

Complexity of Verifying Nonblockingness in Modular Supervisory Control

Tomáš Masopust

Abstract—Complexity analysis becomes a common task in supervisory control. However, many results of interest are spread across different topics. The aim of this paper is to bring several interesting results from complexity theory and to illustrate their relevance to supervisory control by proving new nontrivial results concerning nonblockingness in modular supervisory control of discrete event systems modeled by finite automata.

Index Terms—Discrete event systems; Finite automata; Modular control; Complexity.

I. INTRODUCTION

Nonblockingness is an important property of discrete event systems ensuring that every task can be completed. It has therefore been intensively studied in the literature [1]. An automaton (deterministic or nondeterministic) is *nonblocking* if every sequence of events generated by the automaton can be extended to a marked sequence. Given a set of nonblocking automata, the *modular nonblockingness problem* asks whether the parallel composition of all the automata of the set results in a nonblocking automaton.

The property is easy to verify for a deterministic automaton (DFA) as we discuss in Theorem 2. However, if the automaton is nondeterministic (NFA) or a set of nonblocking DFAs is considered, the verification becomes computationally more demanding. We study the complexity in Theorems 3 and 6, respectively. A result relevant to timed discrete event systems is provided in Theorem 8.

So far, no efficient (polynomial) algorithm for verifying modular nonblockingness is known. In the light of the results of this paper, it is not surprising. The problem is complete for the complexity class for which the experts believe that no efficient algorithms exist. Therefore it is unlikely that there is an efficient algorithm solving the problem in general. However, there can still be optimization methods or algorithms working well for most of the practical cases. For instance, Malik [2] has recently shown that explicit model checking algorithms without any special data structures work well on standard computers for several practical systems with 100 million states.

The aim of this paper is to bring and apply some of the interesting results from automata and complexity theory to the nonblockingness verification problem.

II. PRELIMINARIES

An *alphabet*, Σ , is a finite nonempty set. The elements of an alphabet are called *events*. A *string* over Σ is a finite sequence

Supported by the DFG in Emmy Noether grant KR 4381/1-1 (DIAMOND). T. Masopust (tomas.masopust@tu-dresden.de) is with the Institute of Theoretical Computer Science and cfaed, TU Dresden, Germany

(*catenation*) of events, for example, 001 is a string over $\{0, 1\}$. Let Σ^* denote the set of all finite strings over Σ ; the *empty string* is denoted by ε .

A *nondeterministic finite automaton* (NFA) over an alphabet Σ is a structure $\mathcal{A} = (Q, \Sigma, \delta, I, F)$, where Q is the finite nonempty set of states, $I \subseteq Q$ is the nonempty set of initial states, $F \subseteq Q$ is the set of accepting (marked) states, and $\delta: Q \times \Sigma \rightarrow 2^Q$ is the transition function that can be extended to the domain $2^Q \times \Sigma^*$ by induction. The *language generated by \mathcal{A}* is the set $L(\mathcal{A}) = \{w \in \Sigma^* \mid \delta(I, w) \neq \emptyset\}$ and the *language marked by \mathcal{A}* is the set $L_m(\mathcal{A}) = \{w \in \Sigma^* \mid \delta(I, w) \cap F \neq \emptyset\}$. Equivalently, the transition function is a relation $\delta \subseteq Q \times \Sigma \times Q$. Then the meaning of $\delta(q, a) = \{s, t\}$ is that there are two transitions (q, a, s) and (q, a, t) .

The prefix closure of a language L is the set $\bar{L} = \{w \in \Sigma^* \mid \text{there exists } u \in \Sigma^* \text{ s.t. } wu \in L\}$; L is *prefix-closed* if $L = \bar{L}$. Obviously, $L_m(\mathcal{A}) \subseteq L(\mathcal{A})$ and $L(\mathcal{A})$ is prefix-closed.

NFA \mathcal{A} is *deterministic* (DFA) if it has a unique initial state, $|I| = 1$, and no nondeterministic transitions, $|\delta(q, a)| \leq 1$ for every $q \in Q$ and $a \in \Sigma$. For DFAs, we identify singletons with their elements and simply write p instead of $\{p\}$. Specifically, we write $\delta(q, a) = p$ instead of $\delta(q, a) = \{p\}$.

For every NFA \mathcal{A} there exists a DFA \mathcal{B} such that $L_m(\mathcal{B}) = L_m(\mathcal{A})$ and $L(\mathcal{B}) = L(\mathcal{A})$.

Let Σ and Γ be alphabets, and let $f: \Sigma^* \rightarrow \Gamma^*$ be a map. Then f is a *morphism* (for catenation) if $f(xy) = f(x)f(y)$ for every $x, y \in \Sigma^*$. Let $\Sigma_o \subseteq \Sigma$ be alphabets. A *projection* P from Σ^* to Σ_o^* is a morphism defined by $P(a) = \varepsilon$ for $a \in \Sigma \setminus \Sigma_o$, and $P(a) = a$ for $a \in \Sigma_o$. The action of projection P on a string $w \in \Sigma^*$ is to erase all events from w that do not belong to Σ_o . The *inverse image* of P , denoted P^{-1} , is defined as $P^{-1}(s) = \{w \in \Sigma^* \mid P(w) = s\}$. The definitions can readily be extended to languages.

Let L_i be a language over Σ_i , $i = 1, \dots, n$. The *parallel composition* of $(L_i)_{i=1}^n$ is defined by $\parallel_{i=1}^n L_i = \bigcap_{i=1}^n P_i^{-1}(L_i)$, where P_i is a projection from $(\bigcup_{i=1}^n \Sigma_i)^*$ to Σ_i^* . For $i = 1, 2$, let $\mathcal{A}_i = (Q_i, \Sigma_i, \delta_i, I_i, F_i)$ be NFAs. The *parallel composition of \mathcal{A}_1 and \mathcal{A}_2* is defined as the accessible part of the NFA $(Q_1 \times Q_2, \Sigma_1 \cup \Sigma_2, \delta, I_1 \times I_2, F_1 \times F_2)$, where

$$\delta((x, y), e) = \begin{cases} \delta_1(x, e) \times \delta_2(y, e) & \text{if } e \in \Sigma_1 \cap \Sigma_2 \\ \delta_1(x, e) \times \{y\} & \text{if } e \in \Sigma_1 \setminus \Sigma_2 \\ \{x\} \times \delta_2(y, e) & \text{if } e \in \Sigma_2 \setminus \Sigma_1 \end{cases}$$

The parallel composition of DFAs is a DFA [3]. The relationship between the definitions is $L(\mathcal{A}_1 \parallel \mathcal{A}_2) = L(\mathcal{A}_1) \parallel L(\mathcal{A}_2)$ and $L_m(\mathcal{A}_1 \parallel \mathcal{A}_2) = L_m(\mathcal{A}_1) \parallel L_m(\mathcal{A}_2)$.

