

## REGULAR MODEL CHECKING UPSIDE-DOWN: AN INVARIANT-BASED APPROACH

JAVIER ESPARZA <sup>a</sup>, MICHAEL RASKIN <sup>b</sup>, AND CHRISTOPH WELZEL-MOHR <sup>a</sup>

<sup>a</sup> Technical University of Munich, Boltzmannstraße 3, 85748 Garching b. München  
*e-mail address:* [esparza@in.tum.de](mailto:esparza@in.tum.de), [welzel@in.tum.de](mailto:welzel@in.tum.de)

<sup>b</sup> LaBRI, University of Bordeaux, CNRS - UMR 5800, F-33405 Talence CEDEX  
*e-mail address:* [mraskin@u-bordeaux.fr](mailto:mraskin@u-bordeaux.fr)

**ABSTRACT.** Regular model checking is a technique for the verification of infinite-state systems whose configurations can be represented as finite words over a suitable alphabet. The form we are studying applies to systems whose set of initial configurations is regular, and whose transition relation is captured by a length-preserving transducer. To verify safety properties, regular model checking iteratively computes automata recognizing increasingly larger regular sets of reachable configurations, and checks if they contain unsafe configurations. Since this procedure often does not terminate, acceleration, abstraction, and widening techniques have been developed to compute a regular superset of the reachable configurations.

In this paper, we develop a complementary procedure. Instead of approaching the set of reachable configurations from below, we start with the set of all configurations and approach it from above. We use that the set of reachable configurations is equal to the intersection of all inductive invariants of the system. Since this intersection is non-regular in general, we introduce  $b$ -invariants, defined as those representable by CNF-formulas with at most  $b$  clauses. We prove that, for every  $b \geq 0$ , the intersection of all inductive  $b$ -invariants is regular, and we construct an automaton recognizing it. We show that whether this automaton accepts some unsafe configuration is in EXPSpace for every  $b \geq 0$ , and PSPACE-complete for  $b = 1$ . Finally, we study how large must  $b$  be to prove safety properties of a number of benchmarks.

### 1. INTRODUCTION

Regular model checking (RMC) is a framework for the verification of different classes of infinite-state systems (see, e.g., the surveys [AJNS04, Abd12, AST18, Abd21]). In its canonical version, RMC is applied to systems satisfying the following conditions: configurations can be encoded as words, the set of initial configurations is recognized by a finite automaton

*Key words and phrases:* parameterized verification, finite automata, regular model-checking.

This paper is an extended version of [ERW22b]. In comparison to the proceedings version, this extended version contains complete and detailed proofs, some additional examples, and a more general notion of inductive invariants based on interpreting regular descriptions (see Section 4.2). This work was partially funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 787367 (PaVeS).

$\mathcal{I}$ , and the transition relation is recognized by a length-preserving transducer  $\mathcal{T}$ . RMC algorithms address the problem of, given a regular set of unsafe configurations, deciding if its intersection with the set of reachable configurations is empty or not. In the present paper, we do not consider generalisations to non-length-preserving or non-finite-state-transducer transitions.

The fundamental building block of current RMC algorithms is an automata-theoretic construction that, given a non-deterministic automaton (NFA)  $A$  recognizing a regular set of configurations, produces another NFA recognizing the set of immediate successors (or predecessors) of  $L(A)$  with respect to the transition relation represented by  $\mathcal{T}$ . Therefore, if some unsafe configuration is reachable, one can find a witness by, starting with the automaton  $\mathcal{I}$  for the set of initial configurations, repeatedly adding the set of immediate successors. However, this approach almost never terminates when all reachable configurations are safe. Research on RMC has produced many acceleration, abstraction, and widening techniques to make the iterative computation “jump over the fixpoint” in finite time, and produce an invariant of the system not satisfied by any unsafe configuration (see, e.g., [BJNT00, JN00, DLS01, AJNd02, BLW03, BHV04, BT12, BHRV12, Leg12, CHLR17]).

In this paper, we develop a complementary approach that, starting with the set of all configurations, computes increasingly smaller regular inductive invariants, i.e., sets of configurations closed under the reachability relation and containing all initial configurations. Our main contribution is the definition of a sequence of *regular* inductive invariants that converges (in the limit) to the set of reachable configurations, and for which automata can be directly constructed from  $\mathcal{I}$  and  $\mathcal{T}$ .

While some of the previous work (e.g. using abstraction [BHV04]) does include over-approximation, this comes from replacing automata with smaller more permissive ones and then using acceleration techniques. In contrast, our work directly constructs regular-language-described properties that are satisfied by all the reachable configurations of the original system.

Our starting point is the fact that the set of reachable configurations is equal to the intersection of all inductive invariants. Since this intersection is non-regular in general, we introduce *b-invariants*. An invariant is *b*-bounded, or just a *b*-invariant, if for every  $\ell \geq 0$ , the configurations of length  $\ell$  satisfying the invariant are those satisfying a Boolean formula in conjunctive normal form with at most  $b$  clauses. The atomic propositions are claims of the form “at the position  $i$  there is character  $x$ ”. For example, assume that the configurations of some system are words over the alphabet  $\{a, b, c, d\}$ , and that the configurations of length five where the second letter is an  $a$  or the fourth letter is a  $b$ , and the second letter is a  $b$  or the third is a  $c$ , constitute an inductive invariant. Then this set of configurations is a 2-invariant, represented by the formula  $(a_{2:5} \vee b_{4:5}) \wedge (b_{2:5} \vee d_{3:5})$ . We prove that, for every bound  $b \geq 0$ , the intersection of all inductive *b*-invariants, denoted  $IndInv_b$ , is regular, and recognized by a DFA of double exponential size in  $\mathcal{I}$  and  $\mathcal{T}$ . As a corollary, we obtain that, for every  $b \geq 0$ , deciding if  $IndInv_b$  contains some unsafe configuration is in EXPSPACE. Moreover, the proof inspires a wider class of regular inductive invariants. Introducing this class is a central contribution of this paper which is an extended version of [ERW22b]. To define such an invariant, we pick a finite transducer and a regular language of words, the invariant being the image of the language under the action of the transducer. Generally, we pick a single transducer to analyse a given regular transition system, and study the invariants provided by different regular languages. We show that the EXPSPACE upper bound still holds for the wider class.

In the second part of the paper, we study the special case  $b = 1$  in more detail. We exploit that inductive 1-invariants are closed under union (a special feature of the  $b = 1$  case), and prove that deciding whether  $IndInv_1$  contains some unsafe configuration is PSPACE-complete. The proof also shows that  $IndInv_1$  can be recognized by an NFA of single exponential size in  $\mathcal{I}$  and  $\mathcal{T}$ .

The index  $b$  of a bounded invariant can be seen as a measure of how difficult it is for a human to understand it. So one is interested in the smallest  $b$  such that  $IndInv_b$  is strong enough to prove a given property. In the third and final part of the paper, we experimentally show that for a large number of systems  $IndInv_1$  is strong enough to prove useful safety properties.

*Related work.* The work closest to ours is [ADHR07], which directly computes an overapproximation of the set of reachable configurations of a parameterized system. Contrary to our approach, the paper computes one single approximation, instead of a converging sequence of overapproximations. Further, the method is designed for a model of parameterized systems with existential or universal guarded commands, while our technique can be applied to any model analyzable by RMC. Our work is also related to [CHLR17], which computes an overapproximation using a learning approach, which terminates if the set of reachable configurations is regular; our paper shows that a natural class of invariants is regular, and that automata for them can be constructed explicitly from the syntactic description of the system. This paper generalizes the work of [BEI<sup>+</sup>20, ERW21a, BIS21, ERW21b] on trap invariants for parameterized Petri nets. Trap invariants are a special class of 1-bounded invariants, and the parameterized Petri nets studied in these papers can be modelled in the RMC framework. An alternative to regular model checking are logical based approaches. The invisible invariant method synthesizes candidate invariants from examples, which are then checked for inductiveness [PRZ01]. Our approach does not produce candidates, it generates invariants by construction. Modern tools like Ivy [PMP<sup>+</sup>16, MP20] have verified more complex protocols than the ones in Section 7 using a combination of automation and human interaction. The best way of achieving this interaction is beyond the scope of this paper, which focuses on the foundations of regular model checking.

*Structure of the paper.* Section 2 introduces basic definitions of the RMC framework. Section 3 introduces  $b$ -bound invariants while Section 4 proves regularity of  $IndInv_b$ . Sections 5 and 6 prove the PSPACE-completeness result. Sections 7 and 8 contain some experimental results and conclusions.

## 2. PRELIMINARIES

Given  $n, m \in \mathbb{N}$ , we let  $[n, m]$  denote the set  $\{i \in \mathbb{N} : n \leq i \leq m\}$ .

**Languages and automata.** A language over a finite alphabet  $\Sigma$  is a subset of  $\Sigma^*$ . An element of a language is called a word; for a word  $w$  we denote  $w[i]$  its  $i$ -th letter. Given a language  $L \subseteq \Sigma^*$ , we let  $\bar{L}$  denote the language  $\Sigma^* \setminus L$ . A nondeterministic finite automaton (NFA) is a tuple  $A = \langle Q, \Sigma, \Delta, Q_0, F \rangle$  where  $Q$  is a non-empty finite set of states,  $\Sigma$  is an alphabet,  $\Delta: Q \times \Sigma \rightarrow 2^Q$  is a transition function, and  $Q_0, F \subseteq Q$  are sets of initial and final states, respectively. A run of  $A$  on a word  $w \in \Sigma^\ell$  is a sequence  $q_0 q_1 \dots q_\ell$  of states such that  $q_0 \in Q_0$  and  $q_i \in \Delta(q_{i-1}, w[i])$  for every  $i \in [1, \ell]$ . A run on  $w$  is accepting if  $q_\ell \in F$ , and  $A$  accepts  $w$  if there exists an accepting run of  $A$  on  $w$ . The language recognized by  $A$ , denoted  $L(A)$  or  $L_A$ , is the set of words accepted by  $A$ . We let  $|A|$  denote the number of states of  $A$ . The function  $\delta_A: 2^Q \times \Sigma^* \rightarrow 2^Q$  is defined inductively as follows:  $\delta_A(P, \varepsilon) = P$  and  $\delta_A(P, aw) = \delta_A(\bigcup_{p \in P} \Delta(p, a), w)$ . Observe that  $A$  accepts  $w$  iff  $\delta_A(Q_0, w) \cap F \neq \emptyset$ . An NFA  $A$  is deterministic<sup>1</sup> (DFA) if  $|Q_0| = 1$  and  $|\Delta(q, a)| = 1$  for every  $q \in Q$  and  $a \in \Sigma$ .

**Length-preserving relations and length-preserving transducers.** A *length-preserving relation* over an alphabet  $\Sigma \times \Gamma$  is a language over the alphabet  $\Sigma \times \Gamma$ . We denote elements of  $\Sigma \times \Gamma$  as  $\langle a, b \rangle$  or  $\begin{bmatrix} a \\ b \end{bmatrix}$  where  $a \in \Sigma$  and  $b \in \Gamma$ . Given  $w = w[1] \dots w[\ell] \in \Sigma^*$  and  $u = u[1] \dots u[\ell] \in \Gamma^*$ , we let  $\langle w, u \rangle$  denote the word  $\langle w[1], u[1] \rangle \dots \langle w[\ell], u[\ell] \rangle \in (\Sigma \times \Gamma)^*$ . That is, if we write  $\langle w, u \rangle$  then necessarily  $w$  and  $u$  have the same length. We look at  $\langle w, u \rangle$  as a representation of the pair of words  $(w, u)$ .

**Remark 2.1.** Throughout the paper, we consider only length-preserving relations on words, and so we call them just relations.

The *complement* of a relation  $R \subseteq (\Sigma \times \Gamma)^*$  is the relation  $\bar{R} := \{\langle w, u \rangle \in (\Sigma \times \Gamma)^* \mid \langle w, u \rangle \notin R\}$ . Sometimes we represent relations by infix operators, like  $\rightsquigarrow$ , and then we write  $\text{comp}(\rightsquigarrow)$  to denote the complement. Observe that, by definition, the complement of a relation is a subset of  $(\Sigma \times \Gamma)^*$ , and so it only contains tuple  $\langle w, u \rangle$  where  $|w| = |u|$ . The *join* of two relations  $R_1 \subseteq (\Sigma_1 \times \Gamma)^*$  and  $R_2 \subseteq (\Gamma \times \Sigma_2)^*$ , is the relation  $R_1 \circ R_2 \subseteq (\Sigma_1 \times \Sigma_2)^*$  given by:  $\langle w_1, w_2 \rangle \in R_1 \circ R_2$  if there exists  $w \in \Gamma^*$  such that  $\langle w_1, w \rangle \in R_1$  and  $\langle w, w_2 \rangle \in R_2$ . The *post-image* of a language  $L \subseteq \Sigma^*$  under a relation  $R \subseteq (\Sigma \times \Gamma)^*$  is the language  $L \circ R$  given by:  $w \in L \circ R$  iff there exists  $u \in \Sigma^*$  such that  $u \in L$  and  $\langle u, w \rangle \in R$ . The *pre-image* of  $L$  under  $R$ , denoted  $R \circ L$ , is defined analogously. The *projections* of a relation  $R \subseteq (\Sigma \times \Gamma)^*$  onto its first and second components are the languages  $R|_1 := \{w \in \Sigma^* \mid \exists u \in \Gamma^* : \langle w, u \rangle \in R\}$  and  $R|_2 := \{u \in \Gamma^* \mid \exists w \in \Sigma^* : \langle w, u \rangle \in R\}$ . The *inverse* of a relation  $R \subseteq (\Sigma \times \Gamma)^*$  is the relation  $R^{-1} := \{\langle u, w \rangle \in (\Gamma \times \Sigma)^* \mid \langle w, u \rangle \in R\}$ .

A (length-preserving) *transducer* over  $\Sigma \times \Gamma$  is an NFA with  $\Sigma \times \Gamma$  as alphabet. The (length-preserving) relation recognized by a transducer  $T$ , denoted  $R(T)$  or  $R_T$ , is the set of tuples  $\langle w, u \rangle \in (\Sigma \times \Gamma)^*$  accepted by  $T$ . A relation is *regular* if it is recognized by a transducer. A transducer is *deterministic* if it is a deterministic NFA.

It is folklore that regular relations are closed under Boolean operations (the operations are implemented as for regular languages) and composition, also called join, and that the pre- and post-images of regular languages under regular relations are regular. We sketch the constructions in the following proposition.

**Proposition 2.2.** *Let  $\Sigma, \Gamma$  be finite alphabets.*

<sup>1</sup>Sometimes in the literature this is called “deterministic and complete” with determinism being a weaker property. However, unlike the size gap between NFAs and DFAs, the difference is limited to the presence of a single “useless” bottom state.

- (1) Let  $T, U$  be transducers with  $n_T$  and  $n_U$  states over alphabets  $\Sigma_T \times \Gamma$  and  $\Gamma \times \Sigma_U$ , respectively. There exists a transducer with  $\mathcal{O}(n_T n_U)$  states recognizing  $R(T) \circ R(U)$ .
- (2) Let  $A, B$  be NFAs over  $\Sigma$ , respectively  $\Gamma$ , with  $n_A$  and  $n_B$  states, respectively, and let  $T$  be a transducer over  $\Sigma \times \Gamma$  with  $n_T$  states. There exist NFAs with  $\mathcal{O}(n_A \cdot n_T)$  and  $\mathcal{O}(n_B \cdot n_T)$  states recognizing  $L(A) \circ R(T)$  and  $R(T) \circ L(B)$ , respectively.
- (3) Let  $T$  be a transducer with  $n_T$  states. There exist NFAs with  $n_T$  states recognizing  $R(T)|_1$  and  $R(T)|_2$ .
- (4) Let  $T$  be a transducer with  $n_T$  states. There exists a transducer with  $n_T$  states recognizing  $R(T)^{-1}$ .

*Proof.* 1) Let  $T = \langle Q_T, \Sigma_T \times \Gamma, \Delta_T, Q_{0T}, F_T \rangle$  and  $T_U = \langle Q_U, \Gamma \times \Sigma_U, \Delta_U, Q_{0U}, F_U \rangle$ . Define the transducer  $V := \langle Q_T \times Q_U, \Sigma_T \times \Sigma_U, \Delta, Q_{0T} \times Q_{0U}, F_T \times F_U \rangle$ . For this transducer, we fix  $(\langle q_T, q_U \rangle, \langle \sigma_T, \sigma_U \rangle, \langle q'_T, q'_U \rangle) \in \Delta$  iff there exists  $\gamma \in \Gamma$  such that  $(q_T, \langle \sigma_T, \gamma \rangle, q'_T) \in \Delta_T$  and  $(q_U, \langle \gamma, \sigma_U \rangle, q'_U) \in \Delta_U$ . We have  $R(V) = R(T) \circ R(U)$ .

2) Let  $A = \langle Q_A, \Sigma, \Delta_A, Q_{0A}, F_A \rangle$  and  $T = \langle Q_T, \Sigma \times \Gamma, \Delta_T, Q_{0T}, F_T \rangle$ . Define the NFA  $A' = \langle Q_A \times Q_T, \Gamma, \Delta, Q_{0A} \times Q_{0T}, F_A \times F_T \rangle$  where  $(\langle q_A, q_T \rangle, \gamma, \langle q'_A, q'_T \rangle) \in \Delta$  iff there exists  $\sigma \in \Sigma$  such that  $(q_A, \sigma, q'_A) \in \Delta_A$  and  $(q_T, \langle \sigma, \gamma \rangle, q'_T) \in \Delta_T$ . We have  $L(A') = L(A) \circ R(T)$ . The construction for  $R(T) \circ L(B)$  is analogous.

3) Let  $T = \langle Q_T, \Sigma \times \Gamma, \Delta_T, Q_{0T}, F_T \rangle$ . Define  $A_1 = \langle Q_T, \Sigma, \Delta, Q_{0T}, F_T \rangle$  where  $(q, \sigma, q') \in \Delta$  iff there exists  $\gamma \in \Gamma$  such that  $(q, \langle \sigma, \gamma \rangle, q') \in \Delta_T$ . We have  $L(A_1) = R|_1$ . The NFA  $A_2$  is defined analogously.

4) Let  $T = \langle Q_T, \Sigma \times \Gamma, \Delta_T, Q_{0T}, F_T \rangle$ . Define  $T^{-1} = \langle Q_T, \Gamma \times \Sigma, \Delta', Q_{0T}, F_T \rangle$  where  $\Delta'(q, \langle a, b \rangle) := \Delta_T(q, \langle b, a \rangle)$  for all  $q, a, b$ .  $\square$

**Regular model checking.** Regular model checking (RMC) is a framework for the verification of systems with infinitely many configurations. Each configuration is represented as a finite word over a fixed alphabet  $\Sigma$ . Systems are modelled as regular transition systems of the following form.

**Definition 2.3** (Regular transition systems). A *regular transition system* (RTS) is a triple  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  where  $\Sigma$  is an alphabet,  $\mathcal{I}$  is an NFA over  $\Sigma$ , and  $\mathcal{T}$  is a transducer over  $\Sigma \times \Sigma$ .

Words over  $\Sigma$  are called *configurations*. Configurations accepted by  $\mathcal{I}$  are called *initial*, and pairs of configurations accepted by  $\mathcal{T}$  are called *transitions*. We write  $w \rightsquigarrow u$  to denote that  $\langle w, u \rangle$  is a transition. Observe that  $w \rightsquigarrow u$  implies  $|w| = |u|$ . Given two configurations  $w, u$ , we say that  $u$  is *reachable* from  $w$  if  $w \rightsquigarrow^* u$  where  $\rightsquigarrow^*$  denotes the reflexive and transitive closure of  $\rightsquigarrow$ . The set of reachable configurations of  $\mathcal{R}$ , denoted  $Reach(\mathcal{R})$ , or just  $Reach$  when there is no confusion, is the set of configurations reachable from the initial configurations. In the following, we use  $|\mathcal{R}|$  to refer to  $|\mathcal{I}| + |\mathcal{T}|$ .

**Example 2.4** (Dining philosophers). We model a very simple version of the dining philosophers as an RTS, for use as running example. Philosophers sit at a round table with forks between them. Philosophers can be *thinking* ( $t$ ) or *eating* ( $e$ ). Forks can be *free* ( $f$ ) or *busy* ( $b$ ). A thinking philosopher whose left and right forks are free can *simultaneously* grab both forks—the forks become busy—and start eating. After eating, the philosopher puts both forks to the table and returns to thinking. The model includes two corner cases: a table with one philosopher and one fork, which is then both the left and the right fork (unusable

as it would need to be grabbed twice in a single transition), and the empty table with no philosophers or forks.

We model the system as an RTS  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  over the alphabet  $\Sigma = \{t, e, f, b\}$ . A configuration of a table with  $n$  philosophers and  $n$  forks is represented as a word over  $\Sigma$  of length  $2n$ . Letters at odd and even positions model the current states of philosophers and forks (positions start at 1). For example,  $tftf$  models a table with two thinking philosophers and two free forks. The set of initial configurations is  $L(\mathcal{I}) = (tf)^*$ , and the set of transitions is

$$R(\mathcal{T}) = \begin{bmatrix} t \\ e \end{bmatrix} \begin{bmatrix} f \\ b \end{bmatrix} \begin{bmatrix} x \\ x \end{bmatrix}^* \begin{bmatrix} f \\ b \end{bmatrix} \mid \begin{bmatrix} e \\ t \end{bmatrix} \begin{bmatrix} b \\ f \end{bmatrix} \begin{bmatrix} x \\ x \end{bmatrix}^* \begin{bmatrix} b \\ f \end{bmatrix} \mid \begin{bmatrix} x \\ x \end{bmatrix}^* \left( \begin{bmatrix} f \\ b \end{bmatrix} \begin{bmatrix} t \\ e \end{bmatrix} \begin{bmatrix} f \\ b \end{bmatrix} \mid \begin{bmatrix} b \\ f \end{bmatrix} \begin{bmatrix} e \\ t \end{bmatrix} \begin{bmatrix} b \\ f \end{bmatrix} \right) \begin{bmatrix} x \\ x \end{bmatrix}^*$$

where  $\begin{bmatrix} x \\ x \end{bmatrix}$  stands for the regular expression  $\left( \begin{bmatrix} t \\ t \end{bmatrix} \mid \begin{bmatrix} e \\ e \end{bmatrix} \mid \begin{bmatrix} f \\ f \end{bmatrix} \mid \begin{bmatrix} b \\ b \end{bmatrix} \right)$ . The first two terms of  $R(\mathcal{T})$  describe the actions of the first philosopher, and the second the actions of the others. It is not difficult to show that  $Reach = (t(f|beb))^* | ebt((f|beb)t)^*b$ . These are the configurations where no two philosophers are using the same fork, and fork states match their adjacent philosopher states.

**Safety verification problem for RTSs.** The safety verification problem for RTSs is defined as follows: Given an RTS  $\mathcal{R}$  and an NFA  $\mathcal{U}$  recognizing a set of *unsafe* configurations, decide whether  $Reach(\mathcal{R}) \cap L(\mathcal{U}) = \emptyset$  holds. The problem is known to be undecidable.

**Example 2.5.** A configuration  $w$  of an RTS is *deadlocked* if there is no configuration  $u$  such that  $w \rightsquigarrow u$ . It is easy to see that the set of deadlocked configurations of the dining philosophers of Example 2.4 is

$$Dead = \overline{\Sigma^* f t f \Sigma^*} \cap \overline{\Sigma^* b e b \Sigma^*} \cap \overline{t f \Sigma^* f} \cap \overline{e b \Sigma^* b}.$$

In other words, these are configurations containing neither  $f t f$  not  $b e b$  as a cyclic word. The dining philosophers are deadlock-free iff  $Reach \cap Dead = \emptyset$ , which is the case (recall that philosophers can only grab both forks simultaneously).

### 3. BOUNDED INDUCTIVE SETS OF AN RTS

We present an invariant-based approach to the safety verification problem for RTSs. Fix an RTS  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$ . We introduce an infinite sequence

$$\Sigma^* = IndInv_0 \supseteq IndInv_1 \supseteq IndInv_2 \dots \supseteq Reach$$

of effectively regular inductive invariants of  $\mathcal{R}$  that converges to  $Reach$ , i.e.,  $IndInv_k$  is effectively regular for every  $k \geq 1$ , and  $Reach = \bigcap_{k=0}^{\infty} IndInv_k$ . Section 3.1 recalls basic notions about inductive sets and invariants, and Section 3.2 defines the inductive invariant  $IndInv_b$  for every  $b \geq 0$ .

**3.1. Inductive sets and invariants.** Let  $S \subseteq \Sigma^*$  be a set of configurations.  $S$  is *inductive* if it is closed under reachability, i.e., if  $w \in S$  and  $w \rightsquigarrow u$  implies  $u \in S$ . Observe that inductive sets are closed under union and intersection.  $S$  is an *invariant for length  $\ell$*  if  $\text{Reach} \cap \Sigma^\ell \subseteq S \cap \Sigma^\ell$ , and an *invariant* if  $\text{Reach} \subseteq S$ . Observe that, since  $\mathcal{T}$  is a length-preserving transducer,  $S$  is an invariant iff it is an invariant for every length. Given two invariants  $I_1, I_2$ , we say that  $I_1$  is *stronger* than  $I_2$  if  $I_1 \subset I_2$ .

A set  $S$  is an *inductive invariant* if it is both inductive and an invariant, i.e., if it contains  $\text{Reach}$  and is closed under reachability. Every inductive invariant  $S$  and every initial configuration  $I$  satisfy  $I \in S$ . There is a unique smallest inductive invariant w.r.t. set inclusion, namely the set  $\text{Reach}$  itself.

**Example 3.1.** The set  $I_0 = ((t | e)(f | b))^*$  is an inductive invariant of the dining philosophers. Other inductive invariants are

$$\begin{aligned} I_1 &= \overline{\Sigma^* e f e \Sigma^*}, & I_2 &= \overline{e \Sigma^* e f}, & I_3 &= \overline{\Sigma^* t b t \Sigma^*}, & I_4 &= \overline{t \Sigma^* t b}, \\ I_5 &= \overline{(t \Sigma e \Sigma)^*}, & I_6 &= \overline{(e \Sigma t \Sigma)^*}. \end{aligned}$$

Taking into account that the table is round, these are the sets of configurations without any occurrence of  $e f e$  ( $I_1$  and  $I_2$ ) and  $t b t$  ( $I_3$  and  $I_4$ ) as a cyclic word; as well as without alternation of  $t$  and  $e$  throughout the *entire* configuration ( $I_5$  and  $I_6$ ). Note that the latter condition is vacuously true for an odd number of philosophers.

**3.2. Bounded inductive sets.** Given a length  $\ell \geq 0$ , we represent certain sets of configurations as Boolean formulas over a set  $AP_\ell$  of atomic propositions. More precisely, a Boolean formula over  $AP_\ell$  describes a set containing *some* configurations of length  $\ell$ , and *all* configurations of other lengths.

The set  $AP_\ell$  contains an atomic proposition  $q_{j:\ell}$  for every  $q \in \Sigma$  and for  $j \in [1, \ell]$ . A *formula*  $\varphi$  over  $AP_\ell$  is a positive Boolean combination of atomic propositions of  $AP_\ell$  and the constants *true* and *false*. Formulas are interpreted on configurations. Intuitively, an atomic proposition  $q_{j:\ell}$  states that either the configuration does *not* have length  $\ell$ , or it has length  $\ell$  and its  $j$ -th letter is  $q$ . Formally,  $w \in \Sigma^*$  *satisfies* the single atomic proposition  $q_{j:\ell}$  if either  $|w| \neq \ell$  or  $|w| = \ell$  and  $w[j] = q$ . For non-atomic formulas, i.e.,  $\varphi = \text{true}, \varphi_1 \vee \varphi_2, \varphi_1 \wedge \varphi_2$ , satisfaction is defined as usual. The language  $L(\varphi) \subseteq \Sigma^*$  of a formula is the set of configurations that satisfy  $\varphi$ . We also say that  $\varphi$  *denotes* the set  $L(\varphi)$ . A formula is inductive if it denotes an inductive set.

**Example 3.2.** In the dining philosophers, let  $\varphi = (e_{1:4} \wedge b_{4:4}) \vee f_{2:4}$ . We have

$$L(\varphi) = \epsilon \mid \Sigma \mid \Sigma^2 \mid \Sigma^3 \mid e \Sigma \Sigma b \mid \Sigma f \Sigma \Sigma \mid \Sigma^5 \Sigma^* .$$

Observe that an expression like  $(q_{1:1} \wedge r_{1:2})$  is not a formula because it combines atomic propositions of two different lengths, which is not allowed. Notice also that  $\neg q_{j:\ell}$  is equivalent to  $\bigvee_{r \in \Sigma \setminus \{q\}} r_{j:\ell}$ . Therefore, if we allowed negative atomic propositions, we would still have the same class of expressible predicates on words of a given length (and we would not obtain formulas for the same predicates with fewer clauses.) Abusing language, if  $\varphi$  is a formula over  $AP_\ell$  and  $L(\varphi)$  is an (inductive) invariant, then we also say that  $\varphi$  is an (inductive) invariant. Observe that (inductive) invariants are closed under conjunction and disjunction.

Convention: From now on, “formula” means “positive formula in CNF”.

**Definition 3.3.** Let  $b \geq 0$ . A  $b$ -formula is a formula with at most  $b$  clauses (with the convention that *true* is the only formula with 0 clauses). A set  $S \subseteq \Sigma^*$  of configurations is  $b$ -bounded if for every length  $\ell$  there exists a  $b$ -formula  $\varphi_\ell$  over  $AP_\ell$  such that  $S \cap \Sigma^\ell = L(\varphi_\ell)$ . We abbreviate  $b$ -bounded sets to just  $b$ -sets. We call a  $b$ -bounded invariant a  $b$ -invariant.

Observe that, since one can always add tautological clauses to a formula without changing its language, a set  $S$  is  $b$ -bounded iff for every length  $\ell$  there is a formula  $\varphi_\ell$  with *exactly*  $b$  clauses.

**Example 3.4.** In the dining philosophers, the 1-formulas  $(t_{2i-1:\ell} \vee e_{2i-1:\ell})$  and  $(f_{2i:\ell} \vee b_{2i:\ell})$  are inductive 1-invariants for every even  $\ell \geq 1$  and every  $i \in [1, \ell/2]$ . It follows that the set  $I_0$  of Example 3.1 is an intersection of (infinitely many) such inductive 1-invariants. The same happens for  $I_1, \dots, I_6$ . For example,  $I_1$  is the intersection of all inductive 1-invariants of the form  $(t_{i:\ell} \vee f_{i:\ell} \vee b_{i:\ell} \vee b_{i+1:\ell} \vee t_{i+1:\ell} \vee e_{i+1:\ell} \vee t_{i+2:\ell} \vee f_{i+2:\ell} \vee b_{i+2:\ell})$ , for all  $\ell \geq 1$  and all  $i \in [1, \ell - 1]$ ; inductivity is shown by an easy case distinction.

We introduce the inductive invariants studied in the paper.

**Definition 3.5.** Let  $\mathcal{R} = (\Sigma, \mathcal{I}, \mathcal{T})$  be an RTS, let  $w, u \in \Sigma^*$  be two configurations of  $\mathcal{R}$ , and let  $b \geq 0$ . We say that  $w$  is  $b$ -potentially-reachable from  $u$ , denoted  $u \overset{*}{\rightsquigarrow}_b w$ , if every inductive  $b$ -formula satisfied by  $u$  is also satisfied by  $w$ . In other words, no inductive  $b$ -set provides an explanation why  $w$  should not be reachable from  $u$ . Further, we define  $IndInv_b = \{w \in \Sigma^* \mid u \overset{*}{\rightsquigarrow}_b w \text{ for some initial configuration } u\}$ .

Observe that  $u \overset{*}{\rightsquigarrow} w$  implies  $u \overset{*}{\rightsquigarrow}_b w$  for every  $b \geq 1$ . Indeed, if  $u \overset{*}{\rightsquigarrow} w$  then, by the definition of an inductive formula, every inductive formula satisfied by  $u$  is also satisfied by  $w$ . Therefore, if  $w$  is reachable from  $u$  then it is also  $b$ -potentially reachable from  $u$  for every  $b \geq 0$ . However, the converse does not necessarily hold.

The following proposition shows that  $IndInv_b$  is an inductive invariant for every  $b \geq 0$ , and some fundamental properties.

**Proposition 3.6.** *Let  $\mathcal{R}$  be an RTS. For every  $b \geq 0$ :*

- (1)  $IndInv_b$  is an inductive invariant, and  $IndInv_b \subseteq S$  for every inductive  $b$ -invariant  $S$ .
- (2)  $IndInv_b \supseteq IndInv_{b+1}$ .
- (3)  $Reach = \bigcap_{b=0}^{\infty} IndInv_b$ .

*Proof.* (1) We first show that  $IndInv_b$  is an inductive invariant.

- $IndInv_b$  is inductive. Let  $w \in IndInv_b$  and  $w \rightsquigarrow v$ . Since  $w \in IndInv_b$ , we have  $u \overset{*}{\rightsquigarrow}_b w$  for some initial configuration  $u$ . We prove  $u \overset{*}{\rightsquigarrow}_b v$ . Let  $\varphi$  be an inductive  $b$ -formula satisfied by  $u$ . By definition,  $\varphi$  is also satisfied by  $w$  and, since  $\varphi$  is inductive and  $w \rightsquigarrow v$ , also by  $v$ .
- $IndInv_b \supseteq Reach$ . Since  $IndInv_b$  is inductive, it suffices to show that it contains all initial configurations. This follows immediately from the definition of  $IndInv_b$  and the fact that every configuration is  $b$ -potentially reachable from itself.

For the second part, let  $S$  be an arbitrary  $b$ -invariant, and let  $w \in IndInv_b$ . We prove  $w \in S$ . Let  $\ell$  be the length of  $w$ . Since  $S$  is a  $b$ -invariant, there exists a  $b$ -formula  $\varphi$  over  $AP_\ell$  such that  $S \cap \Sigma^\ell = L(\varphi)$ . So it suffices to prove  $w \models \varphi$ . Since  $w \in IndInv_b$ , there exists an initial configuration  $u$  such that every inductive  $b$ -formula satisfied by  $u$  is also satisfied by  $w$ . In particular, this also holds for  $\varphi$ . Further, since  $S$  is an invariant,  $u \models \varphi$ . So  $w \models \varphi$ , and we are done.



(2)  $IndInv_b \supseteq IndInv_{b+1}$  follows from the fact that, by definition, every  $b$ -formula is also a  $(b+1)$ -formula.

(3) For every  $\ell \geq 0$ , the set  $Reach \cap \Sigma^\ell$  is an inductive invariant for length  $\ell$ . Let  $\varphi_\ell$  be a formula over  $AP_\ell$  such that  $L(\varphi_\ell) \cap \Sigma^\ell = Reach \cap \Sigma^\ell$ , and let  $b_\ell$  be its number of clauses. (Notice that  $\varphi_\ell$  always exists because every subset of  $\Sigma^\ell$  can be expressed as a formula, and every formula can be put in conjunctive normal form.) Then  $\varphi_\ell$  is a  $b_\ell$ -bounded invariant, and so  $L(\varphi_\ell) \supseteq IndInv_{b_\ell}$  for every  $\ell \geq 0$ . So we have  $Reach = \bigcap_{\ell=0}^{\infty} L(\varphi_\ell) \supseteq \bigcap_{b=0}^{\infty} IndInv_b$ . As each  $IndInv_b$  is an invariant,  $IndInv_b \supseteq Reach$  for all  $b$ , thus  $\bigcap_{b=0}^{\infty} IndInv_b \supseteq Reach$  and we are done.  $\square$

By (1), the invariant  $IndInv_b$  is as strong as any  $b$ -invariant. Notice, however, that  $IndInv_b$  needs not be a  $b$ -invariant itself. The reason is that  $b$ -sets are not closed under intersection. Indeed, the conjunction of two formulas with  $b$  clauses is not always equivalent to a formula with  $b$  clauses, one can only guarantee equivalence to a formula with  $2b$  clauses.

**Example 3.7.** Recall that the deadlocked configurations of the dining philosophers are

$$Dead = \overline{\Sigma^* f t f \Sigma^*} \cap \overline{\Sigma^* b e b \Sigma^*} \cap \overline{t f \Sigma^* f} \cap \overline{e b \Sigma^* b}.$$

We prove  $IndInv_1 \cap Dead = \emptyset$ , which implies that the dining philosophers are deadlock-free. Let  $C$  be the set of configurations of  $((t | e)(f | b))^*$  containing no occurrence of  $e f e$  or  $t b t$  as a cyclic word, and without alternation of  $e$  and  $t$  throughout the entire configuration. In Example 3.1, we showed that  $C$  is the intersection of 1-invariants, which implies  $IndInv_1 \subseteq C$ . So it suffices to prove  $C \cap Dead = \emptyset$ . For this, let  $w \in C$ . If  $|w| \leq 3$  the proof is an easy case distinction, so assume  $|w| \geq 4$ . We show that  $w$  as a cyclic word contains an occurrence of  $f t f$  or  $b e b$ , and so it is not a deadlock. If all philosophers are thinking at  $w$ , then, since  $w$  contains no occurrence of  $t b t$ , it contains an occurrence of  $f t f$ .

If some philosopher is eating at  $w$ , assume for a contradiction that there is a deadlock. By this assumption, that philosopher must have a free fork nearby. Without loss of generality, assume that the fork is the next symbol in the configuration. Then we are looking at a  $ef$  fragment. As  $e f e$  is forbidden, the next symbol in the configuration has to be  $t$ . By assumption of a deadlock, the next symbol is  $b$ . We have obtained  $ef t b$  fragment; if we continue, we observe that eating and thinking philosophers always alternate. But this is forbidden by the last pair of invariants; thus the initial assumption of deadlock must be false.

Further, for the dining philosophers we have  $Reach = IndInv_3$ . Apart from some corner cases (e.g. an unsatisfiable invariant for every odd length), the reason is that the 3-formula

$$\begin{aligned} & (t_{i:\ell} \vee f_{i:\ell} \vee b_{i:\ell} \vee b_{i+1:\ell} \vee t_{i+1:\ell} \vee e_{i+1:\ell}) \\ & \wedge (b_{i+1:\ell} \vee t_{i+1:\ell} \vee e_{i+1:\ell} \vee t_{i+2:\ell} \vee f_{i+2:\ell} \vee b_{i+2:\ell}) \\ & \wedge (t_{i:\ell} \vee f_{i:\ell} \vee b_{i:\ell} \vee f_{i+1:\ell} \vee t_{i+1:\ell} \vee e_{i+1:\ell} \vee t_{i+2:\ell} \vee f_{i+2:\ell} \vee b_{i+2:\ell}) \end{aligned}$$

is an inductive 3-invariant for every  $\ell \geq 3$  and every  $i$  from 1 to  $\ell$ , if we interpret indices cyclically modulo  $\ell$ . To verify it, observe that any violation has one of  $ef*$ ,  $*fe$ , or  $ebe$  at the positions  $i$  through  $i+2$ . For either of the former two situations to arise anew, the middle fork needs to be newly freed — requiring the latter situation as a precondition. Conversely, for the latter situation to arise anew, one of the philosophers need to start eating, picking up a free fork next to the other philosopher, who needs to be already eating. Thus any violation of the condition requires a violation on the previous step, too, proving inductivity.

The configurations satisfying this invariant and the inductive 1-invariants  $I_0, \dots, I_6$  of Example 3.4 are the reachable configurations  $Reach = (t(f \mid beb))^* \mid ebt((f \mid beb)t)^*b$ .

Note that the 3-invariant does not rely on the 1-invariants for its inductiveness; we always require the inductive invariants to be inductive on their own independently of each other. The invariant is weaker than it could be because we have obtained it from forbidding the patterns  $ef$ ,  $fe$ , and  $ebe$ . It happens to include some obviously unreachable configurations such as  $ftft$  (note that the forks and the philosophers are swapped, the initial state for two philosophers is  $tftf$ ), but this is not a problem as simpler invariants exclude them.

**Example 3.8.** We construct an (artificial) family  $\{\mathcal{R}_b \mid b \geq 1\}$  of RTSs such that  $IndInv_b \supset Reach(\mathcal{R}_b) = IndInv_{b+1}$ . Fix some  $b \geq 1$ . Let  $\mathcal{R}_b = \{\{0, 1\}, \mathcal{I}, \mathcal{T}\}$  with  $R(\mathcal{T})$  given by the union of the languages

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix}^{k_1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \end{bmatrix}^{k_2} \left( \begin{bmatrix} 0 \\ 0 \end{bmatrix} \mid \begin{bmatrix} 0 \\ 1 \end{bmatrix} \mid \begin{bmatrix} 1 \\ 0 \end{bmatrix} \mid \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right)^*$$

for every  $k_1, k_2 \in \mathbb{N}$  such that  $k_1 + k_2 = b - 1$ . Then every transition of the RTS is of the form  $u \cdot v \rightsquigarrow u' \cdot v'$  where  $|u| = b = |u'|$ , the word  $u$  contains exactly one 1, and the word  $u'$  contains only 0s. If we choose  $0^*$  as the set of initial configurations, then no transition is applicable to any initial configuration, and so  $Reach = 0^*$ . We show that  $IndInv_b \supset IndInv_{b+1} = 0^* = Reach$ .

- $IndInv_{b+1} = 0^*$ . We first claim that, for every  $i > b$ , the formula  $\varphi_i = \bigwedge_{j=1}^b 0_{j:\ell} \wedge 0_{(i-b):\ell}$  is an inductive  $b + 1$ -invariant. It follows immediately from the definitions that  $\varphi_i$  is a  $b + 1$ -invariant. To show that  $\varphi_i$  is inductive, observe that every configuration  $w$  of length  $\ell$  satisfying  $\varphi_i$  is of the form  $w = 0^b w'$  for some word  $w'$ , and so no word  $u$  satisfies  $w \rightsquigarrow u$ . This proves the claim. Now, the only configurations satisfying  $\varphi_i$  for every  $i > b$  are those containing only 0s, and so  $IndInv_{b+1} = 0^*$ .
- $0^b 1 \in IndInv_b$ , and so  $IndInv_b \supset Reach$ . Intuitively, we want to show that every  $b$ -set containing  $0^{b+1}$  contains *some* word of the form  $0^* 1 0^*$ , and all such words can reach  $0^b 1$ , requiring all  $b$ -invariants to contain  $0^b 1$ . It suffices to prove that no inductive  $b$ -set  $\varphi = \bigwedge_{i=1}^b \varphi_i$  separates  $0^{b+1}$  and  $0^b 1$ , i.e., satisfies  $0^{b+1} \models \varphi$  while  $0^b 1 \not\models \varphi$ . For this, let us introduce a family of useful words: for every  $1 \leq i \leq b + 1$ , let  $w_i := 0^{i-1} 1 0^{b-i+1}$ . We have  $w_i \rightsquigarrow 0^b 1$  and so, since  $\varphi$  is inductive and  $0^b 1 \not\models \varphi$ , we get  $w_i \not\models \varphi$ . Intuitively, we exploit the fact that  $w_i$  differs from  $0^{b+1}$  in exactly one position. This implies that, for every  $i$ , some clause of  $\varphi$ , say  $\varphi_j$ , satisfies  $w_i \not\models \varphi_j$  and therefore contains  $0_{i:b+1}$  as an atomic proposition. So we “need” all clauses of the inductive invariant to exclude all  $w_i$ , and have none left to exclude  $0^b 1$ .

More formally, we conduct an induction on  $i$  from 1 to  $b$  to prove that there are distinct  $j_1, \dots, j_b$  such that  $\varphi_{j_i}$  contains the atomic proposition  $0_{i:b+1}$ . Consider the base case: Pick  $j_1$  such that  $w_1 \not\models \varphi_{j_1}$ , which exists because  $w_1 \not\models \varphi$ .  $0^{b+1}$  differs from  $w_1$  only in the first letter, and  $\varphi_{j_1}$  is a disjunction of atomic propositions which is satisfied by  $0^{b+1}$ . This means that  $\varphi_{j_1}$  must contain  $0_{1:b+1}$  because it is the only atomic proposition that is satisfied by  $0^{b+1}$  but not  $w_1$ . For the induction step, one observes that  $w_i$  satisfies the clauses  $\varphi_{j_1}, \dots, \varphi_{j_{i-1}}$  by induction hypothesis. Therefore,  $\varphi$  also contains a new clause  $\varphi_{j_i}$  such that  $w_i \not\models \varphi_{j_i}$ . With the same reasoning as for the base case one concludes that  $0_{i:b+1}$  is an atomic proposition in  $\varphi_{j_i}$ . This concludes the induction and, thus,  $0^b 1 \models \varphi$ .

**Example 3.9.** As a final example, we construct an RTS such that  $Reach$  is a regular language but there is no  $b$  such that  $IndInv_b = Reach$ . Let  $\mathcal{R} = (\{0, 1\}, \mathcal{I}, \mathcal{T})$  be the RTS

with  $L(\mathcal{I}) = 0^*$  and

$$R(\mathcal{T}) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}^* \begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \end{bmatrix}^* \left( \begin{bmatrix} 0 \\ 0 \end{bmatrix} \mid \begin{bmatrix} 0 \\ 1 \end{bmatrix} \mid \begin{bmatrix} 1 \\ 0 \end{bmatrix} \mid \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right).$$

For every length  $\ell$ , the RTS  $\mathcal{R}$  behaves like the RTS  $\mathcal{R}_b$  of Example 3.8 for  $b = \ell - 1$ . Therefore, we have  $Reach = 0^*$ , and so  $Reach$  is regular. Further,  $IndInv_b \neq Reach$  for every  $b \geq 1$  (as they differ when restricted to the length  $b + 1$ ).

#### 4. $IndInv_b$ IS REGULAR FOR EVERY $b \geq 1$

**4.1. Encoding  $b$ -formulas as  $b$ -powerwords.** We introduce an encoding of  $b$ -formulas. We start with some examples. Assume  $\mathcal{R}$  is an RTS with  $\Sigma = \{a, b, c\}$ . We consider formulas over  $AP_3$ , i.e., over the atomic propositions  $\{a_{1:3}, a_{2:3}, a_{3:3}, b_{1:3}, b_{2:3}, b_{3:3}, c_{1:3}, c_{2:3}, c_{3:3}\}$ .

We encode the 1-formula  $(a_{1:3} \vee a_{2:3})$  as the word  $\{a\}\{a\}\emptyset$  of length three over the alphabet  $2^\Sigma$ . Intuitively,  $\{a\}\{a\}\emptyset$  stands for the words of length 3 that have an  $a$  in their first or second position. Similarly, we encode  $(a_{1:3} \vee b_{1:3} \vee b_{3:3})$  as  $\{a, b\}\emptyset\{b\}$ . Intuitively,  $\{a, b\}\emptyset\{b\}$  stands for the set of words of length 3 that have  $a$  or  $b$  as the first letter, or  $b$  as the third letter. Since  $2^\Sigma$  is the powerset of  $\Sigma$ , we call words over  $2^\Sigma$  *powerwords*.

Consider now the 2-formula  $(a_{1:3} \vee b_{1:3} \vee a_{2:3}) \wedge (b_{1:3} \vee b_{3:3} \vee c_{3:3})$ . We put the encodings of its clauses “on top of each other”. Since the encodings of  $(a_{1:3} \vee b_{1:3} \vee a_{2:3})$  and  $(b_{1:3} \vee b_{3:3} \vee c_{3:3})$  are  $\{a, b\}\{a\}\emptyset$  and  $\{b\}\emptyset\{b, c\}$ , respectively, we encode the formula as the word

$$\begin{bmatrix} \{a, b\} \\ \{b\} \end{bmatrix} \begin{bmatrix} \{a\} \\ \emptyset \end{bmatrix} \begin{bmatrix} \emptyset \\ \{b, c\} \end{bmatrix}$$

of length three over the alphabet  $2^\Sigma \times 2^\Sigma = (2^\Sigma)^2$ . We call such a word a *2-powerword*. Similarly, we encode a  $b$ -formula over  $AP_3$  as a  *$b$ -powerword* of length three over the alphabet  $(2^\Sigma)^b$ . In the following, we overload  $\varphi$  to denote both a formula and its encoding as a  $b$ -powerword, and, for example, write

$$\varphi = \begin{bmatrix} X_{11} \\ \dots \\ X_{b1} \end{bmatrix} \cdots \begin{bmatrix} X_{1\ell} \\ \dots \\ X_{b\ell} \end{bmatrix} \quad \text{instead of} \quad \varphi = \bigwedge_{i=1}^b \bigvee_{j=1}^{\ell} \bigvee_{a \in X_{ij}} a_{i:j}$$

where  $X_{ij} \subseteq \Sigma$ . Intuitively, each row  $X_{i1} \cdots X_{i\ell}$  encodes one clause of  $\varphi$ . We also write  $\varphi = \varphi[1] \cdots \varphi[\ell]$  where  $\varphi[i] \in (2^\Sigma)^b$  denotes the  $i$ -th letter of the  $b$ -powerword encoding  $\varphi$ .

Now we show a simple lemma, which, however, provides the key to our results. For every fixed  $b$ , the satisfaction relation  $w \models \varphi$  between configurations and  $b$ -formulas is regular:

**Lemma 4.1.** *Let  $\Sigma$  be an alphabet,  $b \geq 1$ , and  $\Gamma = (2^\Sigma)^b$ . There exists a deterministic transducer  $\mathcal{V}_b$  over the alphabet  $\Sigma \times \Gamma$  with  $2^b$  states such that  $L(\mathcal{V}) = \{\langle w, \varphi \rangle \in (\Sigma \times \Gamma)^* \mid w \models \varphi\}$ .*

*Proof.* A configuration  $w = w[1]w[2] \cdots w[\ell]$  satisfies a  $b$ -formula

$$\varphi = \begin{bmatrix} X_{11} \\ \dots \\ X_{b1} \end{bmatrix} \begin{bmatrix} X_{12} \\ \dots \\ X_{b2} \end{bmatrix} \cdots \begin{bmatrix} X_{1\ell} \\ \dots \\ X_{b\ell} \end{bmatrix}$$

where  $X_{i,j} \subseteq \Gamma$  iff for every index  $1 \leq i \leq b$  there exists a position  $1 \leq j \leq \ell$  in the word  $w$  such that  $w[j] \in X_{ij}$ . We define a transducer  $\mathcal{V}_b$  over  $\Sigma \times \Gamma$  that accepts  $\langle w, \varphi \rangle$  iff this

condition holds. The transducer reads the word

$$\begin{bmatrix} w[1] \\ X_{11} \\ \dots \\ X_{b1} \end{bmatrix} \begin{bmatrix} w[2] \\ X_{12} \\ \dots \\ X_{b2} \end{bmatrix} \dots \begin{bmatrix} w[\ell] \\ X_{1\ell} \\ \dots \\ X_{b\ell} \end{bmatrix}$$

storing in its state the set of indices  $i \in \{1, \dots, b\}$  for which the position  $j$  has already been found. So the states of the transducer are the subsets of  $\{1, \dots, b\}$ , and, given two states  $S, S'$ , there is a transition from  $S$  to  $S'$  labeled by the tuple  $\langle a, X_1, \dots, X_b \rangle$  iff  $S \subseteq S'$  and  $a \in X_i$  holds for every  $i \in S' \setminus S$ . The initial state is the empty set, indicating that no position has been found yet, and the unique final state is the set  $\{1, \dots, b\}$ , indicating that all positions have been found.  $\square$

**4.2. Representations and interpreters.** The transducer  $\mathcal{V}_b$  of Lemma 4.1 recognizes the satisfaction relation  $w \models \varphi$  between configurations and  $b$ -formulas. We generalize this idea. Consider an *arbitrary* transducer  $\mathcal{V}$  over the alphabet  $\Sigma \times \Gamma$ , where  $\Gamma$  is some *arbitrary* alphabet. Now, look at a word  $W \in \Gamma^*$  (we use capital letters  $W, V, U, \dots$  to denote such words) as a *representation* of the set of configurations  $\{w \in \Sigma^* \mid \langle w, W \rangle \in L(\mathcal{V})\}$ . Intuitively,  $W$  is a *name* standing for this set of configurations, and the transducer  $\mathcal{V}$  *interprets* the meaning of  $W$ . For this reason, we call  $\mathcal{V}$  an *interpreter*.

**Definition 4.2** (Interpretation). Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be a regular transition system. An *interpretation* is a pair  $\langle \Gamma, \mathcal{V} \rangle$  where  $\Gamma$  is an alphabet and  $\mathcal{V}$  is a deterministic transducer over  $\Sigma \times \Gamma$ , called the *interpreter*<sup>2</sup>. We call words over  $\Gamma$  *representations*, and use capital letters  $W, V, U, \dots$  to denote them.  $W \in \Gamma^*$  *represents* or *stands for* the set of all configurations  $w \in \Sigma^*$  such that  $\langle w, W \rangle$  is accepted by  $\mathcal{V}$ . We write  $w \models_{\mathcal{V}} W$  to denote that  $w$  is one the words represented by  $W$ .

**Example 4.3.** The transducer of Lemma 4.1 is a particular interpreter with alphabet  $\Gamma = (2^{\Sigma})^b$ . For example, for  $b = 2$  the transducer interprets the word

$$\begin{bmatrix} \{a, b\} \\ \{b\} \end{bmatrix} \begin{bmatrix} \{a\} \\ \emptyset \end{bmatrix} \begin{bmatrix} \emptyset \\ \{b, c\} \end{bmatrix}$$

as the set of configurations that satisfy the formula  $(a_{1:3} \vee b_{1:3} \vee a_{2:3}) \wedge (b_{1:3} \vee b_{3:3} \vee c_{3:3})$ .

In the rest of the section, we define the set  $IndInv_{\mathcal{V}}$  for an arbitrary interpreter  $\mathcal{V}$ , and prove that it is regular. Instantiating  $\mathcal{V}$  as the transducer  $\mathcal{V}_b$  of Lemma 4.1, we obtain as a corollary that  $IndInv_b$  is regular for every  $b \geq 1$ . In order to define  $IndInv_{\mathcal{V}}$ , observe that an interpreter  $\mathcal{V}$  may interpret some representations as inductive sets of configurations, and others as non-inductive sets. We define the set of all representations that are inductive.

**Definition 4.4.** Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be a regular transition system and let  $\mathcal{V}$  be an interpreter. A representation  $W \in \Gamma^*$  is *inductive* if  $u \rightsquigarrow w$  and  $u \models_{\mathcal{V}} W$  implies  $w \models_{\mathcal{V}} W$ . We define the set

$$Ind_{\mathcal{V}} = \{W \in \Gamma^* \mid \text{for every } u, w \in \Sigma^*, \text{ if } u \rightsquigarrow w \text{ and } u \models_{\mathcal{V}} W, \text{ then } w \models_{\mathcal{V}} W\}.$$

We write  $u \overset{*}{\rightsquigarrow}_{\mathcal{V}} w$ , and say that  $w$  is *potentially reachable* from  $u$  with respect to  $\mathcal{V}$  if for every inductive representation  $W \in Ind_{\mathcal{V}}$ , if  $u \models_{\mathcal{V}} W$ , then  $w \models_{\mathcal{V}} W$  too. Further, we define

$$IndInv_{\mathcal{V}} := \{w \in \Sigma^* \mid u \overset{*}{\rightsquigarrow}_{\mathcal{V}} w \text{ for some initial configuration } u\}.$$


<sup>2</sup>Observe that we require the transducer to be deterministic (as is the case in Lemma 4.1).

We have the following fact:

**Fact 4.5.**  $IndInv_{\mathcal{V}} \supseteq Reach$  for every interpreter  $\mathcal{V}$ , i.e.,  $IndInv_{\mathcal{V}}$  is an overapproximation of the set of reachable configurations.

*Proof.* If  $w \in Reach$ , then  $v \rightsquigarrow^* w$  holds for some initial configuration  $v$ . We prove  $v \rightsquigarrow_{\mathcal{V}}^* w$ , which implies  $w \in IndInv_{\mathcal{V}}$ . Since  $v \rightsquigarrow^* w$ , every inductive set containing  $v$  contains  $w$  as well. So, in particular, for every  $W \in Ind_{\mathcal{V}}$ , if  $v \models_{\mathcal{V}} W$ , then  $w \models_{\mathcal{V}} W$ .  $\square$

To illustrate the virtue of this generalization, consider another interpretation for the RTS from Example 3.9. Abstractly speaking, we consider a set of atomic propositions and enforce that none of these propositions is true. The interpreter  $\mathcal{V}$  for this is



where  $M$  are all pairs  $\langle \sigma, U \rangle$  with  $\sigma \notin U$  while  $H$  are all pairs with  $\sigma \in U$ . One can quickly verify that here  $\{1\}^* \subseteq Ind_{\mathcal{V}}$  and, thus,  $Reach = 0^* = IndInv_{\mathcal{V}}$  as there is no 1 in any configuration – a fact that is true initially, and, since then there is no transition applicable, throughout every step.

In the rest of the section, we prove that  $IndInv_{\mathcal{V}}$  is an effectively regular set of configurations for any interpreter  $\mathcal{V}$ . As a first step, we show that  $Ind_{\mathcal{V}}$  is a regular set of representations. This follows immediately from the following proposition, proving that the complement of  $Ind_{\mathcal{V}}$  is regular.

**Lemma 4.6.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be an RTS where  $\mathcal{T}$  has  $n_{\mathcal{T}}$  states, and let  $\mathcal{V}$  be an interpreter with  $n_{\mathcal{V}}$  states. One can effectively compute an NFA with at most  $n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2$  states recognizing the set of representations  $\overline{Ind_{\mathcal{V}}}$ .*

*Proof.* Let  $Id_{\Gamma} = \{\langle W, W \rangle \mid W \in \Gamma^*\}$  be the identity relation on  $\Gamma$ . By definition of  $Ind_{\mathcal{V}}$ , we have

$$\begin{aligned} & \overline{Ind_{\mathcal{V}}} \\ &= \{W \in \Gamma^* \mid \exists u, w \in \Sigma^* \text{ s.t. } u \rightsquigarrow w, u \models_{\mathcal{V}} W \text{ and } w \not\models_{\mathcal{V}} W\} \\ &= \{W \in \Gamma^* \mid \exists u, w \in \Sigma^* \text{ s.t. } \langle u, w \rangle \in R_{\mathcal{T}}, \langle u, W \rangle \in R_{\mathcal{V}}, \langle w, W \rangle \in \overline{R_{\mathcal{V}}}\} \\ &= \{W \in \Gamma^* \mid \exists u, w \in \Sigma^* \text{ s.t. } \langle W, u \rangle \in R_{\mathcal{V}}^{-1}, \langle u, w \rangle \in R(T), \langle w, W \rangle \in \overline{R_{\mathcal{V}}}\} \\ &= \{W \in \Gamma^* \mid \langle W, W \rangle \in R_{\mathcal{V}}^{-1} \circ R(T) \circ \overline{R_{\mathcal{V}}}\} \\ &= ((R_{\mathcal{V}}^{-1} \circ R(T) \circ \overline{R_{\mathcal{V}}}) \cap Id_{\Gamma})|_1. \end{aligned}$$

Since  $\mathcal{V}$  is deterministic, there are transducers for  $R(\mathcal{V})$  and  $comp(R(\mathcal{V}))$  with  $\mathcal{O}(n_{\mathcal{V}})$  states. Applying Proposition 2.2(1) twice and Proposition 2.2(3), we obtain an NFA for  $\overline{Ind_{\mathcal{V}}}$  with at most  $n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2$  states.  $\square$

Using standard automata constructions, we get immediately:

**Lemma 4.7.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be an RTS where  $\mathcal{T}$  has  $n_{\mathcal{T}}$  states, and let  $\mathcal{V}$  be an interpreter with  $n_{\mathcal{V}}$  states. One can effectively compute an DFA with at most  $2^{n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states recognizing the set of representations  $Ind_{\mathcal{V}}$ .*

We prove that the potential reachability relation is effectively regular, that is, the relation is recognized by a transducer that can be effectively constructed from  $\mathcal{T}$  and  $\mathcal{V}$ . Again, we show that its complement is regular.

**Lemma 4.8.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be a regular transition system with  $n_{\mathcal{T}}$  states, and let  $\mathcal{V}$  be an interpreter with  $n_{\mathcal{V}}$  states. One can effectively compute a nondeterministic transducer with at most  $n_{\mathcal{V}}^2 \cdot 2^{n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states recognizing  $\text{comp}(\rightsquigarrow_{\mathcal{V}}^*)$ .*

*Proof.* By definition, we have

$$\begin{aligned} & \text{comp}(\rightsquigarrow_{\mathcal{V}}^*) \\ &= \{ \langle u, w \rangle \in (\Sigma \times \Sigma)^* \mid \exists W \in \text{Ind}_{\mathcal{V}} \text{ s.t. } u \models_{\mathcal{V}} W \text{ and } w \not\models_{\mathcal{V}} W \} \\ &= \left\{ \langle u, w \rangle \in (\Sigma \times \Sigma)^* \mid \exists W \in \text{Ind}_{\mathcal{V}} \text{ s.t. } \langle u, W \rangle \in R_{\mathcal{V}} \text{ and } \langle W, w \rangle \in \overline{R_{\mathcal{V}}^{-1}} \right\}. \end{aligned}$$

Let  $S = \{ \langle W, W \rangle \mid W \in \text{Ind}_{\mathcal{V}} \}$ . We then have  $\text{comp}(\rightsquigarrow_{\mathcal{V}}^*) = (R_{\mathcal{V}} \circ S \circ \text{comp}(R_{\mathcal{V}}^{-1}))$ . Since  $\mathcal{V}$  is deterministic, there is a transducer for  $\text{comp}(R_{\mathcal{V}}^{-1})$  with  $\mathcal{O}(n_{\mathcal{V}})$  states. Further, by Lemma 4.6 there is a DFA for  $\text{Ind}_{\mathcal{V}}$  with at most  $2^{n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states, and so a transducer for  $S$  with the same number of states. Applying Proposition 2.2, we can construct a nondeterministic transducer for  $\text{comp}(\rightsquigarrow_{\mathcal{V}}^*)$  with at most  $n_{\mathcal{V}}^2 \cdot 2^{n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states.  $\square$

Again, we get:

**Lemma 4.9.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be a regular transition system with  $n_{\mathcal{T}}$  states, and let  $\mathcal{V}$  be an interpreter with  $n_{\mathcal{V}}$  states. One can effectively compute a deterministic transducer with at most  $2^{2^{\log(n_{\mathcal{V}}^2)} \cdot n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states recognizing  $\rightsquigarrow_{\mathcal{V}}^*$ .*

We combine the previous results to show that, given an RTS  $\mathcal{R}$  and an interpretation  $\mathcal{V}$ , the set  $\text{IndInv}_{\mathcal{V}}$  of potentially reachable configurations is recognized by an NFA with double exponentially many states in  $\mathcal{T}$  and  $\mathcal{V}$ .

**Theorem 4.10.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be an RTS, and let  $\mathcal{V}$  be an interpretation. Let  $n_{\mathcal{I}}$ ,  $n_{\mathcal{T}}$ , and  $n_{\mathcal{V}}$  be the number of states of  $\mathcal{I}$ ,  $\mathcal{T}$ , and  $\mathcal{V}$ , respectively. Then  $\text{IndInv}_{\mathcal{V}}$  is recognized by an NFA with at most  $n_{\mathcal{I}} \cdot 2^{2^{\log(n_{\mathcal{V}}^2)} \cdot n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states.*

*Proof.* We have  $\text{IndInv}_{\mathcal{V}} = L_{\mathcal{I}} \circ (\rightsquigarrow_{\mathcal{V}}^*)$  by the definition of  $\text{IndInv}_{\mathcal{V}}$ . By Lemma 4.9,  $\rightsquigarrow_{\mathcal{V}}^*$  is recognized by a deterministic transducer with  $2^{2^{\log(n_{\mathcal{V}}^2)} \cdot n_{\mathcal{T}} \cdot n_{\mathcal{V}}^2}$  states. Apply Proposition 2.2.  $\square$

We apply Theorem 4.10 to the interpreter of  $b$ -formulas given in Lemma 4.7, and obtain:

**Corollary 4.11.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be an RTS. Let  $b \geq 1$ , and let  $f(n_{\mathcal{T}}, b) := 2 \cdot n_{\mathcal{T}} \log n_{\mathcal{T}} \cdot 2^{b+1}$ . Then  $\text{IndInv}_b$  (Definition 3.5) is recognized by a DFA with at most  $n_{\mathcal{I}} \cdot 2^{2^{f(n_{\mathcal{T}}, b)}}$  states.*

*Proof.* By Lemma 4.7, there is a deterministic transducer with at most  $2^{2^{(2n_{\mathcal{T}} \log n_{\mathcal{T}} \cdot n_{\mathcal{V}})}}$  states recognizing  $\rightsquigarrow_{\mathcal{V}}^*$ . So, there is a deterministic transducer with at most  $2^{2^{f(n_{\mathcal{T}}, b)}}$  states for  $\rightsquigarrow_b^*$ . Apply Proposition 2.2(2).  $\square$


Given an instance  $\mathcal{R}$ ,  $\mathcal{U}$  of the safety verification problem and a fixed  $b \geq 0$ , if the set  $\text{IndInv}_b$  satisfies  $\text{IndInv}_b \cap L(\mathcal{U}) = \emptyset$ , then  $\mathcal{R}$  is safe. By Corollary 4.11, deciding whether  $\text{IndInv}_b \cap L(\mathcal{U}) = \emptyset$  is in EXSPACE for every fixed  $b$ . Indeed, the theorem and its proof show that there is a DFA recognizing  $\text{IndInv}_b \cap L(\mathcal{U})$  such that one can guess an accepting

path of it, state by state, using exponential space. Indeed, for fixed  $b$ , the interpreter of Lemma 4.1 has a constant number of states  $n_\gamma$ . Therefore, storing one state of the transducer of Lemma 4.9 takes polynomial space, and storing one state of the transducer of Theorem 4.10 takes exponential space. Currently, we do not know if there is a  $b$  such that the problem is EXSPACE-complete for every  $b' \geq b$ . In the next two sections, we show that for  $b = 1$  the problem is actually PSPACE-complete.

5. DECIDING  $IndInv_1 \cap L(\mathcal{U}) = \emptyset$  IS IN PSPACE

We give a non-deterministic polynomial-space algorithm that decides  $IndInv_1 \cap L(\mathcal{U}) = \emptyset$ . As a byproduct, we show that  $IndInv_1$  is recognized by an NFA with a single exponential number of states.

As a running example for the following construction we are introducing a basic token passing protocol: There is a line of agents. Any single agent either holds a token ( $t$ ) or not ( $n$ ). The initial language is  $t n^*$ ; that is, initially there is exactly one token. The transitions of the system allow the token to be passed down the line or, if the token already is at the last position, to be passed back to the front. One can understand this protocol as a token passing algorithm in a circle of agents. The transducer for the transitions looks as follows:



We fix an RTS  $\mathcal{R} = (\Sigma, \mathcal{I}, \mathcal{T})$  for the rest of the section. 1-formulas have a special property: since the disjunction of two clauses is again a clause, the disjunction of two 1-formulas is again a 1-formula. This allows us to define the *separator* of a configuration  $w$ .

**Definition 5.1.** The *separator* of a configuration  $w$ , denoted  $Sep_w$ , is the union of all inductive 1-sets *not* containing  $w$ .

We characterize membership of  $w$  in  $IndInv_1$  in terms of its separator:

**Lemma 5.2.** For every configuration  $w$ , its separator  $Sep_w$  is an inductive 1-set. Further  $w \in IndInv_1$  iff  $Sep_w$  is not an invariant.

*Proof.* Since inductive sets are closed under union,  $Sep_w$  is inductive. Since the disjunction of two clauses is again a clause, the union of two 1-sets of configurations is also a 1-set, and so  $Sep_w$  is an inductive 1-set. For the last part, we prove that  $w \notin IndInv_1$  iff  $Sep_w$  is an invariant. Assume first  $w \notin IndInv_1$ . Then some inductive 1-invariant does not contain  $w$ . Since, by definition,  $Sep_w$  contains this invariant,  $Sep_w$  is also an invariant. Assume now that  $Sep_w$  is an invariant. Then  $Sep_w$  is an inductive 1-invariant, and so  $Sep_w \supseteq IndInv_1$ . Since  $w \notin Sep_w$ , we get  $w \notin IndInv_1$ . □

Our plan for the rest of the section is as follows:

- We introduce the notion of a separation table for a configuration. (Definition 5.3)
- We show that, given a configuration  $w$  and a separation table  $\tau$  for  $w$ , we can construct a 1-formula  $\varphi_{Sep_w^\tau}$  such that  $L(\varphi_{Sep_w^\tau}) = Sep_w$ . (Lemma 5.4)
- We use this result to define a transducer  $T_{sep}$  over  $\Sigma \times 2^\Sigma$  that accepts a word  $\langle w, \varphi \rangle$  iff  $\varphi = \varphi_{Sep_w^\tau}$ . (Proposition 5.6)
- We use  $T_{sep}$  and Proposition 2.2 to define an NFA over  $\Sigma$  that accepts a configuration  $w$  iff  $Sep_w$  is not an invariant, and so, by Lemma 5.2, iff  $w \in IndInv_1$ . (Theorem 5.11)

We present a characterization of  $Sep_w$  in terms of *tables*. Given a transition  $s \rightsquigarrow t$ , we call  $s$  and  $t$  the *source* and *target* of the transition, respectively. A *table* of length  $\ell$  is a sequence  $\tau = s_1 \rightsquigarrow t_1, \dots, s_n \rightsquigarrow t_n$  of transitions of  $\mathcal{R}$  (not necessarily distinct), all of length  $\ell$ .<sup>3</sup> We define the *separation tables* of a configuration  $w$ .

**Definition 5.3.** Let  $w$  be a configuration and let  $\tau = s_1 \rightsquigarrow t_1, \dots, s_n \rightsquigarrow t_n$  be a table, both of length  $\ell$ . For every  $j \in [1, \ell]$ , let  $In(w, \tau)[j] = \{w[j], s_1[j], \dots, s_n[j]\}$  be the set of letters at position  $j$  of  $w$  and of the source configurations  $s_1, \dots, s_n$  of the table.

- $\tau$  is *consistent with  $w$*  if for every  $i \in [1, n], j \in [1, \ell]$ , either  $t_i[j] = w[j]$  or  $t_i[j] = s_{i'}[j]$  for some  $i' < i$ .

(Intuitively:  $\tau$  is consistent with  $w$  if for every position and every target configuration of the table, the letter of the target at that position is either the letter of  $w$ , or the letter of some earlier source configuration, with the choice for different positions made independently.)

- $\tau$  is *complete for  $w$*  if every table  $\tau, s \rightsquigarrow t$  consistent with  $w$  satisfies  $s[j] \in In(w, \tau)[j]$  for every  $j \in [1, \ell]$ .


(Intuitively:  $\tau$  is complete for  $w$  if it cannot be extended by a transition that maintains consistency and introduces a new letter.)

A table is a *separation table of  $w$*  if it is consistent with and complete for  $w$ .

Overall, a separation table represents the following logic. We want an inductive 1-invariant that does not contain  $w$ . We think of the words that must be outside the invariant. As a 1-set is defined via a disjunction of atomic propositions, effectively the question is which letters are excluded at each position. Naturally, all the letters of  $w$  are excluded for their corresponding positions. Moreover, any transition preimage of a word constructed of excluded letters also has to be an excluded word due to the inductive property. Letters of such a preimage are also excluded at their corresponding positions. Note that this means that we use the transitions backwards: we start with  $w$ , then try to add something that can reach  $w$  in one step, etc. For instance in the running example this manifests as follows: Consider the (reachable) configuration  $n t n n n n$  for our running example. The largest statement that is not satisfied by this configuration is  $\{t\} \{n\} \{t\} \{t\} \{t\} \{t\}$  since it contains at every position all letters but the one that is at the same position in the original configuration. We now demonstrate how this statement is refined to become inductive. For this, we show a series of statements and transitions below such that the transition refines the previous statement to the next. In the following table we mark statements with  $\bullet$  and the refining transitions with  $\triangleright$ . Moreover, we mark in red the atoms of the statements that are removed in each step and in the transitions the reason why they are removed.


<sup>3</sup>We call it a table because we visualize  $s_1, t_1, \dots, s_n, t_n$  as a matrix with  $2n$  rows and  $\ell$  columns.





Consistency with  $w$  means that the table contains the transitions implementing this approach. Completeness means that there are no transitions left to add to exclude more letters. Strictly speaking, we do not require each transition in the table to exclude a new letter, but adding transitions without excluding any new letters is useless.

For our running example a separation table for the configuration  $n t n n n n$  is



Observe that every configuration  $w$  has at least one separation table. If there are no transitions with target  $w$ , then the empty table with no transitions is a separation table. Otherwise, starting with any transition  $s \rightsquigarrow w$ , we repeatedly add a transition, maintaining consistency and introducing at least one new letter until no such transition exists. Lack of such transitions implies completeness of the table. This procedure terminates—there are only finitely many transitions between configurations of a fixed length—and yields a separation table.

The next lemma shows how to compute a 1-formula  $\varphi_{Sep_w^\tau}$  such that  $L(\varphi_{Sep_w^\tau}) = Sep_w$  from any separation table  $\tau$  of  $w$ .

**Lemma 5.4.** *Let  $\tau$  be any separation table for a configuration  $w$  of length  $\ell$ . Then  $Sep_w$  is the set of all configurations  $z \in \Sigma^\ell$  such that  $z[j] \notin In(w, \tau)[j]$  for some  $j \in [1, \ell]$ . In particular,  $Sep_w$  is the language of the 1-formula*

$$\varphi_{Sep_w^\tau} := \bigvee_{j=1}^{\ell} \left( \bigvee_{a \notin In(w, \tau)[j]} a_{j:\ell} \right)$$

or, in the powerword encoding, of the formula

$$\varphi_{Sep_w^\tau} = \overline{In(w, \tau)[1]} \cdots \overline{In(w, \tau)[\ell]} .$$

Before we prove this lemma, observe that the separation table for our running example from before would give  $\varphi_{Sep_{n t n n n n}^\tau}$  as the powerword  $\emptyset \emptyset \emptyset \emptyset \emptyset \emptyset$ : In particular, the first letter is  $\emptyset$  because it does not contain  $n$  as the configuration starts with the letter  $n$  but it also does not contain  $t$  as the first transition “removes” it.

*Proof.* We claim that  $\varphi_{Sep_w^\tau}$  denotes an inductive 1-set not containing  $w$ . That  $\varphi_{Sep_w^\tau}$  denotes a 1-set not containing  $w$  follows immediately from the definition. To see that  $\varphi_{Sep_w^\tau}$  denotes an inductive set, assume there is a transition  $s \rightsquigarrow t$  such that  $s \models \varphi_{Sep_w^\tau}$  and  $t \not\models \varphi_{Sep_w^\tau}$ . Since  $t \not\models \varphi_{Sep_w^\tau}$ , all the letters of  $t$  are excluded at their corresponding positions, thus the transition  $s \rightsquigarrow t$  can be added to the table  $\tau$  yielding the table  $\tau, s \rightsquigarrow t$  consistent with  $w$ . Since  $s \models \varphi_{Sep_w^\tau}$ ,  $\tau$  is not complete, contradicting the assumption.

Let us now prove that  $\varphi_{Sep_w^\tau}$  denotes the *largest* inductive 1-set not containing  $w$ . Using the powerword encoding, it is enough to prove that for every position  $j$  and every letter  $x \in In(w, \tau)[j]$ , each inductive 1-set specified by a powerword containing  $x$  at position  $j$  contains the word  $w$ .

Consider the prefixes of  $\tau$ : the empty sequence  $\tau_0$ , and the sequences  $\tau_1 = s_1 \rightsquigarrow t_1$ ,  $\tau_2 = s_1 \rightsquigarrow t_1, s_2 \rightsquigarrow t_2$ , up to  $\tau_n = \tau = s_1 \rightsquigarrow t_1, \dots, s_n \rightsquigarrow t_n$ . All of them are tables consistent with  $w$ . By construction, for every  $j$ , we have  $In(w, \tau_0)[j] \subseteq In(w, \tau_1)[j] \subseteq \dots \subseteq In(w, \tau_n)[j]$ . We prove by induction on  $k$  that, for each position  $j$  and each letter  $x \in In(w, \tau_k)[j]$ , each inductive 1-set specified by a powerword containing  $x$  at position  $j$  contains the word  $w$ . The base is obvious as  $x \in In(w, \tau_0)[j]$  means  $x = w[j]$ .

To prove the induction step, consider some index  $k$ , position  $j$ , and letter  $x \in In(w, \tau_k)[j] \setminus In(w, \tau_{k-1})$ . By definition of  $In(w, \tau_k)$ , this means that  $x = s_k[j]$  (recall that  $s_k \rightsquigarrow t_k$  is the last transition in  $\tau_k$ , and this transition is not present in  $\tau_{k-1}$ ). Consider any inductive 1-set  $S$  containing the language of  $x_{j:\ell}$ . As  $s_k \in S$  and  $S$  is inductive,  $t_k \in S$  holds. Hence, for some  $j'$ , the 1-set  $S$  contains the language of the atomic proposition  $t_k[j']_{j':\ell}$ . By consistency of  $\tau$ , the letter  $t_k[j']$  is in the set  $In(w, \tau_{k-1})[j]$ . Thus consider  $x' = t_k[j']$  and apply the induction hypothesis to  $x'$  at position  $j'$ . We obtain  $w \in S$ .

We have proven that  $\varphi_{Sep_w^\tau}$  denotes a maximal inductive 1-set not containing  $w$ , i.e.  $Sep_w$ .  $\square$

We construct a transducer over the alphabet  $\Sigma \times 2^\Sigma$  that transduces a configuration  $w$  into the formula  $\varphi_{Sep_w^\tau}$  of a table  $\tau$  consistent with and complete for  $w$ . For this, we need the consistency and completeness summaries of a table.

The basic motivation is: locally, there is only a polynomial-size list of options how a transition can look like, and we can always reuse a fragment of a transition with a different continuation as long as transducer states match. So what we most care about is which transitions of the RTS-transition transducer become available at each position, and in which order.

To certify completeness, we need to check a claim about all transitions accepted by the transition transducer; we apply the subset construction and record its execution. We keep track of two sets of states: reachable at all, and reachable while also adding at least one new excluded letter.

**Definition 5.5.** Let  $\tau = s_1 \rightsquigarrow t_1, \dots, s_n \rightsquigarrow t_n$  be a separation table for a configuration  $w$ . The *consistency summary* is the result of applying the following procedure to  $\tau$ :

- Replace each pair of rows  $s_i \rightsquigarrow t_i$  by a sequence of states providing an accepting run of  $\mathcal{T}$  on it.  
(This produces a table with  $n$  rows and  $\ell + 1$  columns, whose entries are states of  $\mathcal{T}$ .)
- In each column, keep the first occurrence of each state, removing the rest.  
(The result is a sequence of columns; the columns of the sequence may have possibly different lengths.)

The *completeness summary* is the sequence  $(Q_0, Q'_0), (Q_1, Q'_1) \dots (Q_\ell, Q'_\ell)$  of pairs of sets of states of  $\mathcal{T}$ , defined inductively as follows for every  $j \in [0, \ell]$ :

- $Q_0$  is the set of initial states and  $Q'_0$  is empty.
- $Q_{j+1}$  is the set of states reachable from  $Q_j$  by means of letters  $[a, b]$  such that  $b \in \text{In}(w, \tau)$ .
- $Q'_{j+1}$  is the set of states reachable from  $Q'_j$  by means of letters  $[a, b]$  such that  $b \in \text{In}(w, \tau)$ , or reachable from  $Q_j$  by means of letters  $[a, b]$  such that  $a \notin \text{In}(w, \tau)$  and  $b \in \text{In}(w, \tau)$ .

To get to the consistency summary of our running example, we start with the accepting runs of the transitions in the separation table and get

▷	$q_0$	$q_2$	$q_3$	$q_3$	$q_3$	$q_3$	$q_3$
▷	$q_0$	$q_4$	$q_4$	$q_4$	$q_4$	$q_4$	$q_5$
▷	$q_0$	$q_1$	$q_1$	$q_1$	$q_1$	$q_2$	$q_3$
▷	$q_0$	$q_1$	$q_1$	$q_1$	$q_2$	$q_3$	$q_3$
▷	$q_0$	$q_1$	$q_1$	$q_2$	$q_3$	$q_3$	$q_3$

By removing all but the first occurrence of each state in each column we get

▷	$q_0$	$q_2$	$q_3$	$q_3$	$q_3$	$q_3$	$q_3$
▷		$q_4$	$q_4$	$q_4$	$q_4$	$q_4$	$q_5$
▷		$q_1$	$q_1$	$q_1$	$q_1$	$q_2$	
▷					$q_2$		
▷					$q_2$		

Since the statement that arises from this table is  $\emptyset \emptyset \emptyset \emptyset \emptyset \emptyset$  the completeness summary becomes

$$\langle \{q_0\}, \emptyset \rangle \langle \{q_1, q_2, q_4\}, \emptyset \rangle \langle Q'', \emptyset \rangle \langle Q'', \emptyset \rangle \langle Q'', \emptyset \rangle \langle Q'', \emptyset \rangle$$

where  $Q'' = \{q_1, q_2, q_3, q_4, q_5\}$  immediately.

Observe that the consistency summary is a sequence  $\alpha = \alpha[1] \dots \alpha[\ell]$  where  $\alpha[i]$  is a sequence of *distinct* states of  $\mathcal{T}$ , i.e., an element of  $Q_T^{n_T}$ , and the completeness summary is a sequence  $\beta = \beta[1] \dots \beta[\ell]$  where  $\beta[i]$  is a pair of sets of states of  $\mathcal{T}$ , i.e., an element of  $2^{Q_T} \times 2^{Q_T}$ . We prove:

**Proposition 5.6.** *There exists a transducer  $T_{sep}$  over the alphabet  $\Sigma \times 2^\Sigma$  satisfying the following properties:*

- *The states of  $T_{sep}$  are elements of  $(Q_T \cup \{\square\})^{n_T} \times (2^{Q_T} \times 2^{Q_T})$  where  $n_T$  is the number of states of  $\mathcal{T}$ .*
- *There is a polynomial time algorithm that, given two states  $q, q'$  of  $T_{sep}$  and a letter  $\langle a, X \rangle \in \Sigma \times 2^\Sigma$  decides whether the triple  $(q, \langle a, X \rangle, q')$  is a transition of  $T_{sep}$ .*
- *$T_{sep}$  recognizes a word  $\langle w, \varphi \rangle$  over  $\Sigma \times 2^\Sigma$  iff  $\varphi = \varphi_{Sep_w^\tau}$ .*

*Proof.* The proof is long. We introduce *local separation refinements* in Definition 5.7, and give an algorithm for computing them in Lemma 5.8. We then show that they can be combined into increasingly longer *fragment separation refinements* (Definition 5.9 and Lemma 5.10). Equipped with these results, we then prove the Proposition.

**Definition 5.7.** A *local separation refinement* for a letter  $c \in \Sigma$  is a sequence of transitions  $\lambda = (q_1, (a_1, b_1), q'_1), \dots, (q_n, (a_n, b_n), q'_n)$  of  $\mathcal{T}$  such that for every  $i \in [1, n]$ , either  $b_i = c$  or  $b_i = a_{i'}$  for some  $i' < i$ . The states  $q_1, \dots, q_n$  and  $q'_1, \dots, q'_n$  are called *incoming* and *outgoing* states, respectively. The *first-appearance lists* of  $\lambda$  are the result of removing from  $q_1, \dots, q_n$  and  $q'_1, \dots, q'_n$  all elements  $q_i$  ( $q'_i$ ) such that  $q_k = q_i$  ( $q'_k = q'_i$ ) for some  $k < i$ .

Observe that the condition on  $b_i$  corresponds to consistency with the configuration (restricted to a single position).

**Lemma 5.8.** *There is a polynomial time algorithm taking an RTS  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$ , a letter  $c \in \Sigma$ , and two lists of distinct states of  $\mathcal{T}$  as input. The algorithm outputs a local separation refinement for  $c$  whose first-appearance lists are equal to the input lists. The algorithm produces empty output if no such refinement exists.*

*Proof.* The algorithm starts with the empty sequence of transitions, and repeatedly adds transitions until the local separation refinement is constructed. The algorithm keeps track of the transitions, the source letters, and the lists of incoming states and outgoing states of the sequence of transitions constructed so far. We call them the sets of *used* transitions, and source letters, and the lists of *used* incoming states, and outgoing states. Initially all these sets and lists are empty. At each moment in time, a target letter is *permissible* if it is either equal to  $c$ , or it has been already used as a source letter. The main loop of the algorithm repeatedly proceeds as follows. First, it checks whether there is a *still unused* transition with a permissible target letter, and whose incoming and outgoing states have already been used (as incoming and outgoing states, respectively). If so, the transition is added to the current sequence. Otherwise, the algorithm checks if there is an unused transition, whose incoming states and outgoing states have already been used, or are equal to the next states in the input lists. If so, the transition is added to the current sequence. The incoming and the outgoing states are added to the corresponding lists of used states if they are not already present.

If a loop iteration fails to add a new transition, then the loop terminates. If, after termination, the lists of used incoming and outgoing states are equal to the input lists, then the algorithm returns the current sequence of transitions, otherwise it returns nothing.

Each iteration of the algorithm runs in polynomial time. Since each iteration but the last adds one transition to the sequence, the algorithm runs in polynomial time.

We now show that if the algorithm returns a sequence, then that sequence is a local separation refinement for  $c$  satisfying the conditions. By construction, after each iteration the current sequence is a local separation refinement; further, the lists of used incoming and outgoing states are first appearances lists, and prefixes of the input lists. Finally, the algorithm returns a sequence only if the lists of used incoming and outgoing states coincide with the input lists.

It remains to show that if there exists a separation refinement for  $c$  satisfying the conditions, then the algorithm returns a sequence. Assume there is a such a refinement, but the algorithm terminated without returning a sequence. Note that if after termination all transitions of the refinement have been used, then all incoming and outgoing states of the input lists have been used too, and so the algorithm would have returned a sequence. So at least one transition of the refinement, was not used by the algorithm. Consider the first such transition, say  $(q, (a, b), q')$ . Since all the previous transitions of the refinement (and possibly some others) were used, after termination all incoming and outgoing states of the input lists before  $q$  and  $q'$  have been used too, and  $b$  has become permissible. Then during the last iteration of the main loop the algorithm was able to use  $(q, (a, b), q')$ , but terminated instead, a contradiction.

This concludes the proof of the lemma.  $\square$

We obtain separation refinements for a word  $w$  by combining local separation refinements for its letters, as follows.

**Definition 5.9.** A *run fragment* of an automaton is a run of the automaton obtained by replacing the set of initial states with the set of all states. (In other words, the condition of starting in an initial state is dropped.) The states starting and ending a run fragment are called the *incoming* and *outgoing* states of the fragment, respectively.

A *fragment separation refinement*  $\tau$  for a word  $w$  of length  $\ell$  is a table of  $\ell$  columns  $\tau[1], \dots, \tau[\ell]$ , where each column is a local separation refinement of the same length. Moreover, the sequences of outgoing and incoming states of adjacent columns coincide; i.e. for each  $j \in [1, \ell]$  the  $i$ -th outgoing state of the  $j$ -th column is equal to the  $i$ -th incoming state of the  $j + 1$ -th column. The incoming and outgoing states of  $\tau$  are the incoming states of  $\tau[1]$  and the outgoing states of  $\tau[\ell]$ , respectively.

We can also split a fragment separation refinement  $\tau$  into rows  $\tau_1, \dots, \tau_n$ . Each row  $\tau_i$  is a run fragment of  $\mathcal{T}$ . We prove that fragment separation refinements can be concatenated.

**Lemma 5.10.** *Let  $\tau_w$  and  $\tau_z$  be fragment separation refinements for words  $w$  and  $z$  such that the first-appearance list of outgoing states of  $\tau_w$  and the first-appearance list of incoming states of  $\tau_z$  coincide. There exists a fragment separation refinement  $\tau_{wz}$  for  $wz$  with the same sets of used source letters (and therefore the same sets of permissible letters) at the corresponding positions. Further,  $\tau_{wz}$  has the same first-appearance list of incoming states as  $\tau_w$ , and the same first-appearances list of outgoing states as  $\tau_z$ .*

*Proof.* We construct  $\tau_{wz}$  row by row. For this we repeatedly choose a row of  $\tau_w$  and a row of  $\tau_z$  such that the outgoing state of  $\tau_w$  and the incoming state of  $\tau_z$  coincide, concatenate them, and add the result to  $\tau_{wz}$ . Rows can be chosen multiple times.

We show how to choose the rows of  $\tau_w$  and  $\tau_z$  so that the final result is a fragment separation refinement of  $\tau_{wz}$  satisfying the conditions of the lemma. It suffices to show how to choose the rows so that they satisfy the following two constraints:

- (1) A row of  $\tau_w$  or  $\tau_z$  can be chosen only if each previous row has been chosen at least once. (Note that once a row is chosen once, it can be chosen again without restrictions.)
- (2) All rows of  $\tau_w$  and  $\tau_z$  are eventually chosen.

In particular, this ensures that  $\tau_{wz}$  has the same first-appearance list of incoming states as  $\tau_w$ , and the same first-appearances list of outgoing states as  $\tau_z$ .

We proceed as follows. At each step, we consider all triples  $(q, A, B)$  where  $A$  is a row of  $\tau_w$  with outgoing state  $q$  and  $B$  is a row of  $\tau_z$  with incoming state  $q$ . (We also formally add a triple  $(\infty, \infty, \infty)$  where each component has to be used last.) We choose any triple such that the addition of  $AB$  to  $\tau_{wz}$  respects the order in the first-appearance lists of incoming states of  $\tau_w$  and outgoing states of  $\tau_z$ . This guarantees that the resulting table satisfies condition 1. It suffices to show that all triples will be used.

Assume the contrary. Let  $\sigma$  be the common list of first-appearances of outgoing states of  $\tau_w$  and incoming states of  $\tau_z$ . Consider the earliest non-chosen rows  $A$  and  $B$  (it is possible that one of them is  $\infty$ , but if  $\infty$  is used, then every row has been used in both refinements). Let  $q$  be the outgoing state of  $A$  and  $r$  be the incoming state of  $B$ . We have  $q \neq r$ , because otherwise the procedure can choose  $(q = r, A, B)$ . Without loss of generality, assume that  $q$  precedes  $r$  in  $\sigma$ . Then  $\tau_z$  has a row  $C$  before  $B$  with incoming state  $q$ . Since  $B$  is the earliest non-chosen row,  $C$  can be chosen. So the triple  $(q, A, C)$  can be chosen, contradicting that  $A$  is the earliest non-chosen row. This concludes the proof of the lemma.  $\square$

We now proceed to prove Proposition 5.6. Let us first construct the transducer  $T_{sep}$ . Recall that a state of the transducer is a pair  $q = (\alpha, \beta)$  where  $\alpha \in (Q_T \cup \{\square\})^{n_T}$ ,  $\beta \in (2^{Q_T} \times 2^{Q_T})$ ,

and  $n_T$  is the number of states of  $\mathcal{T}$ . In every reachable state of  $T_{sep}$ ,  $\alpha$  will be a permutation of a subset of  $Q_T$  followed by some number of  $\square$  symbols. We look at these permutations as first-appearance lists, and so we call  $\alpha$  the first-appearance list of  $q$ . The second component  $\beta$  is the completeness summary letter of  $q$  (see Definition 5.5).

A state  $q = (\alpha, \beta)$  is initial iff

- all the states of  $Q_T$  appearing in  $\alpha$  are initial; and
- $\beta = (Q_{T,0}, \emptyset)$  where  $Q_{T,0}$  are the initial states of  $\mathcal{T}$ .

and final iff

- all the states of  $Q_T$  appearing in  $\alpha$  are final; and
- the second set of  $\beta$  contains no final states.

Let us now define the algorithm recognizing the transitions of  $T_{sep}$ . Given two states  $q = (\alpha, \beta), q' = (\alpha', \beta')$  and a letter  $c$ , the algorithm either outputs a set of letters  $X$ , meaning that  $(q, \langle a, X \rangle, q')$  is a transition, or rejects. The algorithm takes  $\alpha, \alpha'$  and  $c$ , and uses the algorithm from the Lemma 5.8 to construct a local separation refinement for  $c$ , if any exists. If no such refinement exists, the algorithm rejects. Otherwise, the algorithm assigns to  $X$  the complement of the set of permissible letters of this refinement, and checks whether  $\beta, \beta'$  satisfy the definition of a completeness summary (Definition 5.5) with respect to  $X$ . If they do, then the set  $X$  is returned. Clearly, the algorithm runs in polynomial time.

It remains to show that  $T_{sep}$  accepts a word  $\langle w, \varphi \rangle$  over  $\Sigma \times 2^\Sigma$  iff  $\varphi = \varphi_{Sep_w^\tau}$ .

Assume  $\varphi = \varphi_{Sep_w^\tau}$ . Let  $\tau$  be a separation table for  $w$ . By Lemma 5.4, we have  $\varphi = \overline{In(w, \tau)[1]} \cdots \overline{In(w, \tau)[\ell]}$ . By the definition of  $T_{sep}$ , the transducer has a run on  $\langle w, \varphi \rangle$  whose sequence of visited states are the consistency and completeness summaries of  $\tau$  for  $w$ , and the run is accepting.

Assume now that  $T_{sep}$  has an accepting run on  $\langle w, \varphi \rangle$ . We have to show that  $\varphi = \varphi_{Sep_w^\tau}$ . It suffices to show that from the run we can construct a separation table for  $w$ , since all separation tables produce the same formula defining  $Sep_w$ . The procedure to construct a separation table from an accepting run goes as follows<sup>4</sup>. First, the procedure constructs for each transition in the run a local separation refinement, applying Lemma 5.8. Then it repeatedly applies Lemma 5.10 to produce one fragment separation refinement for the complete run. Since, by the definition of  $T_{sep}$ , all states in the initial (resp. final) first-appearance list are initial (resp. final), the rows of the joint fragment separation refinement are accepting runs of  $\mathcal{T}$ , and so transitions of the RTS. Thus we obtain a sequence of transitions, i.e. a table. The table is consistent with the word  $w$ , because this is a local property ensured by local separation refinements and preserved during the fragment merging. To show that the table is complete for  $w$ , observe that, in the completeness summary, for each pair  $(Q_j, Q'_j)$ ,  $Q_j$  is the set of all the states reachable by reading a sequence of pairs of letters with target letter permissible for their positions, and  $Q'_j$  is the subset where at least one of the used source letters read in the process was not permissible. By the definition of the final states of  $T_{sep}$ , no states of  $Q'_\ell$  are final states of  $\mathcal{T}$ . Therefore, no transition of the regular transition system can add a new permissible letter, which is the completeness condition.  $\square$

This concludes the proof of Proposition 5.6. We now prove our main technical result:

<sup>4</sup>Note that this procedure is an inefficient proof of existence, there is no need to run it as a part of the algorithm.

**Theorem 5.11.** *Let  $\mathcal{R} = \langle \Sigma, \mathcal{I}, \mathcal{T} \rangle$  be an RTS. There exists an NFA  $\mathcal{A}_1$  over  $\Sigma$  satisfying the following properties:*

- *The states of  $\mathcal{A}_1$  are elements of  $(Q_T \cup \{\square\})^{n_T} \times (2^{Q_T} \times 2^{Q_T}) \times \{0, 1\} \times Q_I$ .*
- *There is a polynomial time algorithm that, given two states  $q, q'$  of  $\mathcal{A}_1$  and a letter  $a \in \Sigma$  decides whether the triple  $(q, a, q')$  is a transition of  $\mathcal{A}_1$ .*
- *$L(\mathcal{A}_1) = \text{IndInv}_1$*

*Proof.* Let  $T_{sep}$  be the transducer over the alphabet  $\Sigma \times 2^\Sigma$  of Proposition 5.6. By Lemma 5.2,  $w \in \text{IndInv}_1$  iff there exists a 1-formula  $\varphi$  such that  $\langle w, \varphi \rangle \in L(T_{sep})$  and  $\varphi$  is not an invariant, that is, there exists  $u \in L(\mathcal{I})$  such that  $u \not\models \varphi$ . By Lemma 4.1, there is a deterministic transducer  $\bar{\mathcal{V}}$  with two states, say  $\{0, 1\}$ , recognizing the pairs  $\langle \varphi, u \rangle$  such that  $u \not\models \varphi$ . So we get

$$\begin{aligned} \text{IndInv}_1 &= \{w \in \Sigma^* \mid \exists \varphi \in (2^\Sigma)^*, u \in \Sigma^* \text{ s.t. } \langle w, \varphi \rangle \in R(T_{sep}), \langle \varphi, u \rangle \in L(\bar{\mathcal{V}}), u \in L(\mathcal{I})\} \\ &= (R(T_{sep}) \circ L(\bar{\mathcal{V}}) \circ L(\mathcal{I}))|_1. \end{aligned}$$

The automaton  $\mathcal{A}_1$  is obtained from Proposition 2.2. □

Observe that a state of  $\mathcal{A}_1$  can be stored using space linear in  $\mathcal{I}$  and  $\mathcal{T}$ . This yields:

**Corollary 5.12.** *Deciding  $\text{IndInv}_1 \cap L(\mathcal{U}) = \emptyset$  is in PSPACE.*

*Proof.* Guess a configuration  $w$  and an accepting run of  $\mathcal{A}_1$  and  $\mathcal{U}$  on  $w$ , step by step. By Proposition 5.6, this can be done in polynomial space. Apply then  $\text{NPSpace} = \text{PSPACE}$ . □

## 6. DECIDING $\text{IndInv}_1 \cap \mathcal{L}(\mathcal{U}) = \emptyset$ IS PSPACE-HARD.

This section presents the proof of the following lemma:

**Lemma 6.1.** *Given an RTS  $\mathcal{R}$  and an NFA  $\mathcal{U}$ , deciding  $\text{IndInv}_1 \cap \mathcal{L}(\mathcal{U}) = \emptyset$  is PSPACE-hard.*

We reduce from the problem of deciding whether a bounded Turing machine of size  $n$  that can only use  $n$  tape cells accepts when started on the empty tape. This problem (very similar to the acceptance problem for linearly bounded automata) is known to be PSPACE-complete. Given such a machine, we construct a deterministic RTS  $\mathcal{R}$  which, loosely speaking, satisfies the following properties: 1) the execution of  $\mathcal{R}$  from an initial configuration of length  $\Theta(t \cdot n)$  simulates the first  $t$  steps of the computation of the Turing machine, and 2)  $\text{IndInv}_1$  coincides with the set of reachable configurations.

The proof is divided in several parts. We first introduce some notations on Turing machines. Then, we define the RTS  $\mathcal{R}$ , first informally and then formally. Finally, we conduct the reduction.

**Turing machines.** We fix some notations on Turing machines. A Turing machine  $\mathcal{M}$  consists of

- a set of states  $Q$ ,
- an initial state  $q_0 \in Q$  and a final state  $q_f \in Q$ ,
- an input alphabet  $\Sigma$  and a tape alphabet  $\Gamma \supseteq \Sigma$ ,
- a dedicated blank symbol  $B \in \Gamma \setminus \Sigma$ , and
- a transition function  $\delta: Q \times \Gamma \rightarrow Q \times \Gamma \times \{L, R\}$ .

Let  $\mathcal{M}$  be any deterministic linearly bounded Turing machine  $\mathcal{M}$ , meaning that  $\mathcal{M}$  only uses  $|\mathcal{M}|$  tape cells. We construct an instance  $\mathcal{R}, \mathcal{U}$  of the safety verification problem, of size  $O(n)$ , such that  $\mathcal{M}$  accepts the empty word if and only if  $\text{IndInv}_1 \cap L(\mathcal{U}) \neq \emptyset$ .

Let  $n$  denote the size of  $\mathcal{M}$  (and so the number of tape cells) plus 1. We represent a configuration of  $\mathcal{M}$  as a word  $\alpha \in (\Gamma \cup Q)^n$  containing exactly one letter in  $Q$ . That is, by definition  $\alpha = \beta q \eta$  where  $\beta \in \Gamma^*$ ,  $q \in Q$ , and  $\eta \in \Gamma^+$ . If  $\alpha$  is the current configuration of  $\mathcal{M}$ , then  $\mathcal{M}$  is in the state  $q$ , the content of the tape is  $\beta \eta$ , and the head of  $\mathcal{M}$  reads the first letter of  $\eta$ . Observe that  $\beta \eta$  has length  $|\mathcal{M}|$ . The initial configuration  $\alpha_0$  is  $q_0 B^{n-1}$ , that is, the tape is initially empty.

Since  $\mathcal{M}$  is deterministic, there is *exactly* one sequence of configurations  $\alpha_0 \vdash \alpha_1 \vdash \alpha_2 \vdash \dots$  where  $\vdash$  is used to describe that  $\alpha_{i+1}$  is the successor configuration of  $\alpha_i$ . For simplicity, we also allow  $\vdash$  to connect two identical configurations if their state is the final  $q_f$ . Consequently, this sequence of configurations is infinite; either because  $\mathcal{M}$  loops or because  $\mathcal{M}$  “stutters” in a final configuration.

**Description of the RTS.** We construct an RTS that simulates the execution of  $\mathcal{M}$  on the empty word. The alphabet of the RTS consists of  $\Gamma \cup (\Gamma \times Q)$ , and two auxiliary symbols  $\bullet$  and  $\square$ . Using these symbols, the execution of the Turing machine can be encoded as the infinite word

$$\bullet \alpha_0 \bullet \alpha_1 \bullet \alpha_2 \bullet \dots$$

The symbol  $\bullet$  separates the individual configurations of the Turing machine. The set of initial configurations of the RTS contains for every  $m \geq 0$  the configuration

$$\alpha_0 \bullet \underbrace{\square^n \bullet \square^n \bullet \dots \bullet \square^n}_{m \text{ times}}.$$

Intuitively, an initial configuration of the RTS consist of a sequence of  $m + 1$  “pages”. Each page has space to “write” a configuration of the Turing machine. The first page is already filled with the initial configuration of the Turing machine, the others are still “empty”. The transitions of the RTS repeatedly replace the  $\square$ -symbols, from left to right, by the correct letters of the successor configurations  $\alpha_1, \alpha_2, \dots, \alpha_m$  of the Turing machine. That is, repeated application of the transducer eventually reaches the configurations

$$\begin{aligned} &\bullet \alpha_0 \bullet \square^n \bullet \square^n \bullet \dots \bullet \square^n \\ &\bullet \alpha_0 \bullet \alpha_1 \bullet \square^n \bullet \dots \bullet \square^n \\ &\bullet \alpha_0 \bullet \alpha_1 \bullet \alpha_2 \bullet \dots \bullet \square^n \\ &\dots \\ &\bullet \alpha_0 \bullet \alpha_1 \bullet \alpha_2 \bullet \dots \bullet \alpha_m \end{aligned}$$

(with other configurations in between). Note that the symbol at the position  $i + n + 1$  of a word in this sequence is determined by the symbols at positions  $i - 1, i, i + 1$  of the previous word.



In particular, one can define a (partial) function  $\Delta: (\{\bullet\} \cup \Gamma \cup (Q \times \Gamma))^3 \rightarrow \Gamma \cup (Q \times \Gamma)$  which gives the letter at position  $i + n + 1$ , given the values of the symbols at positions  $i - 1, i, i + 1$ . (The function is partial because there are inputs that do not represent a sensible situation; e.g., inputs which contain more than one element of  $Q \times \Gamma$ .)

The transducer of the RTS non-deterministically guesses a position  $i$ , stores the elements at positions  $i - 1, i, i + 1$ , and counts to the position  $i + n + 1$  where it changes  $\square$  to the value dictated by  $\Delta$ .<sup>5</sup> This transducer can be realized with polynomially many states with respect to  $\mathcal{M}$ . By construction, every reachable configuration of the RTS that does not contain any  $\square$ -symbol is a prefix of  $\bullet \alpha_0 \bullet \alpha_1 \bullet \alpha_2 \bullet \alpha_3 \bullet \dots$ . Therefore,  $\mathcal{M}$  accepts the empty word iff the RTS reaches any configuration containing some occurrence of  $q_f$  and no occurrence of  $\square$ .

**Characterizing  $IndInv_1$ .** We prove that a configuration of the RTS with no occurrence of  $\square$  is reachable iff it belongs to  $IndInv_1$ . Therefore,  $\mathcal{M}$  accepts the empty word iff  $IndInv_1$  contains a configuration with no occurrence of  $\square$  and some occurrence of  $q_f$ .

By the definition of the RTS, it suffices to show that every configuration  $w \in IndInv_1$  is of the form

$$\bullet \alpha_0 \bullet \dots \bullet \alpha_{i-1} \bullet \beta \square^k \bullet \square^n \bullet \dots \bullet \square^n \quad (6.1)$$

for some  $i \geq 1$  and  $k \leq n$  where  $\beta$  is the prefix of  $\alpha_i$  of length  $n - k$ . Intuitively, these are the configurations reached by the RTS during the process of “writing down” the execution of the Turing machine from the corresponding initial configuration of the RTS by “filling the  $\square$ s”.

The proof hinges on the introduction of a few formulas from  $IndInv_1$ . To this end, we first observe that every transition of the RTS only changes  $\square$  symbols to other symbols. Therefore, one can immediately see that, for example, the formula that consists of the single atomic proposition  $B_{3:\ell}$  is an inductive 1-invariant, and so an element of  $IndInv_1$  for every  $\ell$ . The reason is that every configuration satisfying  $B_{3:\ell}$  has  $B$  as its third letter, and no transition of the RTS can change it. The same applies for all formulas of the form  $a_{k:\ell}$  such that  $a \neq \square$ , and so, in particular, for all the formulas

$$\bullet_{1:\ell} q_{02:\ell} B_{3:\ell} \dots B_{n+1:\ell} \bullet_{n+2:\ell}$$

So every configuration of  $IndInv_1$  has  $\bullet \alpha_0 \bullet$  as a prefix (note that  $\alpha_0 = q_0 B^{n-1}$ ).

Recall the function  $\Delta$  defined above, and assume  $\Delta(abc) = d$  where  $a, b, c, d \in \{\bullet\} \cup \Gamma \cup (Q \times \Gamma)$ . We claim that the formula

$$\varphi = a_{i-1:\ell} \wedge b_{i:\ell} \wedge c_{i+1:\ell} \rightarrow (d_{i+n+1:\ell} \vee \square_{i+n+1:\ell})$$

is an inductive 1-invariant for all  $i > 1$ . Observe first that  $\varphi$  has indeed one single clause (after applying standard equivalences). Further, since the transducer can only change one single  $\square$  to a letter different from  $\square$ , the formula  $\varphi$  is also inductive. Finally, an inspection of the initial configurations of the RTS shows that they all satisfy  $\varphi$ , which proves the claim. Let  $w$  be any configuration of a length  $\ell$  satisfying the left-hand-side of  $\varphi$ . The only transition that is applicable to  $w$  changes the  $i + n + 1$ -th letter from  $\square$  to  $d_{i+n+1} : \ell$ . Therefore, in all configurations of  $IndInv_1$  of a length at least  $i + n + 1$ , the letter at position  $i + n + 1$  is either  $\square$  or  $d$ . It follows that every configuration with a prefix that does not

<sup>5</sup>One can eliminate this non-deterministic guess by marking the position that is  $n + 1$  steps before the first  $\square$  and moving this marker further one step in each transition. The remaining arguments work analogously. We choose to avoid this improvement for readability.

contain  $\square$  in  $IndInv_1$  is of the form (6.1). Using this and fixing  $\mathcal{U}$  to accept configurations without any  $\square$  but at least one occurrence of  $q_f$  ends the proof.

## 7. HOW LARGE MUST THE BOUND $b$ BE?

The index  $b$  needed to prove a property (i.e., the least  $b$  such that  $IndInv_b$  implies the property) can be seen as a measure of how difficult it is for a human to understand the proof. We use the experimental setup of [BEI<sup>+</sup>20, ERW21a, BIS21, ERW21b], where systems are encoded as WS1S formulas and MONA [EKM98] is used as a computation engine, to show that  $b = 1$  is enough for a substantial number of benchmarks used in the RMC literature. Note that our goal is to evaluate the complexity of invariants needed for systems from diverse domains, not to present a tool ready to verify complex systems.

Our set of benchmarks consists of problems studied in [CHLR17, ADHR07, BEI<sup>+</sup>20, ERW21a, BIS21, ERW21b]. In a first step, we use MONA to construct a minimal DFA for  $IndInv_1$ . For this, we write a WS1S formula  $\Psi_1(w)$  expressing that, for every 1-formula  $\varphi$ , if  $\varphi$  is an inductive invariant, then  $w$  satisfies  $\varphi$ . MONA yields a minimal DFA for the configurations  $w$  satisfying  $\Psi_1$ , which is precisely  $IndInv_1$ . We then construct the formula  $\Psi_1(w) \wedge Unsafe(w)$ , and use MONA to check if it is satisfiable<sup>6</sup>. All files containing the MONA formulas and the results are provided in [ERW22a]. The results are shown in Figure 1. The first column gives the name of the example. In the second and third column, we give the number of states of the minimal DFA for  $\mathcal{L}_I$  and  $\mathcal{L}_T$ , which we also compute via MONA. In the next column, we give the size of the minimal DFA for  $IndInv_1$ . The fifth column reports whether a property is implied by  $IndInv_1$  (indicated by  $\checkmark$ ) or not (indicated by  $\times$ ). For the cache coherence protocols, we replace  $\checkmark$  with  $k/m$  to state that  $k$  of  $m$  custom safety properties can be established. The last column gives the total running time of MONA<sup>7</sup>. As we can see,  $IndInv_1$  is strong enough to satisfy 46 out of 59 properties.

In the second step, we have studied some of the cases in which  $IndInv_1$  is not strong enough. A direct computation of the automaton for  $IndInv_2$  from a formula  $\Psi_2(w)$  using MONA fails. (A computation based on the automata construction of Section 3 might yield better results and will be part of our future work.) Using a combination of the automatic invariant computation method of [BEI<sup>+</sup>20, ERW21a, BIS21, ERW21b] and manual inspection of the returned invariants, we can report some results for some examples.

*Examples for  $IndInv_b$  with  $b > 1$ .* Figure 1 contains two versions of the dining philosophers in which philosophers take one fork at a time. All philosophers but one are right-handed, i.e., take their right fork first, and the remaining philosopher is left-handed. If the forks “know” which philosopher has taken them (i.e., if they have states  $b_\ell$  and  $b_r$  indicating that the left or the right philosopher has the fork), then deadlock-freedom can be proved using  $IndInv_1$ . If the states of the forks are just “free” and “busy”, then proving deadlock-freedom requires  $IndInv_3$ , and in fact  $Reach = IndInv_3$  holds. We show how to establish this using the technique of [ERW21b] and some additional reasoning in Appendix A.

The Berkeley and Dragon cache coherence protocols are considered as parameterized systems in [Del03]. For both examples,  $IndInv_1$  is too coarse to establish all desired

<sup>6</sup>The second formula  $\Psi_1(w) \wedge Unsafe(w)$  being unsatisfiable suffices for verification purposes, but we use  $\Psi_1(w)$  to obtain information on the size of the minimal DFA.

<sup>7</sup>As reported by MONA.

Figure 1: Experimental results of using  $IndInv_1$  as abstraction of the set of reachable configurations.

System	$ \mathcal{L}_I $	$ \mathcal{L}_T $	$ IndInv_1 $	Properties	time (ms)
Bakery	5	9	8	deadlock mutual exclusion	✓ ✓ < 1
Burns	5	9	6	deadlock mutual exclusion	✓ ✓ < 1
Dijkstra	6	24	22	deadlock mutual exclusion	✓ ✓ 1920
Dijkstra (ring)	6	17	17	deadlock mutual exclusion	✓ × 2
D. cryptographers	6	69	11	one cryptographer paid no cryptographer paid	✓ ✓ 5
Herman, linear	6	7	6	deadlock at least one token	× ✓ < 1
Herman, ring	6	7	7	deadlock at least one token	✓ ✓ < 1
Israeli-Jafon	6	21	7	deadlock at least one token	✓ ✓ < 1
Token passing	6	7	7	at most one token	✓ < 1
Lehmann-Rabin	5	15	13	deadlock	✓ 1
LR phils.	6	14	15	deadlock	× 2
LR phils.(with $b_\ell$ and $b_r$ )	5	14	9	deadlock	✓ 1
Atomic D. phil.	5	12	20	deadlock	✓ 5
Mux array	6	7	8	deadlock mutual exclusion	✓ × < 1
Res. allocator	5	9	8	deadlock mutual exclusion	✓ × < 1
Berkeley	5	19	9	deadlock custom properties	✓ 2/3 1
Dragon	5	26	11	deadlock custom properties	✓ 6/7 3
Firefly	5	18	7	deadlock custom properties	✓ 0/4 1
Illinois	5	25	14	deadlock custom properties	✓ 0/2 1
MESI	5	13	7	deadlock custom properties	✓ 2/2 < 1
MOESI	5	13	10	deadlock custom properties	✓ 7/7 1
Synapse	5	16	7	deadlock custom properties	✓ 2/2 1

consistency assertions. In Appendix B, we describe the formalization of both examples and show that  $IndInv_2$  suffices to obtain the missing assertions.

## 8. CONCLUSION

We have introduced a regular model checking paradigm that approaches the set of reachable configurations from above. As already observed in [ADHR07, CHLR17], such an approach does not require widening or acceleration techniques, as is the case when approaching from

below. The main novelty with respect to [ADHR07, CHLR17] is the discovery of a natural sequence of regular invariants converging to the set of reachable configurations.

Our new paradigm raises several questions. The first one is the exact computational complexity of checking emptiness of the intersection  $IndInv_b$  and the unsafe configurations. We have shown PSPACE-completeness for  $b = 1$ , and we conjecture that the problem is already EXPSPACE-complete for all  $b \geq 2$ . We also think that the CEGAR techniques used in [ERW21b, ERW21a] can be extended to the RMC setting, allowing one to compute intermediate regular invariants between  $IndInv_b$  and  $IndInv_{b+1}$ . Another interesting research venue is the combination with acceleration or widening techniques, and the application of learning algorithms, like the one of [CHLR17]. Currently these techniques try to compute some inductive regular invariant, or perhaps one described by small automata, which may lead to invariants difficult to interpret by humans. A better approach might be to stratify the search, looking first for invariants for small values of  $b$ .

#### ACKNOWLEDGMENTS

We are grateful to Ahmed Bouajjani for fruitful discussions. We thank the anonymous reviewers of this and the earlier versions of the paper for their valuable suggestions on presentation.

#### REFERENCES

- [Abd12] Parosh Aziz Abdulla. Regular model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):109–118, 2012. doi:10.1007/S10009-011-0216-8.
- [Abd21] Parosh Aziz Abdulla. Regular model checking: Evolution and perspectives. In Ernst-Rüdiger Olderog, Bernhard Steffen, and Wang Yi, editors, *Model Checking, Synthesis, and Learning - Essays Dedicated to Bengt Jonsson on The Occasion of His 60th Birthday*, volume 13030 of *Lecture Notes in Computer Science*, pages 78–96. Springer, 2021. doi:10.1007/978-3-030-91384-7\_5.
- [ADHR07] Parosh Aziz Abdulla, Giorgio Delzanno, Noomene Ben Henda, and Ahmed Rezine. