sublanguages based on abstracted ones if MNOC holds. The above argument applied to Theorem 2 means that it holds even without the commutativity of the diagram in Fig. 1.

6.2 Cyber-Attacked Discrete-Event Systems

Another example of set-valued (nondeterministic) partial observations has been studied in [18], where the observation channel is attacked by a malicious intruder who can modify the information from sensors. This is represented by a nondeterministic resulting (corrupted) observation map $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$, including deletion, insertion, and replacement attacks. The authors proposed necessary and sufficient conditions on specification languages for the existence of a supervisor under cyber-attacks, ensuring that the upper bound language is equal to the specification. The key condition in the case of sensor attacks is the so-called CA-observability, which is another instance of ND-observability.

Moreover, let us recall from [18] that $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$ is defined as a morphism. Specifically, it is assumed that an attacker can change every event $\sigma \in \Sigma_o^a \subseteq \Sigma_o$ to any string in a language $A_\sigma \subseteq \Sigma_o^*$. For simplicity, we assume that the attack does not distinguish between different transitions under the same event. Events $\sigma \in \Sigma_o^a$ are referred to as attackable events. It is assumed that only observable events can be attacked since the attacker does not have more sensors than those already present in the system. After a string $s = \sigma_1 \sigma_2 \dots \sigma_{|s|} \in L(G)$ occurs in G , the set of possible strings after attacks, denoted by $\theta_A(s)$, where $\theta_A : L(G) \rightarrow 2^{\Sigma_o^*}$, is calculated as

$$\theta_A(s) = L_1 L_2 \dots L_{|s|},$$

where

$$L_i = \begin{cases} \{\sigma_i\}, & \text{if } \sigma_i \notin \Sigma_o^a \\ A_{\sigma_i}, & \text{if } \sigma_i \in \Sigma_o^a \end{cases}. \quad (4)$$

The resulting nondeterministic partial observations for systems under cyber-attacks are given by $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$, defined as $\theta(s) = P(\theta_A(s))$. Note that since θ_A is defined as a morphism, the overall nondeterministic observation $\theta = P \circ \theta_A$ is also a morphism, being the composition of two morphisms. This means that we have two stronger conditions: strong CA-observability and CA-normality by specializing the ND-observability and ND-normality to this map θ . According to Proposition 2, CA-normality implies CA-observability. The results presented in this paper naturally apply to the specialized map θ above. For example, supremal CA-normal sublanguages can be computed based on abstractions under the presented conditions in Theorem 2. However, there is one restriction we should be aware of, namely the commutativity of the diagram in Fig. 1. Three sufficient conditions are stated to check the commutativity of this diagram in Proposition 3, and whether the first and the second conditions hold depends on the specific attacks.

Below, we provide an illustrative example, where only nondeterministic replacements are possible attacks. Specifically, we consider as θ_A the replacement morphism $R : \Sigma_o^* \rightarrow 2^{\Sigma_o^*}$, generated by $R : \Sigma_o \rightarrow 2^{\Sigma_o}$, where observable events $a \in \Sigma_o^a$ that are not fixed points of R are replaced by some events in $R(a) = A_a \subseteq \Sigma_o$. In this case, we need to study the commutativity of the diagram in Fig. 1, which will enable us to apply the results from Section 5 to compute the supervisor based on the abstracted plant.

Since the replacement is defined as a morphism, ensuring commutativity of the diagram for replacements requires assuming consistency similar to conditions (i) and (ii) in Proposition 3. Specifically, condition (i), $R(\varepsilon) = \{\varepsilon\}$, is always satisfied by the definition of possible attacks that only react to observations. Hence, we only need to guarantee that (ii), $R_{hi} \beta(a) = \beta_o R(a)$, where R_{hi} is a restriction of R on Σ_{hi}^* , i.e., $R_{hi} : (\Sigma_o \cap \Sigma_{hi})^* \rightarrow 2^{(\Sigma_o \cap \Sigma_{hi})^*}$. Two cases need to be distinguished: For $a \in \Sigma_{hi}$, condition (ii) amounts to $R_{hi}(a) = \beta_o R(a)$. This holds if, for $a \in \Sigma_{hi}$, we have $R(a) \subseteq \Sigma_{hi}$, i.e., $R(\Sigma_{hi}) \subseteq \Sigma_{hi}$. This implies that the replacement should preserve the abstraction status of events, meaning that high-level events should only be replaced by high-level events. This restricts the choice of suitable abstractions. For $a \in \Sigma \setminus \Sigma_{hi}$, we have $\beta(a) = \varepsilon$. Thus, condition (ii) becomes $R_{hi}(\varepsilon) = \{\varepsilon\} = \beta_o R(a)$. This holds if, for $a \in \Sigma \setminus \Sigma_{hi}$, we have $R(a) \subseteq \Sigma \setminus \Sigma_{hi}$. Note that in the case of deterministic replacement attacks, these two conditions can be unified to $a \in \Sigma_{hi}$ iff $R(a) \in \Sigma_{hi}$, which is a restriction on the choice of the abstraction.

