




From Natural Projection to Partial Model Checking and Back

Gabriele Costa¹(✉) , David Basin², Chiara Bodei³, Pierpaolo Degano³,
and Letterio Galletta³

¹ DIBRIS, Università degli Studi di Genova, Genoa, Italy
gabriele.costa@unige.it

² Department of Computer Science, ETH Zurich,
Zürich, Switzerland
basin@inf.ethz.ch

³ Dipartimento di Informatica, Università di Pisa,
Pisa, Italy
{chiara,degano,galletta}@di.unipi.it



Abstract. Specification decomposition is a theoretically interesting and practically relevant problem for which two approaches were independently developed by the control theory and verification communities: *natural projection* and *partial model checking*. In this paper we show that, under reasonable assumptions, natural projection reduces to partial model checking and, when cast in a common setting, the two are equivalent. Aside from their theoretical interest, our results build a bridge whereby the control theory community can reuse algorithms and results developed by the verification community. In addition, we present an algorithm and a tool for the partial model checking of finite-state automata that can be used as an alternative to natural projection.

1 Introduction

System verification requires comparing a system's behavior against a specification. When the system consists of several components, we can distinguish between *local* and *global* specifications. A local specification applies to a single component, whereas a global specification should hold for the entire system. Since these specifications are needed to reason at different levels of abstraction, both of them are often present.

Ideally, we would like to freely pass from local to global specifications and vice versa. Most specification formalisms natively support specification composition. Logical conjunction, set intersection, and the synchronous product of automata are all examples of operators for composing specifications. Unfortunately, the same does not hold for specification decomposition: obtaining local specifications from a global one is, in general, extremely complex, as illustrated below.

G. Costa and L. Galletta—The current address of the authors is IMT Lucca.

© The Author(s) 2018

D. Beyer and M. Huisman (Eds.): TACAS 2018, LNCS 10805, pp. 344–361, 2018.

https://doi.org/10.1007/978-3-319-89960-2_19

Example 1. Consider the classical submodule construction problem (SCP) [27]: given a system and a global specification, find a submodule whose composition with the system leads to a global behavior conformant to the global specification. For instance, imagine that we aim to control the usage of a buffer of size n , where two agents, A and B , can insert and remove items. Now assume that A 's behavior is to “insert one item when the buffer is empty and delete one item when it is full”, while B 's behavior is unknown. If we want to prevent buffer overflow and underflow, some questions arise about B . For example, are there compatible behaviors for B ? Is there a *most general* one? How could we effectively compute it? These questions require decomposing the buffer overflow/underflow specification so that it only refers to B , while exploiting the known structure of A . \square

Over the past decades, researchers have investigated methods for decomposing specifications. Interestingly, different communities have tackled this problem independently, each considering specification formalisms and assumptions appropriate for their application context. In particular, important results were obtained in two distinct fields: *control theory* and *formal verification*.

In control theory, *natural projection* [33] is used for simplifying systems built from multiple components, modeled as automata. It is often applied component-wise to synthesize local controllers from a global specification of asynchronous *discrete-event system* [9], namely the *controller synthesis problem* (CSP). Briefly, local controllers guarantee that the global specification is not violated by interacting only with a single component of a system. The local controllers can be used to implement distributed control systems [34,35] by composing them in parallel with other sub-systems.

In the formal verification community, *partial model checking* [1] was proposed as a technique for mitigating the state explosion problem when verifying large systems built from many parallel processes. Partial model checking tackles this problem by decomposing a specification, given as a formula of the μ -calculus [22], using a *quotienting* operator, thereby supporting the analysis of the individual processes independently. Quotienting carries out a partial evaluation of a specification while *preserving the model checking problem*. Thus, a system built by composing two modules satisfies a specification if and only if one of the modules

Table 1. Summary of existing results on natural projection and partial model checking. Notice that the algorithm in [32] runs in PTIME on a specific class of discrete-event systems.

	NATURAL PROJECTION	PARTIAL MC
Spec. Lang.	FSA [19,31]	μ -calculus [1,3]
Theory	FSA [19,31]	LTS [1,3]
Complexity	EXPTIME [14,32]	EXPTIME [1,3]
Tools	TCT [13], IDES3 [30], DESTool [28]	mCRL2 [18], CADP [23], MuDiv [2]

satisfies the specification after quotienting against the other [1]. This may reduce the problem size, resulting in smaller models and hence faster verification.

Table 1 summarizes some relevant facts about the two approaches and we refer the reader to Sect. 5 for a more detailed analysis. Since natural projection and partial model checking apply to different formalisms, they cannot be directly compared without defining a common framework (see below). For instance, a relevant question is comparing how specifications grow under the two approaches. Although it is known that both may lead to exponential growth (see [3, 21, 32]), these results apply in one case to finite-state automata (FSAs) and in the other case to μ -calculus formulae.

Over the past few years, there has been a substantial cross-fertilization between the two research communities [12]. For instance, methods for synthesizing controllers using partial model checking are given in [6, 25]. The authors of [15, 17] propose similar techniques, but they use fragments of the μ -calculus and CTL*, respectively.

We follow here the suggestion of Ehlers et al. [12], who advocate formally connecting these two approaches so as to make them interchangeable. In their words:

“Such a formal bridge should be a source of inspiration for new lines of investigation that will leverage the power of the synthesis techniques that have been developed in these two areas. [...] It would be worthwhile to develop case studies that would allow a detailed comparison of these two frameworks in terms of plant and specification modeling, computational complexity of synthesis, and implementation of derived supervisor/controller.”

As for the first remark, we show that, under reasonable assumptions, natural projection reduces to partial model checking and, when cast in a common setting, they are equivalent. To this end, we first define a common theoretical framework for both. In particular, we slightly extend both the notion of natural projection and the semantics of the μ -calculus in terms of the satisfying traces. These extensions allow us to apply natural projection to the language denoted by a specification. In addition, we extend the main property of the quotienting operator by showing that it corresponds to the natural projection of the language denoted by the specification, and vice versa (Theorem 4).

We provide additional results that contribute to the detailed comparison referred to in the second remark. Namely we define a new algorithm for partial model checking directly on LTSs (rather than on the μ -calculus). We prove our algorithm correct with respect to the traditional quotienting rules and we show it runs in polynomial time, like the algorithms based on natural projection. We have implemented this algorithm in a tool, which is available online [11]. Along with the tool, we developed several case studies illustrating its application to the synthesis of both submodules and local controllers.

Structure of the paper. Section 2 presents our unified theoretical framework for natural projection and partial model checking. Section 3 contains our main

theoretical results, in particular Theorem 4 on the equivalence of quotienting and natural projection. In Sect. 4 we introduce a novel quotienting algorithm, discuss its properties, and apply it to our running example. In Sect. 5 we briefly survey the related literature and in Sect. 6 we draw conclusions. The formal proofs together with the correctness and the complexity of our algorithm, and our experimental results are available at https://github.com/SCPTeam/pests/blob/master/proofs_and_experiments.pdf.

2 A General Framework

In this section we cast both natural projection and partial model checking in a common framework. We start with a running example: a scenario inspired by [10,33], which is an instance of Example 1.

Example 2 (Running example). A drone package delivery (DPD) system relies on unmanned aerial vehicles (UAVs) to transport goods within a given area. Drones interact with docking stations where they can pick up (action u) or deposit (action d) an item. These actions are only observable to the docking station. Additional interactions are represented by the two control actions s (for synchronize) and t (for terminate). An action t takes place when UAVs are requested to leave the station, e.g., due to a maintenance operation, while s is used for the global synchronization of both the docking station and UAVs.

Figure 1 depicts a transition system that encodes the specification of an n -position buffer $P(n)$ handled by a docking station. Intuitively, UAVs cannot perform d actions when the buffer is full (state p_n) and u actions when the buffer is empty (state p_0). Since synchronization actions s and t are immaterial, they label self-loops. \square

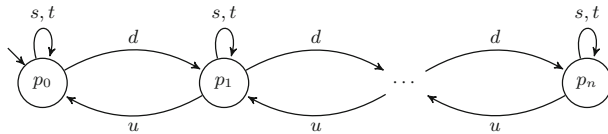


Fig. 1. An n -positions buffer specification $P(n)$.

2.1 Language Semantics Versus State Semantics

Natural projection is commonly defined over (sets of) *words* [33]. Words are finite sequences of actions, i.e., symbols labeling each transition between two states of a finite-state automaton (FSA). The language of an FSA is the set of all words that label a sequence of transitions from an initial state to some distinguished state, like a final or marking state. We call the function \mathcal{L} that maps each FSA to the corresponding *language semantics*. Given a system T and a specification S , both FSAs, then T is said to satisfy S whenever $\mathcal{L}(T) \subseteq \mathcal{L}(S)$.

For partial model checking, the specification S is defined by a formula of the μ -calculus. In this case, the standard interpretation is given by a *state semantics*, i.e., a function that given a system T and a formula Φ returns the set of states of T that satisfy Φ . Usually, T is given as a labeled transition system (LTS). An LTS is similar to an FSA, but with a weaker notion of acceptance, where all states are final. A set of evaluation rules formalizes whether a state satisfies a formula or not. Given an LTS T and a μ -calculus formula Φ , we say that T satisfies Φ whenever its initial state does.

The language semantics of temporal logics is strictly less expressive than the state-based one [16]. A similar fact holds for FSA and regular expressions [5]. Here we use a semantics from which both the state-based and the language semantics can be obtained.

2.2 Operational Model and Natural Projection

We now generalize slightly the existing approaches based on partial model checking and on supervisory control theory used for locally verifying global properties of discrete event systems. We then constructively prove that the two approaches are equally expressive so that techniques from one can be transferred to the other. To this end, we consider models expressed as (finite) labeled transition systems, which describe the behavior of discrete systems. In particular, we restrict ourselves here to deterministic transition systems.

Definition 1. A (deterministic) labeled transition system (LTS) is a tuple $A = (S_A, \Sigma_A, \rightarrow_A, \iota_A)$ where S_A is a finite set of states (with ι_A the initial state), Σ_A is a finite set of action labels, and $\rightarrow_A: S_A \times \Sigma_A \rightarrow S_A$ is the transition function. We write $t = s \xrightarrow{a} s'$ to denote a transition, whenever $\rightarrow_A(a, s) = s'$, and we call s the source state, a the action label, and s' the destination state.

A trace $\sigma \in \mathcal{T}$ of an LTS A is either a single state s or a sequence of transitions $t_1 \cdot t_2 \cdot \dots$ such that for each t_i , its destination is the source of t_{i+1} (if any). When unnecessary, we omit the source of t_{i+1} , and write a trace as the sequence $\sigma = s_0 \mathbf{a}_1 s_1 \mathbf{a}_2 s_2 \dots \mathbf{a}_n s_n$, alternating elements of S_A and Σ_A (written in boldface for readability). Finally, we denote by $\llbracket A, s \rrbracket$ the set of traces of A starting from state s and we write $\llbracket A \rrbracket$ for $\llbracket A, \iota_A \rrbracket$, i.e., for those traces starting from the initial state ι_A . □

Example 3. Consider again our running example. Figure 2 depicts the LTSs A and B . A models an UAV that deposits (d) two items in the buffer and performs a synchronization action (s). Optionally, A can also leave the docking station (t). In contrast, B repeatedly picks an item and synchronizes. Both A and B may also leave the docking station (t). Notice that the traces $\llbracket A \rrbracket$ starting from the initial state of A are $\llbracket A \rrbracket = \{q_0, q_0 \mathbf{d} q_1, q_0 \mathbf{t} q_3, q_0 \mathbf{d} q_1 \mathbf{d} q_2, q_0 \mathbf{d} q_1 \mathbf{d} q_2 \mathbf{s} q_0, \dots\}$. In contrast, the traces starting from the initial state of B are

$$\llbracket B \rrbracket = \{r_0, r_0 \mathbf{u} r_1, r_0 \mathbf{t} r_2, r_0 \mathbf{u} r_1 \mathbf{s} r_0, r_0 \mathbf{u} r_1 \mathbf{s} r_0 \mathbf{u} r_1, \dots\}. \quad \square$$

Typically, a system, or *plant* in control theory terms, consists of multiple interacting components running in parallel. Below we rephrase the synchronous

product of [33]. Intuitively, when two LTSs are put in parallel, each proceeds asynchronously, except on those actions they share, upon which they synchronize.

Definition 2. Given two LTSs A and B such that $\Sigma_A \cap \Sigma_B = \Gamma$, the synchronous product of A and B is $A \parallel B = (S_A \times S_B, \Sigma_A \cup \Sigma_B, \rightarrow_{A \parallel B}, \langle \iota_A, \iota_B \rangle)$, where $\rightarrow_{A \parallel B}$ is as follows:

$$\begin{aligned} \langle s_A, s_B \rangle &\xrightarrow{a}_{A \parallel B} \langle s'_A, s_B \rangle \text{ if } s_A \xrightarrow{a}_A s'_A \text{ and } a \in \Sigma_A \setminus \Gamma \\ \langle s_A, s_B \rangle &\xrightarrow{b}_{A \parallel B} \langle s_A, s'_B \rangle \text{ if } s_B \xrightarrow{b}_B s'_B \text{ and } b \in \Sigma_B \setminus \Gamma \\ \langle s_A, s_B \rangle &\xrightarrow{\gamma}_{A \parallel B} \langle s'_A, s'_B \rangle \text{ if } s_A \xrightarrow{\gamma}_A s'_A, s_B \xrightarrow{\gamma}_B s'_B, \text{ and } \gamma \in \Gamma. \end{aligned} \quad \square$$

Example 4. Consider again the LTSs A and B of Fig. 2. Their synchronous product $A \parallel B$ (with $\Gamma = \{s, t\}$) is depicted in Fig. 3. □



Fig. 2. From left to right, two UAVs adding to (A) and removing from (B) the buffer.

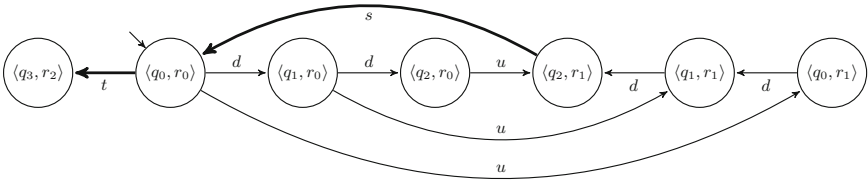


Fig. 3. Synchronous product $A \parallel B$ where bold transitions denote synchronous moves.

Next, we generalize the natural projection on languages, which is a kind of inverse operation of the synchronous product of two LTSs. Given a computation of $A \parallel B$, the natural projection extracts the relevant trace of one of the LTSs, including the synchronized transitions (see the second case below). Note that, unlike the definition given, for example in [33], our definition returns a sequence of transitions, including both states and actions. We also define the inverse projection in the expected way.

Definition 3. For LTSs A and B with $\Gamma = \Sigma_A \cap \Sigma_B$, the natural projection on A of a trace σ , in symbols $P_A(\sigma)$, is defined as follows:

$$\begin{aligned} P_A(\langle s_A, s_B \rangle) &= s_A \\ P_A(\langle s_A, s_B \rangle \xrightarrow{a}_{A\parallel B} \langle s'_A, s'_B \rangle \cdot \sigma) &= s_A \xrightarrow{a}_A s'_A \cdot P_A(\sigma) && \text{if } a \in \Sigma_A \\ P_A(\langle s_A, s_B \rangle \xrightarrow{b}_{A\parallel B} \langle s_A, s'_B \rangle \cdot \sigma) &= P_A(\sigma) && \text{if } b \in \Sigma_B \setminus \Gamma. \end{aligned}$$

Natural projection on second component B is analogously defined. We extend the natural projection to sets of traces in the usual way: $P_A(T) = \{P_A(\sigma) \mid \sigma \in T\}$.

The inverse projection of a trace σ over an LTS $A \parallel B$, in symbols $P_A^{-1}(\sigma)$, is defined as $P_A^{-1}(\sigma) = \{\sigma' \mid P_A(\sigma') = \sigma\}$. Its extension to sets is $P_A^{-1}(T) = \bigcup_{\sigma \in T} P_A^{-1}(\sigma)$. \square

Example 5. Consider the two traces $\sigma_1 = \langle q_0, r_0 \rangle \mathbf{d}\langle q_1, r_0 \rangle \mathbf{u}\langle q_1, r_1 \rangle \mathbf{d}\langle q_2, r_1 \rangle \mathbf{s}\langle q_0, r_0 \rangle$ and $\sigma_2 = \langle q_0, r_0 \rangle \mathbf{u}\langle q_0, r_1 \rangle \mathbf{d}\langle q_1, r_1 \rangle \mathbf{d}\langle q_2, r_1 \rangle \mathbf{s}\langle q_0, r_0 \rangle$. We have that the projections $P_A(\sigma_1) = P_A(\sigma_2) = q_0 \mathbf{d}q_1 \mathbf{d}q_2 \mathbf{s}q_0 \in \llbracket A \rrbracket$, and $\sigma_1, \sigma_2 \in P_B^{-1}(q_0 \mathbf{d}q_1 \mathbf{d}q_2 \mathbf{s}q_0)$. Also notice that all the traces of the form $(\langle q_0, r_0 \rangle \mathbf{d})^* \sigma_1$ belong to $P_B^{-1}(q_0 \mathbf{d}q_1 \mathbf{d}q_2 \mathbf{s}q_0)$. \square

Two classical properties [33] concerning the interplay between the synchronous product and the natural projection hold, the proofs of which are trivial.

Fact 1. $P_A(\llbracket A \parallel B \rrbracket) \subseteq \llbracket A \rrbracket$ and $\llbracket A \parallel B \rrbracket = P_B^{-1}(\llbracket A \rrbracket) \cap P_A^{-1}(\llbracket B \rrbracket)$.

2.3 Equational μ -calculus and Partial Model Checking

Below, we recall the variant of μ -calculus commonly used in partial model checking called *modal equations* [1]. A specification is given as a sequence of modal equations, and one is typically interested in the value of the top variable that is the simultaneous solution to all the equations. Equations have variables on the left-hand side and assertions on the right-hand side. Assertions are built from the boolean constants *ff* and *tt*, variables x , boolean operators \wedge and \vee , and modalities for necessity $[\cdot]$ and possibility $\langle \cdot \rangle$. Equations also have fix-point operators (minimum μ and maximum ν) over variables x , and can be organized in equation systems.

Definition 4 (Syntax of the μ -calculus). Given a set of variables $x \in X$ and an alphabet of actions $a \in \Sigma$, assertions $\varphi, \varphi' \in \mathcal{A}$ are given by the syntax:

$$\varphi ::= \text{ff} \mid \text{tt} \mid x \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi' \mid [a]\varphi \mid \langle a \rangle \varphi.$$

An equation is $x =_\pi \varphi$, where $\pi \in \{\mu, \nu\}$, μ denotes a minimum fixed point equation, and ν a maximum one. An equation system Φ is a possibly empty sequence (ϵ) of equations, where each variable x occurs in the left-hand side of at most a single equation. Thus Φ is given by

$$\Phi ::= x =_\pi \varphi; \Phi \mid \epsilon.$$

A top assertion $\Phi \downarrow x$ projects the simultaneous solution of an equation system Φ onto the top variable x . \square

We define the semantics of modal equations in terms of the traces of an LTS by extending the usual state semantics of [1] as follows. First, given an assertion φ , its state semantics $\|\varphi\|_\rho$ is given by the set of states of an LTS that satisfy φ in the context ρ , where the function ρ assigns meaning to variables. The boolean connectives are interpreted as intersection and union. The possibility modality $\|\langle a \rangle \varphi\|_\rho$ (respectively, the necessity modality $\|[a]\varphi\|_\rho$) denotes the states for which some (respectively, all) of their outgoing transitions labeled by a lead to states that satisfy φ . For more details on μ -calculus see [8, 22].

Definition 5 (Semantics of the μ -calculus [1]). *Let A be an LTS, and $\rho : X \rightarrow 2^{S_A}$ be an environment that maps variables to sets of A 's states. Given an assertion φ , the state semantics of φ is the mapping $\|\cdot\| : \mathcal{A} \rightarrow (X \rightarrow 2^{S_A}) \rightarrow 2^{S_A}$ inductively defined as follows.*

$$\begin{aligned} \|\text{ff}\|_\rho &= \emptyset & \|\text{tt}\|_\rho &= S_A & \|x\|_\rho &= \rho(x) \\ \|\varphi \wedge \varphi'\|_\rho &= \|\varphi\|_\rho \cap \|\varphi'\|_\rho & \|[a]\varphi\|_\rho &= \{s \in S_A \mid \forall s'.s \xrightarrow{a} s' \Rightarrow s' \in \|\varphi\|_\rho\} \\ \|\varphi \vee \varphi'\|_\rho &= \|\varphi\|_\rho \cup \|\varphi'\|_\rho & \|\langle a \rangle \varphi\|_\rho &= \{s \in S_A \mid \exists s'.s \xrightarrow{a} s' \wedge s' \in \|\varphi\|_\rho\} \end{aligned}$$

We extend the state semantics from assertions to equation systems. First we introduce some auxiliary notation. The empty mapping is represented by $[\]$, $[x \mapsto U]$ is the environment where U is assigned to x , and $\rho \circ \rho'$ is the mapping obtained by composing ρ and ρ' . Given a function $f(U)$ on the powerset of S_A , let $\pi U.f(U)$ be the corresponding fixed-point. We now define the semantics of equation systems by:

$$\begin{aligned} \|\epsilon\|_\rho &= [\] \\ \|x = \pi U.\varphi; \Phi\|_\rho &= R(U^*) \quad \text{where } U^* = \pi U.\|\varphi\|_{\rho \circ R(U)} \\ &\text{and } R(U) = [x \mapsto U] \circ \|\Phi\|_{\rho \circ [x \mapsto U]}. \end{aligned}$$

Finally, for top assertions, let $\|\Phi \downarrow x\|$ be a shorthand for $\|\Phi\|_{[\]}(x)$. □

Note that whenever we apply function composition \circ , its arguments have disjoint domains. Next, we present the trace semantics: a trace starting from a state s satisfies φ if s does.

Definition 6. *Given an LTS A , an environment ρ , and a state $s \in S_A$, the trace semantics of an assertion φ is a function $\langle\langle \cdot \rangle\rangle : \mathcal{A} \rightarrow S_A \rightarrow (X \rightarrow 2^{S_A}) \rightarrow \mathcal{T}$, which we also extend to equation systems, defined as follows.*

$$\langle\langle \varphi \rangle\rangle_\rho^s = \begin{cases} \llbracket A, s \rrbracket & \text{if } s \in \|\varphi\|_\rho \\ \emptyset & \text{otherwise} \end{cases} \quad \langle\langle \Phi \rangle\rangle_\rho = \lambda x. \bigcup_{s \in \|\Phi\|_\rho(x)} \llbracket A, s \rrbracket.$$

We write $\langle\langle \Phi \downarrow x \rangle\rangle$ in place of $\langle\langle \Phi \rangle\rangle_{[\]}(x)$. □

Example 6. Consider $\Phi \downarrow x$ where $\Phi = \{x =_\mu [d]y \wedge \langle u \rangle \text{tt}; y =_\nu \langle d \rangle x \vee \langle s \rangle x\}$. We compute $\|\Phi \downarrow x\|$ with respect to $A \parallel B$. $\|\Phi \downarrow x\| = U^* = \mu U.F(U)$, where $F(U) = \|[d]y \wedge \langle u \rangle \text{tt}\|_{[x \mapsto U, y \mapsto G(U)]}$ and $G(U) = \nu U'.\|\langle d \rangle x \vee \langle s \rangle x\|_{[x \mapsto U, y \mapsto U']} =$

$\|\langle d \rangle x \vee \langle s \rangle x\|_{[x \mapsto U]}$ (since y does not occur in the assertion). Following the Knaster-Tarski theorem, we compute $U^* = \bigcup^n F^n(\emptyset)$:

1. $G(\emptyset) = \|\langle d \rangle x \vee \langle s \rangle x\|_{[x \mapsto \emptyset]} = \emptyset$ and $U^1 = F(\emptyset) = \|[d]y \wedge \langle u \rangle tt\|_{[x \mapsto \emptyset, y \mapsto \emptyset]} = \{\langle q_2, r_0 \rangle\}$ (i.e., the only state that admits u but not d).
2. $G(\{\langle q_2, r_0 \rangle\}) = \|\langle d \rangle x \vee \langle s \rangle x\|_{[x \mapsto \{\langle q_2, r_0 \rangle\}]} = \{\langle q_1, r_0 \rangle\}$ (since $\langle q_1, r_0 \rangle \xrightarrow{d} \langle q_2, r_0 \rangle$) and $U^2 = F(\{\langle q_2, r_0 \rangle\}) = \|[d]y \wedge \langle u \rangle tt\|_{[x \mapsto \{\langle q_2, r_0 \rangle\}, y \mapsto \{\langle q_1, r_0 \rangle\}]} = \{\langle q_0, r_0 \rangle, \langle q_2, r_0 \rangle\}$.
3. $G(U^2) = \|\langle d \rangle x \vee \langle s \rangle x\|_{[x \mapsto \{\langle q_0, r_0 \rangle, \langle q_2, r_0 \rangle\}]} = \{\langle q_1, r_0 \rangle, \langle q_2, r_1 \rangle\}$ and $U^3 = F(U^2) = \|[d]y \wedge \langle u \rangle tt\|_{[x \mapsto U^2, y \mapsto G(U^2)]} = \{\langle q_0, r_0 \rangle, \langle q_2, r_0 \rangle\}$.

Since $U^2 = U^3$, we have obtained the fixed point U^* . Finally, we can compute $\langle\langle \Phi \downarrow x \rangle\rangle$, which amounts to $\llbracket A, \langle q_0, r_0 \rangle \rrbracket \cup \llbracket A, \langle q_2, r_0 \rangle \rrbracket$. \square

We now define when an LTS satisfies an equation system. Recall that $\llbracket A \rrbracket$ stands for $\llbracket A, \iota_A \rrbracket$.

Definition 7. *An LTS A satisfies a top assertion $\Phi \downarrow x$, in symbols $A \models_s \Phi \downarrow x$, if and only if $\iota_A \in \|\Phi \downarrow x\|$. Moreover, let $A \models_\sigma \Phi \downarrow x$ if and only if $\llbracket A \rrbracket \subseteq \langle\langle \Phi \downarrow x \rangle\rangle$. \square*

The following fact relates the notion of satisfiability defined in terms of state semantics (\models_s) with the one based on trace semantics (\models_σ); its proof is immediate by Definition 6.

Fact 2. *$A \models_s \Phi \downarrow x$ if and only if $A \models_\sigma \Phi \downarrow x$.*

As previously mentioned, partial model checking is based on the *quotienting* operation \parallel . Roughly, the idea is to specialize the specification of a composed system on a particular component. Below, we define the *quotienting* operation [1] on the LTS $A \parallel B$. Quotienting reduces $A \parallel B \models_s \Phi$ to solving $B \models_s \Phi \downarrow x \parallel_B A$. Note that each equation of the system Φ gives rise to a system of equations, one for each state s_i of A , all of the same kind, minimum or maximum (thus forming a π -block [3]). This is done by introducing a fresh variable x_{s_i} for each state s_i . Intuitively, the equation $x_{s_i} = \pi \varphi \parallel_{\Sigma_B} s_i$ represents the requirements on B when A is in state s_i . Since the occurrence of the variables on the right-hand side depends on the A 's transitions, $\Phi \downarrow x \parallel_B A$ embeds the behavior of A .

Definition 8. *Given a top assertion $\Phi \downarrow x$, we define the quotienting of the assertion on an LTS A with respect to an alphabet Σ_B as follows.*

$$\Phi \downarrow x \parallel_{\Sigma_B} A = (\Phi \parallel_{\Sigma_B} A) \downarrow x_{\iota_A}, \text{ where}$$

$$\epsilon \parallel_{\Sigma_B} A = \epsilon \quad (x = \pi \varphi; \Phi) \parallel_{\Sigma_B} A = \begin{cases} x_{s_1} = \pi \varphi \parallel_{\Sigma_B} s_1 \\ \vdots \\ x_{s_n} = \pi \varphi \parallel_{\Sigma_B} s_n \end{cases} ; \Phi \parallel_{\Sigma_B} A \quad (\forall s_i \in S_A)$$

$$x \parallel_{\Sigma_B} s = x_s \quad tt \parallel_{\Sigma_B} s = tt \quad ff \parallel_{\Sigma_B} s = ff$$

$$\begin{aligned}
 \varphi \vee \varphi' //_{\Sigma_B} s &= \varphi //_{\Sigma_B} s \vee \varphi' //_{\Sigma_B} s & \varphi \wedge \varphi' //_{\Sigma_B} s &= \varphi //_{\Sigma_B} s \wedge \varphi' //_{\Sigma_B} s \\
 \langle a \rangle \varphi //_{\Sigma_B} s &= \bigvee_{s \xrightarrow{a} s'} \varphi //_{\Sigma_B} s' & [a] \varphi //_{\Sigma_B} s &= \bigwedge_{s \xrightarrow{a} s'} \varphi //_{\Sigma_B} s' & \text{if } a \in \Sigma_A \setminus \Gamma \\
 \langle b \rangle \varphi //_{\Sigma_B} s &= \langle b \rangle (\varphi //_{\Sigma_B} s) & [b] \varphi //_{\Sigma_B} s &= [b] (\varphi //_{\Sigma_B} s) & \text{if } b \in \Sigma_B \setminus \Gamma \\
 \langle \gamma \rangle \varphi //_{\Sigma_B} s &= \bigvee_{s \xrightarrow{\gamma} s'} \langle \gamma \rangle (\varphi //_{\Sigma_B} s') & [\gamma] \varphi //_{\Sigma_B} s &= \bigwedge_{s \xrightarrow{\gamma} s'} [\gamma] (\varphi //_{\Sigma_B} s') & \text{if } \gamma \in \Gamma.
 \end{aligned}$$

□

Example 7. Consider the top assertion $\Phi \downarrow x$ of Example 6 and the LTSs A and B of Example 3. Quotienting $\Phi \downarrow x$ against B , we obtain $\Phi //_{\Sigma_A} B \downarrow x_{r_0}$ where

$$\Phi //_{\Sigma_A} B = \begin{cases} x_{r_0} =_{\mu} [d]y_{r_0} \wedge tt \\ x_{r_1} =_{\mu} [d]y_{r_1} \wedge ff \\ x_{r_2} =_{\mu} [d]y_{r_2} \wedge ff \\ y_{r_0} =_{\nu} \langle d \rangle x_{r_0} \vee ff \\ y_{r_1} =_{\nu} \langle d \rangle x_{r_1} \vee \langle s \rangle x_{r_0} \\ y_{r_2} =_{\nu} \langle d \rangle x_{r_2} \vee ff \end{cases} = \begin{cases} x_{r_0} =_{\mu} [d]y_{r_0} \\ x_{r_1} =_{\mu} ff \\ x_{r_2} =_{\mu} ff \\ y_{r_0} =_{\nu} \langle d \rangle x_{r_0} \\ y_{r_1} =_{\nu} \langle d \rangle x_{r_1} \vee \langle s \rangle x_{r_0} \\ y_{r_2} =_{\nu} \langle d \rangle x_{r_2} \end{cases} = \begin{cases} x_{r_0} =_{\mu} [d]y_{r_0} \\ y_{r_0} =_{\nu} \langle d \rangle x_{r_0}. \end{cases}$$

The leftmost equations are obtained by applying the rules of Definition 8. Then we simplify on the right-hand sides. Finally we reduce the number of equations by removing those unreachable from the top variable x_{r_0} . For a detailed description of our simplification strategies we refer the reader to [3]. By proceeding as in Example 6, we obtain $\{q_0, q_2, q_3\}$ as the fixpoint. □

3 Unifying the Logical and the Operational Approaches

In this section we prove the equivalence between natural projection and partial model checking (Theorem 4). To start, we introduce an auxiliary definition that roughly acts as a quotienting of an environment ρ . Below, we will write $\bigoplus_{i \in I} \rho_i$ for the finite composition of functions ρ_i over the elements of an index set I .

Definition 9. Given a synchronous product $A \parallel B$, we define $\Delta_B(\cdot) : (X \rightarrow 2^{S_A \times S_B}) \rightarrow (X_{S_A} \rightarrow 2^{S_B})$ as

$$\Delta_B(\rho) = \bigoplus_{x \in \text{Dom}(\rho)} \bigoplus_{s_A \in S_A} [x_{s_A} \mapsto U_B^x(s_A)], \text{ where } U_B^x(s_A) = \{s_B \mid \langle s_A, s_B \rangle \in \rho(x)\}.$$

□

A technical lemma follows. Intuitively, quotienting an assertion (and an environment) preserves the semantics, i.e., a state $\langle s_A, s_B \rangle$ satisfies φ if and only if s_B satisfies the quotient of φ on B . Indeed, the following statement can be rewritten as $\|\varphi //_{\Sigma_B} s_A\|_{\Delta_B(\rho)} = \{s_B \mid \langle s_A, s_B \rangle \in \|\varphi\|_{\rho}\}$.

Lemma 1. For all A, B, ρ and φ on $A \parallel B$, $\langle s_A, s_B \rangle \in \|\varphi\|_{\rho} \iff s_B \in \|\varphi //_{\Sigma_B} s_A\|_{\Delta_B(\rho)}$.

We next extend Lemma 1 to a system of equations, providing an alternative view of quotienting an assertion on a component of a synchronous product.

Lemma 2. For all A, B, ρ and Φ on $A \parallel B$, $\Delta_B(\|\Phi\|_\rho) = \|\Phi\|_{\Sigma_B A} \parallel_{\Delta_B(\rho)}$.

The following corollary is immediate (recall that x_{s_A} is the variable corresponding to the quotient of x on s_A).

Corollary 1. For all A, B, ρ, x and Φ on $A \parallel B$,

$$\langle s_A, s_B \rangle \in \|\Phi\|_\rho(x) \iff s_B \in \|\Phi\|_{\Sigma_B A} \parallel_{\Delta_B(\rho)}(x_{s_A}).$$

We next establish the correspondence between quotienting and natural projection.

Theorem 3. For all A, B, x and Φ on $A \parallel B$, $\langle\langle \Phi \downarrow x \parallel_{\Sigma_B A} \rangle\rangle = P_B(\langle\langle \Phi \downarrow x \rangle\rangle)$.

The following theorem states that the synchronous product of two LTSs satisfies a global equation system if and only if its components satisfy their quotients, i.e., their local assertions.

Theorem 4. For all A, B, x and Φ on $A \parallel B$,

$$A \parallel B \models_\varsigma \Phi \downarrow x \quad (\varsigma \in \{s, \sigma\})$$

if and only if any of the following equivalent statements holds:

1. $A \models_\varsigma \Phi \downarrow x \parallel_{\Sigma_A} B$
2. $B \models_\varsigma \Phi \downarrow x \parallel_{\Sigma_B} A$
3. $A \models_\sigma P_A(\langle\langle \Phi \downarrow x \rangle\rangle)$
4. $B \models_\sigma P_B(\langle\langle \Phi \downarrow x \rangle\rangle)$.

4 Quotienting Finite-State Systems

In this section we present an algorithm for quotienting a finite-state system defined as an LTS. Afterwards, we prove its correctness with respect to the standard quotienting operator and we study its complexity. Finally, we apply it to our working example to address three problems: verification, submodule construction, and controller synthesis.

4.1 Quotienting Algorithm

Our algorithm consists of two procedures that are applied sequentially. The first, called **quotient** (Table 2), builds a non-deterministic transition system from two LTSs, i.e., a specification P and an agent A . Moreover, it takes as an argument the alphabet of actions of the new transition system. Non-deterministic transition systems have a distinguished label λ , and serve as an intermediate representation. The states of the resulting transition system include all the pairs of states of P and A , except for those that denote a violation of P (line 1). The transition relation (line 3) is defined using the quotienting rules from Sect. 2. Also, note that the relation \rightarrow is restricted to the states of S (denoted \rightarrow_S).

The second procedure, called **unify** (given in Table 3) translates a non-deterministic transition system back to an LTS. Using closures over λ , **unify**

groups states of the transition system. This process is similar to the standard subset construction algorithm [19], except that we put an $a \in \Sigma_B \setminus \Gamma$ transition between two groups Q and M only if (i) M is the intersection of the λ -closures of the states reachable from Q with an a transition and (ii) all the states of Q admit at least an a transition leading to a state of M (\wedge -move). Procedure **unify** works as follows. Starting from the λ -closure of B 's initial state (line 1) it repeats a partition generation cycle (lines 4–13). Each cycle removes an element Q from the set S of the partitions to be processed. Then, for all the actions in $\Sigma_B \setminus \{\lambda\}$, a partition M is computed by \wedge -move (line 7). If the partition is nonempty, a new transition is added from Q to M (line 9). Also, if M is a freshly generated partition, i.e., $M \notin R$, it is added to both S and R (line 10). The procedure terminates when no new partitions are generated.

Table 2. The quotienting algorithm.

```

Begin proc quotient
  input P =  $\langle S_P, \Sigma_P, \rightarrow_P, i_P \rangle$ 
  input A =  $\langle S_A, \Sigma_A, \rightarrow_A, i_A \rangle$ 
  input  $\Sigma_B$ 

  1:  $S := (S_P \times S_A) \setminus \bigcup_{a \in \Sigma_A} \{(s_P, r_A) \mid s_P \xrightarrow{a} r_A \wedge r_A \xrightarrow{a} r'_A\}$ 
  2:  $i := (i_P, i_A)$ 
  3:  $\rightarrow := \bigcup_{s_P} \left\{ \begin{array}{l} \bigcup_{a \in \Sigma_A \setminus \Gamma} \{(s_P, r_A), \lambda, (s'_P, r'_A)\} \mid s_P \xrightarrow{a} r'_P \wedge r_A \xrightarrow{a} r'_A\} \\ \bigcup_{a \in \Sigma_B \setminus \Gamma} \{(s_P, r_A), a, (s'_P, r'_A)\} \mid s_P \xrightarrow{a} r'_P\} \\ \bigcup_{a \in \Gamma} \{(s_P, r_A), a, (s'_P, r'_A)\} \mid s_P \xrightarrow{a} r'_P \wedge r_A \xrightarrow{a} r'_A\} \end{array} \right.$ 
  4: B :=  $\langle S, \Sigma_B, \rightarrow_S, i \rangle$ 
  5: output unify(B)
End proc
    
```

Our quotienting algorithm is correct with respect to the quotienting operator and runs in PTIME. Intuitively, we avoid an exponential blow-up in our contribution (in contrast to Table 1) since we only consider deterministic transition systems. Notice that a determinization step for non-deterministic transition systems is exponential in the worst case.

4.2 Prototype and Application to the Running Example

We implemented the algorithm presented above as part of a tool suite for the partial evaluation of finite state models called the *partial evaluator of simple transition systems* (PESTS).¹ We applied the prototype to some case studies,

¹ The tools in our library work on FSA.

Table 3. The unification algorithm.

Begin proc unify input $B = \langle S_B, \Sigma_B, \rightarrow_B, i_B \rangle$ 1: $I := \lambda\text{-close}(\{i_B\})$ 2: $R, S := \{I\}$ 3: $\rightarrow := \emptyset$ 4: while $S \neq \emptyset$ do 5: $Q := \text{pick\&remove}(S)$ 6: for each $a \in \Sigma_B \setminus \{\lambda\}$ 7: $M := \wedge\text{-move}(Q, a)$ 8: if $M \neq \emptyset$ then 9: $\rightarrow := \rightarrow \cup \{(Q, a, M)\}$ 10: if $M \notin R$ then $S := S \cup \{M\}; R := R \cup \{M\}$ end if 11: end if 12: end for 13: end while 14: output $\langle R, \Sigma_B \setminus \{\lambda\}, \rightarrow, I \rangle$ end proc	Begin proc $\wedge\text{-move}$ input Q input a 1: $M := \lambda\text{-close}(\bigcap_{q \in Q} \{q' q \xrightarrow{a}_B q'\})$ 2: output M end proc
---	---

including a real world one based on a flexible manufacturing system.² For the sake of presentation we only show here the application to the running example. In particular, we leverage our algorithm to address three different problems: (i) reducing the verification of a parallel composition to that of a single component, (ii) synthesizing a submodule that respects a global specification (SCP), and (iii) synthesizing a controller for a given component (CSP).

Verification. Here we want to check whether $A \parallel B \not\models_s P(2)$. To do this we follow the approach of [1], i.e., we start by quotienting the specification $P(2)$ against A (see Fig. 2). The result is a two-state specification P' having a single transition labeled with t . Clearly $B \not\models_s P'$ and a counterexample is the trace $\sigma = \mathbf{r_0ur_1sr_0tr_2}$, as $\sigma \in \llbracket B \rrbracket$ while $\sigma \notin \llbracket P' \rrbracket$. As a consequence $A \parallel B \not\models_s P$. Intuitively, this is because after the two d actions, A performs a single s , which is insufficient to delimit a “safety zone” for actions u by B (which might occur too late, e.g., after another d by A). Thus, P' does not allow an s that might permit A to carry out the third d move before a u action.

Fixing the example. Given to the previous reasoning, we cannot synthesize meaningful submodules and controllers starting from A . To fix our example, we therefore replace A with A' , as depicted in Fig. 4. A' resembles A but has an extra state q_4 that enables a second s transition. Intuitively, it represents the “safety zone” just described.

SCP. We now apply the quotienting algorithm to A' in the case of buffer sizes 2 and 3 to construct the submodules that comply with $P(2)$ and $P(3)$, respectively.

² *FlexFact* <http://www.rt.eei.uni-erlangen.de/FGdes/productionline.html>.

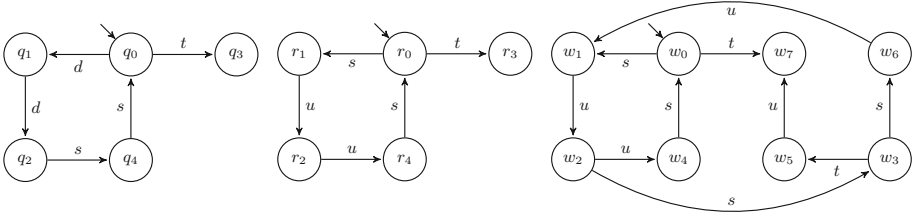


Fig. 4. Graphical representation of A' , B_2 and B_3 (from left to right).

Thus, we set $\Sigma_B = \{u, s, t\}$. In this way, the quotienting algorithm generates a component that not only synchronizes through actions s and t , but also performs actions u autonomously. The resulting agents B_2 and B_3 appear in Fig. 4.

The agent B_2 synchronizes on the first s transition of A' to ensure that both the d actions have been performed. Afterwards, the buffer must be cleared (two u actions) before synchronizing again on s (thereby permitting A' to cycle). Synchronizing on t is also possible. In this case, no further modifications of the buffer happen.

With a buffer of size 3, the agent B_3 is more complex. Intuitively, A' can perform its two d actions only when the buffer contains at most one item. Thus, B_3 has two loops. The inner loop (passing through the states $w_0 w_1 w_2 w_4$) is analogous to that of B_2 (where two u actions are performed in sequence) and, if completed, it empties the buffer. Moreover, the specification includes an external loop ($w_2 w_3 w_6 w_1$) that removes two elements from the full buffer of size 3. As expected, the two cycles can be combined. Finally, notice that the action t can occur under two conditions: if the buffer contains no items (w_0) or exactly 1 item (w_3). In the second case, a final u action (w_5) can occur.

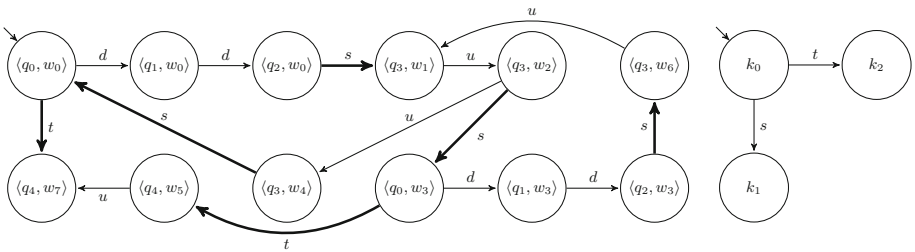


Fig. 5. Graphical representation of $A' \parallel B_3$ and C (from left to right).

CSP. We consider now the problem of synthesizing a controller for $A' \parallel B_3$ (see Fig. 5). In particular, we want a controller to enforce $P(2)$ on it.³ To generate the controller, we run the quotienting algorithm with $\Sigma_B = \{s, t\}$, i.e., we force the

³ Notice that $A' \parallel B_3$ does not comply with $P(2)$ as B_3 was synthesized from $P(3)$.

algorithm to build a component that can only synchronize on the *controllable* actions s and t . The resulting controller is depicted on the right of Fig. 5. Intuitively, the controller only admits two operations: either t or s . The first case is when A' and B_3 terminate (state $\langle q_4, w_7 \rangle$). Otherwise, only a single s action can occur. In fact, after one s action, the target reaches $\langle q_3, w_1 \rangle$ completely filling in the stack with two d actions. The system can then reach both $\langle q_3, w_2 \rangle$ and $\langle q_3, w_4 \rangle$. Since an s action leads from $\langle q_3, w_2 \rangle$ to $\langle q_0, w_3 \rangle$, where the system can perform other two d transitions, it is not allowed.

5 Related Work

Natural projection is mostly used by the community working on control theory and discrete-event systems. In the 1980s, the seminal works by Wonham et al. (e.g., [34, 35]) exploited natural projection-based algorithms for synthesizing both local and global controllers. Along this line of research, other authors proposed extensions and refinements, relying on natural projection (e.g., see [13, 14, 24, 32]).

Partial model checking has been successfully applied to the synthesis of controllers. Given an automaton representing a plant and a μ -calculus formula, Basu and Kumar [6] compute the quotient of the specification with respect to the plant. The satisfiability of the resulting formula is checked using a tableau that also returns a valid model yielding the controller. Their tableau works similarly to our quotienting algorithm, but applies to a more specific setting, as they are interested in generating controllers. In contrast, Martinelli and Matteucci [26] use partial model checking to generate a control process for a partially unspecified system in order to guarantee the compliance against a μ -calculus formula. The generated controller takes the form of an edit automaton [7].

Some authors proposed techniques based on the formal verification of temporal logics for addressing CSP. Arnold et al. [4] were among the first to control a deterministic plant with a μ -calculus specification. Also Ziller and Schneider [36] and Riedwge and Pinchinat [29] reduce the problem of synthesizing a controller to check the satisfiability of a formula of (a variant of) the μ -calculus. A similar approach was presented by Jiang and Kumar [20] and Gromyko et al. [17]. Similarly to [29, 36], [20] puts forward an approach that reduces the problem of synthesizing a controller to that of checking a CTL* formula's satisfiability. In contrast, [17] proposes a method based on symbolic model checking to synthesize a controller. Their approach applies to a fragment of CTL.

6 Conclusion

Our work goes in the same direction of [12] and provides results that build a new bridge between supervisory control theory and formal verification. In particular, we have formally established the relationship between partial model checking and natural projection by reducing natural projection to partial model checking and proving them equivalent under common assumptions. This equivalence helps

explain why some authors use partial model checking and others use natural projection to synthesize controllers. We have also developed a working prototype that we have applied to our running example and to a realistic case study.

Besides establishing an interesting and novel connection, our work also opens new directions for investigation. Since natural projection is related to language theory in general, there could be other application fields where partial model checking can be used as an alternative. The original formulation of partial model checking applies to the μ -calculus, while our quotienting algorithm works on LTSs. To the best of our knowledge, no quotienting algorithms exist for formalisms with a different expressive power, such as LTL or CTL.

Data Availability Statement and Acknowledgments. All the experiments and artifacts generated during the current study are available in the github repository: <https://doi.org/10.6084/m9.figshare.5918707.v1>.

This work was partially supported by SNSF funded project IZK0Z2_168370 “Enforceable Security Policies in Fog Computing” and by CINI Cybersecurity National Laboratory within the project FilieraSicura: Securing the Supply Chain of Domestic Critical Infrastructures from Cyber Attacks (www.filierasicura.it) funded by CISCO Systems Inc. and Leonardo SpA.

References

1. Andersen, H.R.: Partial model checking (extended abstract). In: Proceedings of Tenth Annual IEEE Symposium on Logic in Computer Science, pp. 398–407. IEEE Computer Society Press (1995)
2. Andersen, H.R., Lind-Nielsen, J.: MuDiv: a tool for partial model checking. Demo Presentation at CONCUR (1996)
3. Andersen, H.R., Lind-Nielsen, J.: Partial model checking of modal equations: a survey. *Int. J. Softw. Tools Technol. Transf.* **2**(3), 242–259 (1999)
4. Arnold, A., Vincent, A., Walukiewicz, I.: Games for synthesis of controllers with partial observation. *Theor. Comput. Sci.* **1**(303), 7–34 (2003)
5. Baeten, J.C.M., Luttik, B., Muller, T., Van Tilburg, P.: Expressiveness modulo bisimilarity of regular expressions with parallel composition. *Mathe. Struct. Comput. Sci.* **26**, 933–968 (2016)
6. Basu, S., Kumar, R.: Quotient-based approach to control of nondeterministic discrete-event systems with μ -calculus specification (2006). <http://home.eng.iastate.edu/~rkumar/PUBS/acc06-muctrl.pdf>
7. Bauer, L., Ligatti, J., Walker, D.: More enforceable security policies. In: Foundations of Computer Security, Copenhagen, Denmark, July 2002
8. Bradfield, J., Stirling, C.: Modal Mu-calculi. In: Handbook of Modal Logic, vol. 3. Elsevier Science (2006)
9. Cassandras, C.G., Lafortune, S.: Introduction to Discrete Event Systems. Kluwer, Boston (1999)
10. Open Fog Consortium: Out of the Fog: Use Case Scenarios. Supply Chain, High-Scale Drone Package Delivery, July 2016. <https://www.openfogconsortium.org/wp-content/uploads/OpenFog-Transportation-Drone-Delivery-Use-Case.pdf>. Accessed Jan 2017

11. Costa, G., Basin, D., Bodei, C., Degano, P., Galletta, L.: Pests: Partial evaluator of simple transition systems. GitHub, February 2018. <https://github.com/SCPTeam/pests>, <https://doi.org/10.6084/m9.figshare.5918707.v1>
12. Ehlers, R., Lafortune, S., Tripakis, S., Vardi, M.: Bridging the gap between supervisory control and reactive synthesis: case of full observation and centralized control. *IFAC Proc. Volumes* **47**(2), 222–227 (2014)
13. Feng, L., Wonham, W.M.: TCT: a computation tool for supervisory control synthesis. In: *Proceedings of 2006 8th International Workshop on Discrete Event Systems*, pp. 388–389 (2006)
14. Feng, L., Wonham, W.M.: On the computation of natural observers in discrete-event systems. *Discret. Event Dyn. Syst.* **20**(1), 63–102 (2010)
15. Feuillade, G., Pinchinat, S.: Modal specifications for the control theory of discrete event systems. *Discret. Event Dyn. Syst.* **17**(2), 211–232 (2007)
16. Giacobazzi, R., Ranzato, F.: States vs. traces in model checking by abstract interpretation. In: Hermenegildo, M.V., Puebla, G. (eds.) *SAS 2002*. LNCS, vol. 2477, pp. 461–476. Springer, Heidelberg (2002). <https://doi.org/10.1007/3-540-45789-5.32>
17. Gromyko, A., Pistore, M., Traverso, P.: A tool for controller synthesis via symbolic model checking. In: *8th International Workshop on Discrete Event Systems*, pp. 475–476, July 2006
18. Groote, J.F., Mousavi, M.R.: *Modeling and Analysis of Communicating Systems*. The MIT Press, Cambridge (2014)
19. Hopcroft, J.E., Motwani, R., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*, 3rd edn. Addison-Wesley Longman Publishing Co., Inc., Boston (2006)
20. Jiang, S., Kumar, R.: Supervisory control of discrete event systems with ctl^* temporal logic specifications. *SIAM J. Control Optim.* **44**(6), 2079–2103 (2006)
21. Jirásková, G., Masopust, T.: On a structural property in the state complexity of projected regular languages. *Theoret. Comput. Sci.* **449**, 93–105 (2012)
22. Kozen, D.: Results on the propositional mu-calculus. *Theor. Comput. Sci.* **27**, 333–354 (1983)
23. Lang, F., Mateescu, R.: Partial model checking using networks of labelled transition systems and boolean equation systems. In: Flanagan, C., König, B. (eds.) *TACAS 2012*. LNCS, vol. 7214, pp. 141–156. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-28756-5.11>
24. Lin, F., Wonham, W.M.: Decentralized supervisory control of discrete-event systems. *Inf. Sci.* **44**(3), 199–224 (1988)
25. Martinelli, F., Matteucci, I.: Synthesis of local controller programs for enforcing global security properties. In: *3rd International Conference on Availability, Reliability and Security (ARES)*, pp. 1120–1127, March 2008
26. Martinelli, F., Matteucci, I.: A framework for automatic generation of security controller. *Softw. Test. Verif. Reliab.* **22**(8), 563–582 (2012)
27. Merlin, P., Bochmann, G.V.: On the construction of submodule specifications and communication protocols. *ACM Trans. Program. Lang. Syst. (TOPLAS)* **5**(1), 1–25 (1983)
28. Moor, T., Schmidt, K., Perk, S.: libFAUDES – an open source C++ library for discrete event systems. In: *9th International Workshop on Discrete Event Systems*, pp. 125–130, May 2008
29. Riedweg, S., Pinchinat, S.: Quantified Mu-calculus for control synthesis. In: Rován, B., Vojtáš, P. (eds.) *MFCS 2003*. LNCS, vol. 2747, pp. 642–651. Springer, Heidelberg (2003). <https://doi.org/10.1007/978-3-540-45138-9.58>

30. Rudie, K., Grigorov, L.: Integrated Discrete-Event Systems (IDES), Department of Electrical and Computer Engineering, Queen's University in Kingston, ON, Canada. <https://qshare.queensu.ca/Users01/rudie/www/software.html>. Accessed Feb 2017
31. Su, R., Wonham, W.M.: Global and local consistencies in distributed fault diagnosis for discrete-event systems. *IEEE Trans. Autom. Control* **50**(12), 1923–1935 (2005)
32. Wong, K.C.: On the complexity of projections of discrete-event systems. In: *Proceedings of IEEE Workshop on Discrete Event Systems*, pp. 201–208 (1998)
33. Wonham, W.M.: Supervisory control of discrete-event systems, Department of Electrical and Computer Engineering, University of Toronto, ON, Canada. <http://www.control.toronto.edu/DES>. Accessed Feb 2017
34. Wonham, W.M., Ramadge, P.J.: On the supremal controllable sublanguage of a given language. In: *Proceedings of the 23rd IEEE Conference on Decision and Control*, pp. 1073–1080, December 1984
35. Wonham, W.M., Ramadge, P.J.: Modular supervisory control of discrete-event systems. *Math. Control Signals Syst.* **1**(1), 13–30 (1988)
36. Ziller, R., Schneider, K.: Combining supervisor synthesis and model checking. *ACM Trans. Embed. Comput. Syst.* **4**(2), 331–362 (2005)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