An NFA \mathcal{A} is *nonblocking* if $L_m(\mathcal{A}) = L(\mathcal{A})$. The inclusion $L_m(\mathcal{A}) \subseteq L(\mathcal{A})$ always holds.

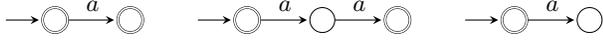


Fig. 1. DFAs \mathcal{A}_1 and \mathcal{A}_2 and their parallel composition $\mathcal{A}_1 \parallel \mathcal{A}_2$

To show that a composition of nonblocking automata can be blocking, let \mathcal{A}_1 and \mathcal{A}_2 be DFAs over $\{a\}$ depicted in Fig. 1. Both automata are nonblocking but their parallel composition is blocking, because a cannot be extended to a marked string.

We now briefly recall the basic notions of complexity theory. For all unexplained notions, the reader is referred to the literature [4], [5].

There are two complexity measures: space and time. The class $\text{NSPACE}(f(n))$ denotes the class of all problems decidable by a nondeterministic Turing Machine (TM) (a nondeterministic algorithm) in space $O(f(n))$ for an input of size n . The class $\text{NL} = \text{NSPACE}(\log n)$ is thus the class of all problems decidable by a nondeterministic TM in logarithmic space, and $\text{PSPACE} = \text{NPSpace} = \cup_{k \in \mathbb{N}} \text{NSPACE}(n^k)$ is the class of all problems decidable by a (nondeterministic) TM in polynomial space. The space required to store the input and output is not considered in space complexity.

The class P (NP) denotes the class of all problems decidable by a (nondeterministic) TM in polynomial time.

The hierarchy of classes is $\text{NL} \subseteq \text{P} \subseteq \text{NP} \subseteq \text{PSPACE}$. Even though $\text{NL} \subsetneq \text{PSPACE}$, the strictness of any other inclusion is unknown. The (non)strictness of these inclusions is the most interesting and important open problem of complexity theory.

The classes NL and NP are defined in terms of a nondeterministic TM (a nondeterministic algorithm). Although for every nondeterministic TM there is an equivalent deterministic TM, the difference is in complexity. A typical nondeterministic step of a nondeterministic algorithm is “choose $x \in X$ ”. Deterministically, one can imagine to check all the possibilities for x one by one. Nondeterministically, the situation is different. There are two basic views how a nondeterministic algorithm performs a nondeterministic step. The first view is that the algorithm “guesses” the right value of x that eventually leads to a success, if such a value exists. The other view is that the algorithm makes a copy of itself for every nondeterministic step with different value of x in every copy. For a nondeterministic step “choose $x \in \{1, \dots, 100\}$ ”, 100 copies of the algorithm would be created, where the value of x in the i th copy is $x = i$. The nondeterministic algorithm is successful if at least one of the copies is successful. In this view, the time (space) complexity is the maximum of time (space) required by a copy.

Example 1: Let $G = (V, E, s, t)$ be a directed graph with $s, t \in V$ the source and target nodes. The *graph reachability* problem asks whether the target node t is reachable from the source node s in G . The problem belongs to NL [4].

To show this, we describe a nondeterministic algorithm (Algorithm 1) that solves the graph reachability problem in logarithmic space. Algorithm 1 is nondeterministic because of the nondeterministic step on line 3. Following the first view, the algorithm correctly guesses the edges that lead from node s to node t , if such a path exists. Following the second view, the algorithm forks for every possible edge on line 3. If any

of the copies ever reaches t , the copy returns `true`, which is then the overall answer. The variable *numSteps* counts the number of steps and terminates the cycle if it is bigger than the number of nodes. This is fine because if there is a path from s to t , then there is a path of length at most $|V| - 1$.

Algorithm 1: (Graph reachability)

Input : A directed graph $G = (V, E, s, t)$
Output: `true` iff t is reachable from s in G

- 1 $k := s$; *numSteps* := 0
- 2 **repeat**
- 3 choose k' such that $(k, k') \in E$
- 4 $k := k'$; *numSteps* := *numSteps* + 1
- 5 **until** $k = t$ **or** *numSteps* > $|V| - 1$
- 6 **if** $k = t$ **then return** `true`
- 7 **return** `false`

It remains to show that Algorithm 1 works in logarithmic space. Since the input is not considered in space complexity, the only space required by the algorithm is the space to store $k, k', |V| - 1$ and *numSteps*. However, *numSteps* is a binary number, bounded by $|V|$, which requires at most $\lceil \log |V| \rceil$ digits. Similarly, k is a pointer to the position in the input, where the actual value of k is stored. Thus, it is again a binary number with at most $\lceil \log(|V| + |E|) \rceil$ digits. Similarly for k' and $|V| - 1$. \square

A problem is PSPACE -complete if it can be solved using only polynomial space (*membership* in PSPACE) and if every problem that can be solved in polynomial space can be reduced (transformed) to it in polynomial time (*PSPACE-hardness*). PSPACE -complete problems are therefore the hardest problems in PSPACE . Similarly for the other complexity classes, with only a different requirement on the reduction. Namely, to prove NL -hardness, the reduction has to be in deterministic logarithmic space, and to prove NP -hardness, the reduction has to be in polynomial time (as well as for PSPACE -hardness).

For instance, satisfiability of formulae in conjunctive normal form¹ (3CNF) is an NP -complete problem [6]. Therefore, by definition, any problem in NP can be reduced to 3CNF in polynomial time. We show in Theorem 8 that the *One-shared-event DFA modular nonblockingness* (1SE-DFA-MN) problem is in NP , hence reducible to 3CNF in polynomial time.

The membership in NP gives an upper bound on the complexity of 1SE-DFA-MN, which can still be polynomially or even trivially solvable. To rule out this possibility, we further show that 1SE-DFA-MN is NP -hard (and hence NP -complete) by reducing 3CNF to 1SE-DFA-MN. Then, consequently, any problem in NP can be reduced to the 1SE-DFA-MN problem in polynomial time. Hence, from the complexity point of view, both problems are equally difficult.

¹A (boolean) formula consists of variables, operators conjunction, disjunction and negation, and parentheses. A formula is satisfiable if there is an assigning of 1 (*true*) and 0 (*false*) to its variables making it *true*. A literal is a variable or its negation. A clause is a disjunction of literals. A formula is in 3-cnf if it is a conjunction of clauses, each clause with three literals. For instance, $\varphi = (x \vee y \vee z) \wedge (\neg x \vee y \vee z)$ is a formula in 3-cnf with two clauses $x \vee y \vee z$ and $\neg x \vee y \vee z$. Given a formula in 3-cnf, the 3CNF problem asks whether the formula is satisfiable. The formula φ is satisfiable for, e.g., $(x, y, z) = (0, 1, 0)$.

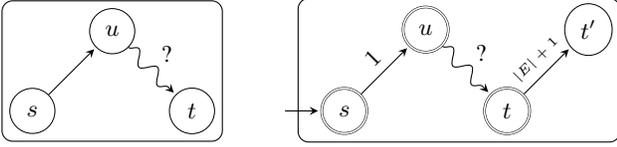


Fig. 2. Graph $G = (V, E, s, t)$ and its corresponding DFA \mathcal{A} of Theorem 2

III. COMPLEXITY OF NONBLOCKINGNESS

Let $\mathcal{A} = (Q, \Sigma, \delta, I, F)$ be an NFA. We define the size of \mathcal{A} as $|\mathcal{A}| = |Q| + |\Sigma| + |\delta| + |I| + |F|$.

A DFA is nonblocking iff from every state a marked state is reachable (in other words, every state is reachable and co-reachable). This property can be tested in linear time using the computation of strongly connected components [7]. From the complexity point of view, under the assumption that $NL \neq P$, a stronger result can be shown.

Theorem 2 (DFA-nonblockingness): Given a DFA \mathcal{A} , the problem whether $\overline{L_m(\mathcal{A})} = L(\mathcal{A})$ is NL-complete.

Proof: The membership of DFA-nonblockingness in NL follows from Algorithm 3 below for $n = 1$.

We now show that DFA-nonblockingness is NL-hard by reducing *graph non-reachability* [4] to DFA-nonblockingness. Namely, let $G = (V, E, s, t)$ be a directed graph with s, t in V . We construct a DFA \mathcal{A} from G in logarithmic space such that t is not reachable from s in G iff \mathcal{A} is nonblocking.

Let $\mathcal{A} = (V \cup \{t'\}, \Sigma, \delta, s, V)$, where δ is defined as the relation E with every transition under a unique label, and a transition under a new label is added from t to the new non-marked state t' . This reduction (transformation) of G to \mathcal{A} can be done in logarithmic space and is performed by Algorithm 2, where $\Sigma = \{1, 2, \dots, |E| + 1\}$. If the algorithm reads a node v in V , it outputs state v . Then it prints state t' . After this part, it has printed the state set of \mathcal{A} . It only needs to store a pointer (of logarithmic size) to the position of the input currently read. Then the algorithm counts from 1 to $|E| + 1$ and outputs the numbers, that is, it prints the alphabet of \mathcal{A} . For this, two numbers, i and $|E| + 1$ with at most $\lceil \log(|E| + 1) \rceil$ digits are stored. Then, it reads the input again (using the pointer as above) and uses a counter c (with at most $\lceil \log(|E| + 1) \rceil$ digits) to print, for every edge (u, v) in E , the corresponding transition (u, c, v) of δ . Finally, it prints the transition $(t, |E| + 1, t')$, state s , and all $v \in V$. After this, the output contains the DFA \mathcal{A} . The reduction uses logarithmic space to produce the output. Recall that the size of the input and output is not considered in space complexity.

Algorithm 2: (Reduction of a graph to a DFA)

Input : A directed graph $G = (V, E, s, t)$
Output: The DFA $\mathcal{A} = (V \cup \{t'\}, \Sigma, \delta, s, V)$
1 **for** each $v \in V$ **do** output v
2 output t'
3 **for** $i = 1, \dots, |E| + 1$ **do** output i
4 $c := 1$; **for** each $(u, v) \in E$ **do** {output (u, c, v) ; $c++$ }
5 output (t, c, t') ; output s
6 **for** each $v \in V$ **do** output v

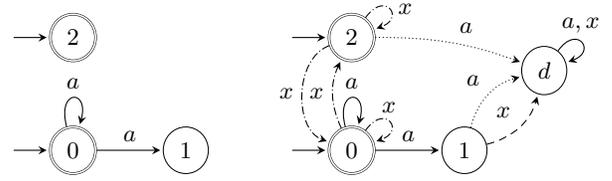


Fig. 3. A nonblocking NFA \mathcal{B} with $L_m(\mathcal{B}) = L(\mathcal{B}) = \{a\}^*$ and its corresponding NFA \mathcal{A} constructed in the proof of Theorem 3

It is not difficult to see that t is reachable from s in G iff t is accessible from the initial state s in \mathcal{A} . Namely, if t is not accessible in \mathcal{A} , then all accessible states are marked and the language of \mathcal{A} is nonblocking. If t is accessible in \mathcal{A} , then so is t' , which is not marked and makes thus the language of \mathcal{A} blocking; cf. Fig. 2 for an illustration. ■

Therefore, to check nonblockingness of a DFA is easy. This is, however, not true for NFAs. An NFA can be nonblocking even if there is a state from which no marked state is reachable, cf. Fig. 3.

Theorem 3 (NFA-nonblockingness): Given an NFA \mathcal{A} , the problem whether $\overline{L_m(\mathcal{A})} = L(\mathcal{A})$ is PSPACE-complete.

Proof: To show that the problem is in PSPACE, let $\mathcal{A} = (Q, \Sigma, \delta, I, F)$ be an NFA. Let \mathcal{D} be a DFA obtained from \mathcal{A} by the standard subset construction [5]. States of \mathcal{D} are subsets of states of \mathcal{A} , and \mathcal{A} is nonblocking iff \mathcal{D} is nonblocking. To check nonblockingness of \mathcal{D} in polynomial space, \mathcal{D} cannot be computed and stored, because it may require exponential space in the size of \mathcal{A} . Instead, we use the on-the-fly technique that keeps only a small part of \mathcal{D} in memory and re-computes the required parts on request. Namely, for every state $X \subseteq Q$ of \mathcal{D} , we check that X is reachable from the initial state of \mathcal{D} (in the way depicted in Example 1). If so, we guess a marked state Y of \mathcal{D} , that is, $Y \cap F \neq \emptyset$, and check that Y is reachable from X . This principle is generalized in Algorithm 3 below. At any time during the computation, the algorithm stores only a constant number of states of \mathcal{D} , which are subsets of the state set Q of \mathcal{A} . Therefore, the algorithm uses space polynomial in the size of \mathcal{A} and the problem is thus in PSPACE.

To show that NFA-nonblockingness is PSPACE-hard, we reduce the NFA *universality problem* [6] to it. NFA universality asks, given an NFA \mathcal{B} over Σ , whether $L_m(\mathcal{B}) = \Sigma^*$.

Let $\mathcal{B} = (Q, \Sigma, \delta_B, I, F)$ be an NFA, and let d be a new non-marked state. We “complete” \mathcal{B} in the sense that if, for an event a in Σ , no a -transition is defined in a state q , we add the transition (q, a, d) to the transition relation, see the dotted transitions in Fig. 3. Let $x \notin \Sigma$ be a new event. State d contains self-loops under all events of $\Sigma \cup \{x\}$. For each non-marked state p , we add the transition (p, x, d) , see the dashed transitions in Fig. 3, while for each marked state p , we add the transitions (p, x, i) for every initial state i in I , cf. the dot-dash transitions in Fig. 3. Let $\mathcal{A} = (Q \cup \{d\}, \Sigma \cup \{x\}, \delta_A, I, F)$ denote the resulting NFA. Notice that $L_m(\mathcal{B}) \subseteq L_m(\mathcal{A})$. We now show that \mathcal{B} is universal iff \mathcal{A} is nonblocking.

If \mathcal{B} is universal, that is, $L_m(\mathcal{B}) = \Sigma^*$, we show that \mathcal{A} is nonblocking by showing that $L_m(\mathcal{A}) = L(\mathcal{A})$. It is sufficient to show that $L(\mathcal{A}) \subseteq L_m(\mathcal{A})$. Let $w \in L(\mathcal{A})$. We proceed by induction on the number of occurrences of event x in w . If x

does not occur in w , then $w \in \Sigma^* = L_m(\mathcal{B}) \subseteq L_m(\mathcal{A})$. Thus, assume that $w = w_1 x w_2$ with $w_1 \in \Sigma^*$ and $w_2 \in (\Sigma \cup \{x\})^*$. Since $w_1 \in \Sigma^* = L_m(\mathcal{B}) \subseteq L_m(\mathcal{A})$, there is a path in \mathcal{A} from an initial state i in I to a marked state f in F labeled by w_1 . By construction, $I \subseteq \delta_{\mathcal{A}}(i, w_1 x)$, since $\delta_{\mathcal{A}}(f, x) = I$. By the induction hypothesis, $\delta_{\mathcal{A}}(I, w_2) \cap F \neq \emptyset$, hence $w \in L_m(\mathcal{A})$. Thus, $L_m(\mathcal{A}) = L(\mathcal{A})$, which means that \mathcal{A} is nonblocking.

If \mathcal{B} is not universal, that is, $L_m(\mathcal{B}) \neq \Sigma^*$, then there exists a w in Σ^* such that $\delta_{\mathcal{A}}(I, w) \cap F = \emptyset$, since for any w over Σ , $\delta_{\mathcal{A}}(I, w) \cap F = \delta_{\mathcal{B}}(I, w) \cap F$. By the construction, $\delta_{\mathcal{A}}(I, wx) = \{d\}$, from which no marked state is reachable. Since \mathcal{A} is complete, wx belongs to $L(\mathcal{A})$, therefore \mathcal{A} is blocking. ■

The situation with NFAs is even worse as shown now.

Theorem 4 (NFA-prefix-closed): Given an NFA \mathcal{A} , the problem whether $L_m(\mathcal{A})$ is prefix-closed is PSPACE-complete.

Proof: Let \mathcal{A} be an NFA. Then $L_m(\mathcal{A}) = \overline{L_m(\mathcal{A})}$ iff the DFA \mathcal{D} obtained from \mathcal{A} by the standard subset construction has no reachable and co-reachable non-marked states. Since the class PSPACE is closed under complement, we can check the opposite – a nondeterministic algorithm guesses a subset of non-marked states of \mathcal{A} and verifies, using the on-the-fly technique, that they form a reachable and co-reachable state in \mathcal{D} . The NFA-prefix-closed problem is thus in PSPACE.

To show PSPACE-hardness, Hunt III and Rosenkrantz [8] have shown that a property R of languages over $\{0, 1\}$ such that $R(\{0, 1\}^*)$ is true and there exists a regular language that is not expressible as a quotient $x \setminus L = \{w \mid xw \in L\}$, for some L for which $R(L)$ is true, is as hard as to decide “ $= \{0, 1\}^*$ ”. Since prefix-closedness is such a property (the class of prefix-closed languages is closed under quotient) and universality is PSPACE-hard for NFAs, the result implies that the NFA-prefix-closed problem is PSPACE-hard. ■

These results justify why the attention is mostly focused on DFAs rather than NFAs. In the rest of the paper, we also focus on DFAs, unless stated otherwise.

A. Modular Nonblockingness Problem

We now focus on the modular nonblockingness problem. The simplest case is that there is no interaction between the different subsystems. The following result is well known.

Theorem 5: Let J be a finite set, and let \mathcal{A}_j be a nonblocking NFA over Σ_j , for $j \in J$. If the alphabets are pairwise disjoint, that is, $\Sigma_j \cap \Sigma_{j'} = \emptyset$ for $j \neq j'$, then the parallel composition $\overline{L_m(\|_{j \in J} \mathcal{A}_j)} = L(\|_{j \in J} \mathcal{A}_j)$ is nonblocking. ■

In many complex systems, it is however the case that there are events shared between the subsystems. In such a case, nonblockingness is in general PSPACE-complete [9]. A more fine-grained complexity can be distinguished based on the following criteria. Let $(\mathcal{A}_i)_{i=1}^n$ be DFAs:

- 1) The number of DFAs is not restricted.
- 2) The number of DFAs is restricted by a function $g(m)$, that is, $n \leq g(m)$, where m is the length of the encoding of the DFAs $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$.
- 3) The number of DFAs is restricted by a constant k , that is, $n \leq k$.

Case 2 is the most general one and deserves a discussion. Assume, for example, that our encoding of \mathcal{A}_i requires $c > 1$ bits and that the encoding of $\mathcal{A}_1, \dots, \mathcal{A}_n$ requires $m = n \cdot c$ bits. If $g(m) = m$, then $n \leq g(nc)$ for every $n \geq 1$, which results in the non-restricted case 1. If $g(m) = k$, for a constant k , then $n \leq g(nc)$ iff $n \leq k$, which results in the restriction of case 3. See also Remark 7 below.

We can now prove the following result.

Theorem 6 (DFA $g(m)$ -bounded modular nonblockingness): Given nonblocking DFAs $(\mathcal{A}_i)_{i=1}^n$ with \mathcal{A}_i over Σ_i , $2 \leq n \leq g(m)$, where m is the length of an encoding of the sequence of DFAs $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$. The problem whether $\overline{L_m(\|_{i=1}^n \mathcal{A}_i)} = L(\|_{i=1}^n \mathcal{A}_i)$ is NSPACE($g(m) \log m$)-complete.

Algorithm 3: Is $\mathcal{A} = \|_{i=1}^n \mathcal{A}_i$ nonblocking?

Input : Encoding of $\mathcal{A}_1, \dots, \mathcal{A}_n$ of size m
Output: yes iff $\|_{i=1}^n \mathcal{A}_i$ is nonblocking

- 1 **for each** $(p_1, \dots, p_n) \in \times_{i=1}^n Q_i$ **do**
- 2 **if** (p_1, \dots, p_n) is reachable from the init st. of \mathcal{A} **then**
- 3 choose $(s_1, \dots, s_n) \in \times_{i=1}^n F_i$
- 4 $k_i := p_i$, for $i = 1, 2, \dots, n$
- 5 **repeat**
- 6 choose $a \in \Sigma$
- 7 $k_i := \delta(k_i, a)$, for $i = 1, 2, \dots, n$
- 8 **until** $k_i = s_i$, for $i = 1, 2, \dots, n$
- 9 **return** yes

Proof: Let $(\mathcal{A}_i)_{i=1}^n$ over $(\Sigma_i)_{i=1}^n$ be nonblocking DFAs. Algorithm 3 solves the $g(m)$ -bounded modular nonblockingness problem. It works as follows: for every reachable state (p_1, \dots, p_n) of \mathcal{A} (lines 1-2), the algorithm nondeterministically chooses a marked state (s_1, \dots, s_n) of \mathcal{A} (line 3) that is reachable from state (p_1, \dots, p_n) (lines 4-8). The algorithm returns yes iff there is such a marked state for every reachable state, hence iff \mathcal{A} is nonblocking. (Compared to Example 1, we omitted the counter *numSteps* for simplicity. It should be clear how the counter is introduced to make the algorithm always terminate.)

During the computation, the algorithm stores only a constant number of n -tuples of pointers (t_1, \dots, t_n) . The space used is therefore $O(n \log m)$. Since n is bounded by $g(m)$, the space used by Algorithm 3 is $O(g(m) \log m)$, hence the problem is in NSPACE($g(m) \log m$).

To prove hardness, we reduce the NSPACE($g(m) \log m$)-complete *finite DFA intersection problem (DFA-int)* [10] to our problem. DFA-int asks, given DFAs $(\mathcal{B}_i)_{i=1}^n$ with $2 \leq n \leq g(m)$, where m is the length of the encoding of the sequence of $\mathcal{B}_1, \dots, \mathcal{B}_n$, whether $\cap_{i=1}^n L_m(\mathcal{B}_i) = \emptyset$. The DFAs $\mathcal{B}_1, \dots, \mathcal{B}_n$ are over a common alphabet Σ .

We now describe a deterministic logarithmic-space reduction from DFA-int to DFA $g(m)$ -bounded modular nonblockingness. Notice that $n \geq 2$. Let $x \notin \Sigma$ be a new event.

We construct \mathcal{A}_1 from \mathcal{B}_1 by adding two new states d_1 and d'_1 and x -transitions from every marked state of \mathcal{B}_1 to d_1 , and from d_1 to d'_1 , see an illustration in Fig. 4. All states of \mathcal{A}_1 , but d_1 , are marked, that is, $L_m(\mathcal{A}_1) = L(\mathcal{B}_1) \cup L_m(\mathcal{B}_1)xx$ and $L(\mathcal{A}_1) = L(\mathcal{B}_1) \cup L_m(\mathcal{B}_1)\{x, xx\}$.

For every $i \geq 2$, we construct \mathcal{A}_i from \mathcal{B}_i by adding a new state d_i and x -transitions from every marked state of \mathcal{B}_i to d_i . All states of \mathcal{A}_i are marked, hence $L_m(\mathcal{A}_i) = L(\mathcal{A}_i) = L(\mathcal{B}_i) \cup L_m(\mathcal{B}_i)x$.

We thus have that $L_m(\|_{i=1}^n \mathcal{A}_i) = \cap_{i=1}^n L(\mathcal{B}_i)$ and

$$L(\|_{i=1}^n \mathcal{A}_i) = \cap_{i=1}^n L(\mathcal{B}_i) \cup \cap_{i=1}^n L_m(\mathcal{B}_i)x.$$

We show $\overline{L_m(\|_{i=1}^n \mathcal{A}_i)} = L(\|_{i=1}^n \mathcal{A}_i)$ iff $\cap_{i=1}^n L_m(\mathcal{B}_i) = \emptyset$.

If $\cap_{i=1}^n L_m(\mathcal{B}_i) = \emptyset$, then $L(\|_{i=1}^n \mathcal{A}_i) = \cap_{i=1}^n L(\mathcal{B}_i) = L_m(\|_{i=1}^n \mathcal{A}_i)$.

If $\cap_{i=1}^n L_m(\mathcal{B}_i) \neq \emptyset$, then there exists $w \in \cap_{i=1}^n L_m(\mathcal{B}_i)$, hence $wx \in L(\|_{i=1}^n \mathcal{A}_i) \setminus \overline{L_m(\|_{i=1}^n \mathcal{A}_i)}$, because $L_m(\|_{i=1}^n \mathcal{A}_i)$ does not contain any string with event x . ■

Remark 7: Let k be a constant. If for every m , $g(m) \leq k$, then Algorithm 3 uses space $O(\log m)$, hence DFA k -bounded modular nonblockingness is in $\text{NL} = \text{NSPACE}(\log m)$ and it is NL -complete. If $g(m) \leq \log^k m$, then DFA $(\log^k m)$ -bounded modular nonblockingness is $\text{NSPACE}(\log^{k+1} m)$ -complete. It is called a poly-logarithmic space complexity.

B. One-Shared-Event Modular Nonblockingness

We now focus on the case where exactly one event is shared. An application of this case is, for example, in the Brandin and Wonham modular framework for timed discrete event systems [11], where only one event simulating the tick of a global clock is shared and all the other events are local [12]. Unless $\text{NP} = \text{PSPACE}$, nonblockingness is computationally easier in this case.

Let \mathcal{A} be an NFA over Σ and P be a projection from Σ^* . Then $P(\mathcal{A})$ is a DFA such that $L_m(P(\mathcal{A})) = P(L_m(\mathcal{A}))$ and $L(P(\mathcal{A})) = P(L(\mathcal{A}))$, called an *observer*; cf. [3], [13] for a construction. In the worst case, $P(\mathcal{A})$ has exponentially many states compared to \mathcal{A} [14], [15].

Theorem 8 (One-shared-event DFA modular nonblockingness): Given $n \geq 2$ nonblocking DFAs $(\mathcal{A}_i)_{i=1}^n$ with \mathcal{A}_i over Σ_i such that $|\cup_{i \neq j} (\Sigma_i \cap \Sigma_j)| = 1$. The problem to decide whether $L_m(\|_{i=1}^n \mathcal{A}_i) = L(\|_{i=1}^n \mathcal{A}_i)$ is NP -complete.

Proof: Let $(\mathcal{A}_i)_{i=1}^n$ over $(\Sigma_i)_{i=1}^n$ be nonblocking DFAs, and let a be the only event such that $\Sigma_i \cap \Sigma_j = \{a\}$, for every $i \neq j$. Let P be a projection from $(\cup_{i=1}^n \Sigma_i)^*$ to $\{a\}^*$. Let m be the maximum number of states of all \mathcal{A}_i 's. Then $P(\mathcal{A}_i)$ is a DFA with at most 2^m states, hence the composition $\mathcal{A} = \|_{i=1}^n P(\mathcal{A}_i)$ has at most 2^{mn} states, each of the form (X_1, \dots, X_n) , where X_i is a subset of states of \mathcal{A}_i .

Let δ denote the transition function of \mathcal{A} and q_0 its initial state. Notice that \mathcal{A} is a sequence of transitions possibly with a cycle at the end. Then \mathcal{A} is nonblocking iff there exist $k \leq 2^{mn}$ and $k < \ell \leq 2^{mn+1}$ such that $\delta(q_0, a^k)$ is an accepting state of \mathcal{A} and either $\delta(q_0, a^{k+1})$ is not defined or $\delta(q_0, a^k) = \delta(q_0, a^\ell)$. We now show how to check this property in nondeterministic polynomial time.

The nondeterministic algorithm guesses k and ℓ in binary, requiring at most $mn + 1$ digits each. To compute the states $\delta(q_0, a^k)$ and $\delta(q_0, a^\ell)$ in polynomial time, the algorithm proceeds as follows.

Let \mathcal{A}'_i denote the NFA obtained from \mathcal{A}_i by replacing each transition (s, b, t) with the transition $(s, P(b), t)$, and by

eliminating the ε -transitions afterwards. This can be computed in polynomial time [16] and is often used as the middle step in the computation of the observer; namely, it preserves the languages. Then \mathcal{A}'_i is over $\{a\}$ and has the same states as \mathcal{A}_i . Let J_i denote the set of all initial states of \mathcal{A}'_i ; it is computed in polynomial time as the set of all states of \mathcal{A}_i reachable under $\Sigma \setminus \{a\}$ from the initial state of \mathcal{A}_i (it is also the initial state of $P(\mathcal{A}_i)$, that is, $q_0 = (J_1, \dots, J_n)$).

The transition relation of \mathcal{A}'_i can be represented as a binary matrix M_i , where for states s, t of \mathcal{A}'_i , $M_i[s, t] = 1$ iff (s, a, t) is a transition in \mathcal{A}'_i . For $k \geq 2$, let M_i^k be the multiplication of M_i with itself k times. Then $M_i^k[s, t]$ is the number of paths of length k from s to t in \mathcal{A}'_i [17]. Let $\delta_{\mathcal{A}'_i}$ denote the transition function of \mathcal{A}'_i . Then $\delta_{\mathcal{A}'_i}(q_i, a^k) = \{t \mid M_i^k[q_i, t] > 0\}$ (if it is empty, the transition is undefined). The size of matrix M_i^k is polynomial in the number of states of \mathcal{A}_i and can be computed in time logarithmic in k by fast matrix multiplication: $M_i^2 = M_i \times M_i$, $M_i^4 = M_i^2 \times M_i^2$, $M_i^8 = M_i^4 \times M_i^4$, \dots .

To compute $\delta(q_0, a^k)$, we compute M_i^k , for $i = 1, \dots, n$, in polynomial time. Then the state $\delta(q_0, a^k) = (\delta_{\mathcal{A}'_1}(J_1, a^k), \dots, \delta_{\mathcal{A}'_n}(J_n, a^k))$ and it is marked iff every $\delta_{\mathcal{A}'_i}(J_i, a^k)$ contains a marked state of \mathcal{A}'_i . It should now be clear how to check, in polynomial time, that either $\delta(q_0, a^{k+1})$ is not defined or $\delta(q_0, a^k) = \delta(q_0, a^\ell)$, cf. also Example 10 and Remark 11.

To show NP -hardness, we reduce 3CNF to our problem and use the construction of [18]. Let φ be a formula in 3CNF (see footnote 1 on page 2) with n distinct variables and m clauses, and let C_k be the set of literals in the k th clause, $1 \leq k \leq m$. The assignment to the variables is represented as a binary vector of length n . Let p_1, \dots, p_n denote the first n prime numbers. For a natural number z congruent with 0 or 1 modulo p_i , for all $i = 1, \dots, n$, z satisfies φ if the assignment $(z \bmod p_1, \dots, z \bmod p_n)$ satisfies φ .

For $u = 1, \dots, n$ and $j = 2, \dots, p_u - 1$, let $\mathcal{B}'_{u,j}$ denote a nonblocking DFA such that

$$L_m(\mathcal{B}'_{u,j}) = 0^j \cdot (0^{p_u})^*.$$

Then $\cup_{u=1}^n \cup_{j=2}^{p_u-1} L_m(\mathcal{B}'_{u,j}) = \{0^z \mid \exists u \leq n, z \not\equiv 0 \pmod{p_u} \text{ and } z \not\equiv 1 \pmod{p_u}\}$ is the set of all natural numbers that *do not encode an assignment* to the variables.

For each clause C_k , we construct a nonblocking DFA \mathcal{B}'_k such that if $0^z \in L_m(\mathcal{B}'_k)$ and z is an assignment, then z does not assign value 1 to any literal in C_k . For example, if $C_k = \{x_r, \neg x_s, x_t\}$, for $1 \leq r, s, t \leq n$ and r, s, t distinct, let z_k be the unique integer such that $0 \leq z_k < p_r p_s p_t$, $z_k \equiv 0 \pmod{p_r}$, $z_k \equiv 1 \pmod{p_s}$, and $z_k \equiv 0 \pmod{p_t}$. Then

$$L_m(\mathcal{B}'_k) = 0^{z_k} \cdot (0^{p_r p_s p_t})^*.$$

Let $\mathcal{B}_1, \dots, \mathcal{B}_\ell$ denote all the DFAs $\mathcal{B}'_{u,j}$ and \mathcal{B}'_k constructed above, and let \mathcal{A}_i denote \mathcal{B}_i with the sets of marked and non-marked states exchanged, that is, $L_m(\mathcal{A}_i) = 0^* \setminus L_m(\mathcal{B}_i)$. Note that all \mathcal{B}_i and \mathcal{A}_i are nonblocking and their generated languages are 0^* .

Now, φ is satisfiable if and only if there exists z such that z encodes an assignment to φ , i.e., $0^z \notin \cup_{u=1}^n \cup_{j=2}^{p_u-1} L_m(\mathcal{B}'_{u,j})$, and z satisfies every clause C_k , that is, $0^z \notin L_m(\mathcal{B}'_k)$ for all $k = 1, \dots, m$. This is iff $0^z \in \cap_{i=1}^\ell L_m(\mathcal{A}_i) = L_m(\|_{i=1}^\ell \mathcal{A}_i)$.

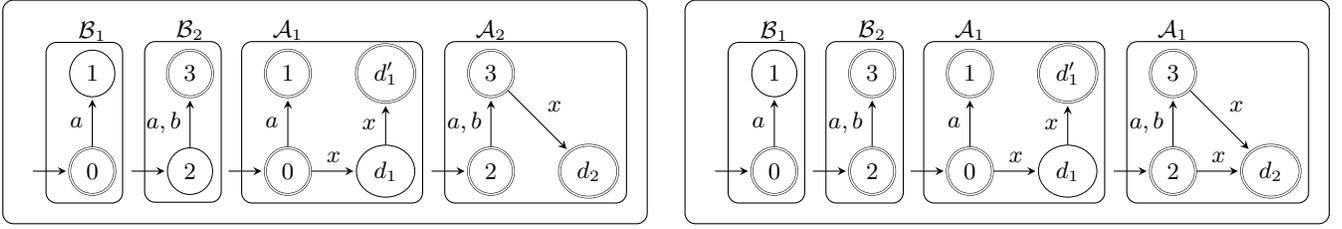


Fig. 4. DFAs \mathcal{B}_1 and \mathcal{B}_2 over $\Sigma = \{a, b\}$ and the corresponding DFAs \mathcal{A}_1 and \mathcal{A}_2 of Theorem 6; a nonblocking instance (left) with $\overline{L_m(\mathcal{A}_1 \parallel \mathcal{A}_2)} = L(\mathcal{A}_1 \parallel \mathcal{A}_2) = \{\varepsilon, a\}$ and a blocking instance (right) with $\overline{L_m(\mathcal{A}_1 \parallel \mathcal{A}_2)} = \{\varepsilon, a\}$ and $L(\mathcal{A}_1 \parallel \mathcal{A}_2) = \{\varepsilon, a, x\}$

We show that $\|_{i=1}^{\ell} \mathcal{A}_i$ is nonblocking iff $L_m(\|_{i=1}^{\ell} \mathcal{A}_i) \neq \emptyset$.

If $L_m(\|_{i=1}^{\ell} \mathcal{A}_i) = \emptyset$, then $\|_{i=1}^{\ell} \mathcal{A}_i$ is blocking, because $\varepsilon \in L(\|_{i=1}^{\ell} \mathcal{A}_i)$.

If $0^z \in L_m(\|_{i=1}^{\ell} \mathcal{A}_i)$, then z satisfies φ . For a natural number c , the number $z + c \cdot \prod_{i=1}^n p_i$ also satisfies φ : indeed, if $z \equiv x_i \pmod{p_i}$, then $(z + c \cdot \prod_{i=1}^n p_i) \equiv x_i \pmod{p_i}$, for all i . Thus, $0^z (0^{\prod_{i=1}^n p_i})^* \subseteq L_m(\|_{i=1}^{\ell} \mathcal{A}_i)$. Since for every $0^s \in L(\|_{i=1}^{\ell} \mathcal{A}_i)$, there exists c such that $s \leq z + c \cdot \prod_{i=1}^n p_i$, we have that $\|_{i=1}^{\ell} \mathcal{A}_i$ is nonblocking. ■

Remark 9: If the number of DFAs in Theorem 8 is at most k , for a constant k , the problem is NL-complete. The membership in NL is by Theorem 6 and NL-hardness by Theorem 2.

Example 10: We illustrate the polynomial computation used in the proof of Theorem 8 for $n = 1$. Its generalization to $n > 1$ is straightforward. Let $\mathcal{A}_1 = (\{1, 2, 3, 4\}, \{a, b\}, \{(1, a, 2), (2, a, 1), (2, b, 3), (3, a, 4), (4, a, 1)\}, 1, \{1\})$ be a DFA, and let \mathcal{A}'_1 , depicted in Fig. 5, be the NFA obtained from \mathcal{A}_1 by

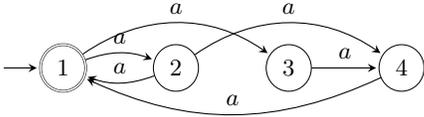


Fig. 5. The NFA \mathcal{A}'_1 of Example 10

renaming b -transitions to ε -transitions, and by the elimination of ε -transitions afterwards. Then the 4×4 transition matrix M_1 and its 4th power M_1^4 are

$$M_1 = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad M_1^4 = \begin{pmatrix} 1 & 2 & 2 & 2 \\ 3 & 1 & 1 & 2 \\ 1 & 0 & 0 & 2 \\ 2 & 1 & 1 & 0 \end{pmatrix}.$$

The reader may verify that $\delta(\{1\}, a^4) = \{1, 2, 3, 4\}$, where δ is the transition relation of the observer $P(\mathcal{A}_1)$.

Remark 11: The number $M_i^k[s, t]$ represents the number of paths from state s to state t of length k . This information is not important for us. The information we need is whether there is a path, i.e., $M_i^k[s, t] > 0$, or not, i.e., $M_i^k[s, t] = 0$. The numbers $M_i^k[s, t]$ may become large and affect thus the complexity. To keep the complexity polynomial (the numbers small), the $+$ operation in the definition of matrix multiplication is replaced by max operation. This minor trick keeps the matrices M_i^k binary, while providing the same information [17].

IV. CONCLUSION

The theoretical results do not seem very optimistic. However, there are techniques to reduce the size of an automaton,

which allows to handle large automata that appear in practical applications. A well-known technique is the BDD diagrams [19]. Another technique is the state-tree structures [20] or the method using extended finite-state machines and abstractions [21].

ACKNOWLEDGMENT

The author gratefully acknowledges very useful suggestions and comments of the anonymous referees.

REFERENCES

- [1] S. Mohajerani, R. Malik, and M. Fabian, "A framework for compositional synthesis of modular nonblocking supervisors," *IEEE Trans. Autom. Control*, vol. 59, pp. 150–162, 2014.
- [2] R. Malik, "Programming a fast explicit conflict checker," in *WODES*, 2016, pp. 438–443.
- [3] C. Cassandras and S. Lafortune, *Introduction to discrete event systems*, 2nd ed. Springer, 2008.
- [4] C. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
- [5] M. Sipser, *Introduction to the theory of computation*, 2nd ed. Thompson Course Technology, 2006.
- [6] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [7] R. Tarjan, "Depth-first search and linear graph algorithms," *SIAM J. Comput.*, vol. 1, pp. 146–160, 1972.
- [8] H. B. Hunt III and D. J. Rosenkrantz, "Computational parallels between the regular and context-free languages," *SIAM J. Comput.*, vol. 7, pp. 99–114, 1978.
- [9] K. Rohloff and S. Lafortune, "PSPACE-completeness of modular supervisory control problems," *Discrete Event Dyn. Syst.*, vol. 15, pp. 145–167, 2005.
- [10] M. Holzer and M. Kutrib, "Descriptive and computational complexity of finite automata – A survey," *Inform. and Comput.*, vol. 209, pp. 456–470, 2011.
- [11] B. Brandin and W. Wonham, "Supervisory control of timed discrete-event systems," *IEEE Trans. Autom. Control*, vol. 39, pp. 329–342, 1994.
- [12] G. Schafaschek, M. Queiroz, and J. Cury, "Local modular supervisory control of timed discrete-event systems," in *WODES*, 2014, pp. 271–277.
- [13] W. Wonham, "Supervisory control of discrete-event systems," 2009.
- [14] K. Wong, "On the complexity of projections of discrete-event systems," in *WODES*, 1998, pp. 201–206.
- [15] G. Jirásková and T. Masopust, "On a structural property in the state complexity of projected regular languages," *Theoret. Comput. Sci.*, vol. 449, pp. 93–105, 2012.
- [16] J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
- [17] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. MIT Press, 2009.
- [18] L. Stockmeyer and A. Meyer, "Word problems requiring exponential time: Preliminary report," in *STOC*, 1973, pp. 1–9.
- [19] R. Bryant, "Graph-based algorithms for boolean function manipulation," *IEEE Trans. Comput.*, vol. 35, pp. 677–691, 1986.
- [20] C. Ma and W. Wonham, "Nonblocking supervisory control of state tree structures," *IEEE Trans. Autom. Control*, vol. 51, pp. 782–793, 2006.
- [21] S. Mohajerani, R. Malik, and M. Fabian, "A framework for compositional nonblocking verification of extended finite-state machines," *Discrete Event Dyn. Syst.*, vol. 26, pp. 33–84, 2016.