Example 2 Consider the plant $L = \{\overline{a\sigma b}, \overline{\tau a b}, \overline{d\sigma b}\}$

with unobservable events $\Sigma_{uo} = \{\tau, \sigma\}$ and the abstraction β given by $\Sigma_{hi} = \Sigma \setminus \{\sigma\}$. Since $\Sigma_o \subseteq \Sigma_{hi}$, both MNOC and NOC conditions hold by Theorem 4. The specification is $K_{hi} = \{ab, \tau ab, d\}$, meaning that b should be disabled after d but should be retained after a as well after τa . Let $\Sigma_o^a = \{a\}$ and $R(a) = A_a = \{a, d\}$, indicating that a may (nondeterministically) be replaced by d . The diagram commutes for the chosen abstraction because, for $a \in \Sigma \setminus \Sigma_{hi}$, we have $R(a) \subseteq \Sigma \setminus \Sigma_{hi}$, and for $a \in \Sigma_{hi}$, we have also $R(a) \subseteq \Sigma_{hi}$. Thus, the considered replacement attack commutes with the abstraction. Despite $K_{hi} = \overline{K_{hi}}$ being observable w.r.t. $\beta(L)$ and P , in the presence of the replacement attack R , K_{hi} is not CA-observable w.r.t. $\beta(L)$ and $\theta = P \circ R$, where $P : \Sigma^* \rightarrow \Sigma_o^*$ is the standard projection (partial observation). For example, consider $s = ab \in K_{hi}$, but $s' = db \in \beta(L) \setminus K_{hi}$, even though $s' \in \theta^{-1}\theta(s)$ because $db \in \theta(s) \cap \theta(s')$, which violates CA-normality of K_{hi} w.r.t. $\beta(L)$ and θ . ~~Since θ is a morphism for concatenation, Proposition 2 implies that K_{hi} is not CA-observable w.r.t. $\beta(L)$ and θ either.~~

Given that the replacement attack R , and hence θ , commutes with the abstraction β , both Theorem 1 and Theorem 2 can be directly applied. In accordance with Theorem 1, $K_{hi} \parallel L_m$ is not CA-normal w.r.t. L and θ either. Theorem 2 provides a method to compute the supremal CA-normal sublanguage, denoted $\sup N_{CA}$, based on the abstracted plant $\beta(L)$. Specifically, we have: $\sup N_{CA}(K_{hi} \parallel L, L, \theta) = \sup N_{CA}(K_{hi}, \beta(L), \theta_{hi}) \parallel L$. In the given example, using Theorem 2, it can be checked these two languages equal to $\{a\sigma, \tau a, \tau d, d\sigma\}$. This means that b needs to be disabled not only after observing a but also after observing d , which may be replaced by a . Note that, by Proposition 2, CA-normal languages are CA-observable because, by definition, $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$, and the replacement map R is a morphism.

Finally, we recall that CA-controllability of [18] is just standard controllability w.r.t. a larger set of uncontrollable events. Thus, our results can be combined with the well-known ones dealing with controllability [10].

7 Conclusions

In this paper, for DESs with nondeterministic observations, we take the first step in studying the conditions under which supervisors can be computed on simplified (abstracted) plants and then implemented in the original system without changing the closed-loop language. We show that the commutativity of the diagram in Fig. 1 and whether the nondeterministic observation map θ is a morphism, are the fundamental problems in employing hierarchical control for networked and cyber-attacked DESs. Investigating general nondeterministic observations and subsequently specializing in their special instances is the appropriate approach. This is because directly investigating networked or cyber-attacked DESs

would be significantly more technical compared to specializing in the unified results advocated in this paper. In future work, we plan to extend the results presented in this paper to modular, and more generally, to coordination supervisory control. This will involve considering the projection to local events, potentially enriched with coordinator events, as a form of abstraction.

Appendix

In this appendix, we extend the well-known distributivity of inverse projections and the (half) distributivity of projections to accommodate nondeterministic observations. This is needed to extend Theorem 2 to non-prefix-closed specifications K_{hi} .

Lemma 4 Let $\theta : \Sigma^* \rightarrow 2^{\Sigma_o^*}$ be a nondeterministic observation map and $\beta : \Sigma^* \rightarrow \Sigma_{hi}^*$ be an abstraction projection such that the diagram in Fig. 1 commutes. Then we have:

(i) : For any $M \subseteq \Sigma_{hi}^*$ and $N \subseteq \Sigma^*$, we have $\theta(M \parallel N) \subseteq \theta_{hi}(M) \parallel \theta(N)$;

(ii) : For any $M \subseteq (\Sigma_{hi} \cap \Sigma_o)^*$ and $N \subseteq \Sigma_o^*$, we have $\theta^{-1}(M \parallel N) \subseteq \theta_{hi}^{-1}(M) \parallel \theta^{-1}(N)$.

Proof: (i) : Let $M \subseteq \Sigma_{hi}^*$ and $N \subseteq \Sigma^*$ and let $t \in \theta(M \parallel N)$. This means there exists $s \in M \parallel N = \beta^{-1}(M) \cap N$ such that $t \in \theta(s)$. Then, $t \in \theta(N)$ and $t \in \theta(\beta^{-1}(M))$. It remains to show that $\beta_o(t) \in \theta_{hi}(M)$, because then $t \in \beta_o^{-1}\theta_{hi}(M) \cap \theta(N) = \theta_{hi}(M) \parallel \theta(N)$, which was to be shown. Note that $t \in \theta\beta^{-1}(M)$ means that $t \in \theta(w)$ for some $w \in \Sigma^*$ such that $\beta(w) \in M$. Thus, we have $\beta_o(t) \in \beta_o\theta(w) = \theta_{hi}\beta(w) \subseteq \theta_{hi}(M)$.

(ii) : Let $M \subseteq (\Sigma_{hi} \cap \Sigma_o)^*$ and $N \subseteq \Sigma_o^*$ and let $s \in \theta^{-1}(M \parallel N)$. Then $\theta(s) \cap (M \parallel N) \neq \emptyset$. We need to show that $s \in \theta_{hi}^{-1}(M) \parallel \theta^{-1}(N)$, i.e., $s \in \beta^{-1}\theta_{hi}^{-1}(M) \cap \theta^{-1}(N)$, where $\beta^{-1} \circ \theta_{hi}^{-1} = (\theta_{hi} \circ \beta)^{-1}$. Otherwise stated, it must be shown that $\theta_{hi}\beta(s) \cap M \neq \emptyset$ and $\theta(s) \cap N \neq \emptyset$. From $\theta(s) \cap (M \parallel N) = \theta(s) \cap \beta_o^{-1}(M) \cap N \neq \emptyset$, we immediately have $\theta(s) \cap N \neq \emptyset$. Now let us consider the string $\beta(s)$. Since $\theta(s) \cap \beta_o^{-1}(M) \neq \emptyset$, we have, by applying the natural projection β_o and noting $\beta\beta^{-1}(M) = M$, that $\beta_o\theta(s) \cap M \neq \emptyset$. Note that $\beta_o\theta(s) = \theta_{hi}\beta(s)$. Thus, $\theta_{hi}\beta(s) \cap M \neq \emptyset$ as well. ■

It should be obvious that both points in Lemma 4 can be combined. Thus, for languages M and N , we have $\theta^{-1}\theta(M \parallel N) \subseteq \theta_{hi}^{-1}\theta_{hi}(M) \parallel \theta^{-1}\theta(N)$.

We now show that MNOC guarantees that supremal ND-normal supervisors, which are computed based on abstraction, can be applied to the original system while maintaining their supremal ND-normal properties.

Theorem 5 Assume that the diagram in Fig. 1 commutes. Let G be a nonblocking DFA and let $K_{hi} \subseteq \beta(L_m)$. If L is MNOC w.r.t. β , θ , and θ_{hi} , and $\text{supN}_{\text{ND}}(K_{hi}, \beta(L), \theta_{hi})$ and L_m are nonconflicting, then

$$\text{supN}_{\text{ND}}(K_{hi} \| L_m, L, \theta) = \text{supN}_{\text{ND}}(K_{hi}, \beta(L), \theta_{hi}) \| L_m.$$

Proof: Let $\text{supN}_{\text{ND}} = \text{supN}_{\text{ND}}(K_{hi}, \beta(L), \theta_{hi})$.

($\subseteq$) : The proof of this inclusion is exactly the same as in the prefix-closed case (Theorem 2).

($\supseteq$) : We show that $\text{supN}_{\text{ND}}(K_{hi}, \beta(L), \theta_{hi}) \| L_m$ is ND-normal w.r.t. L and θ . We derive that

$$\begin{aligned} & \theta^{-1} \theta (\text{supN}_{\text{ND}} \| L_m) \cap L \\ &= \theta^{-1} \theta (\text{supN}_{\text{ND}} \| L_m) \| \overline{L_m} \quad (\text{by } G \text{ being nonblocking}) \\ &\subseteq \theta_{hi}^{-1} \theta_{hi} (\text{supN}_{\text{ND}}) \| \theta^{-1} \theta (\overline{L_m}) \| \overline{L_m} \\ &\quad (\text{by Lemma 4 and the nonconflictingness}) \\ &= \theta_{hi}^{-1} \theta_{hi} (\text{supN}_{\text{ND}}) \| \overline{L_m} \\ &\quad (\text{by Eq. (1) and } \overline{L_m} \subseteq \theta^{-1} \theta (\overline{L_m})) \\ &= \theta_{hi}^{-1} \theta_{hi} (\text{supN}_{\text{ND}}) \| \overline{L_m} \| \beta(\overline{L_m}) \\ &\quad (\text{by } \overline{L_m} = \overline{L_m} \| \beta(\overline{L_m})) \\ &= \theta_{hi}^{-1} \theta_{hi} (\text{supN}_{\text{ND}}) \| \overline{L_m} \| \beta(L) \\ &= \text{supN}_{\text{ND}} \| \overline{L_m} \quad (\text{by the ND-normality of } \text{supN}_{\text{ND}}) \\ &= \text{supN}_{\text{ND}} \| L_m. \quad (\text{by the nonconflictingness}) \end{aligned}$$

Thus, $\text{supN}_{\text{ND}}(K_{hi} \| L_m, L, \theta) \supseteq \text{supN}_{\text{ND}} \| L_m$ holds. ■

References

- [1] Jean-Pierre Aubin. *Set-valued analysis*. Springer, 1999.
- [2] Robert D. Brandt, Vijay Garg, Ratnesh Kumar, Feng Lin, Steven I. Marcus, and W. Murray Wonham. Formulas for calculating supremal controllable and normal sublanguages. *Systems & Control Letters*, 15(2):111–117, 1990.
- [3] Christos G. Cassandras and Stéphane Lafortune. *Introduction to discrete event systems*. Springer, 2021.
- [4] Christoforos N. Hadjicostis, Stéphane Lafortune, Feng Lin, and Rong Su. Cybersecurity and supervisory control: A tutorial on robust state estimation, attack synthesis, and resilient control. In *Proc. IEEE 61st Conference on Decision and Control*, pages 3020–3040, 2022.
- [5] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 2006.
- [6] Jan Komenda and Tomáš Masopust. Hierarchical supervisory control under partial observation: Normality. *IEEE Transactions on Automatic Control*, 68(12):7286–7298, 2023.
- [7] Jan Komenda and Tomáš Masopust. Supervisory control of modular discrete-event systems under partial observation: Normality. *IEEE Transactions on Automatic Control*, 69(6):3796–3807, 2024.
- [8] Feng Lin. Control of networked discrete event systems: Dealing with communication delays and losses. *SIAM Journal on Control and Optimization*, 52(2):1276–1298, 2014.
- [9] Feng Lin and W. Murray Wonham. Decentralized supervisory control of discrete-event systems. *Information sciences*, 44(3):199–224, 1988.
- [10] Klaus Schmidt and Christian Breindl. Maximally permissive hierarchical control of decentralized discrete event systems. *IEEE Transactions on Automatic Control*, 56(4):723–737, 2011.
- [11] Klaus Schmidt, Thomas Moor, and Sebastian Perk. Nonblocking hierarchical control of decentralized discrete event systems. *IEEE Transactions on Automatic Control*, 53(10):2252–2265, 2008.
- [12] Shaolong Shu and Feng Lin. Decentralized control of networked discrete event systems with communication delays. *Automatica*, 50(8):2108–2112, 2014.
- [13] Shaolong Shu and Feng Lin. Deterministic networked control of discrete event systems with nondeterministic communication delays. *IEEE Transactions on Automatic Control*, 62(1):190–205, 2017.
- [14] Shaolong Shu and Feng Lin. Predictive networked control of discrete event systems. *IEEE Transactions on Automatic Control*, 62(9):4698–4705, 2017.
- [15] Kai Cheung Wong and W. Murray Wonham. Hierarchical control of discrete-event systems. *Discrete Event Dynamic Systems*, 6:241–273, 1996.
- [16] W. Murray Wonham and Kai Cai. *Supervisory control of discrete-event systems*. Springer, 2019.
- [17] Pan Xu, Shaolong Shu, and Feng Lin. Nonblocking and deterministic decentralized control for networked discrete event systems under communication delays. *Discrete Event Dynamic Systems*, 31:295–315, 2021.
- [18] Shengbao Zheng, Shaolong Shu, and Feng Lin. Modeling and control of discrete event systems under joint sensor-actuator cyber attacks. *IEEE Transactions on Control of Network Systems*, 11(2):782–794, 2024.
- [19] Hao Zhong and W. Murray Wonham. On the consistency of hierarchical supervision in discrete-event systems. *IEEE Transactions on automatic Control*, 35(10):1125–1134, 1990.
- [20] Yuting Zhu, Liyong Lin, Ruochen Tai, and Rong Su. Overview of networked supervisory control with imperfect communication channels. *Discrete Event Dynamic Systems*, 33(1):25–61, 2023.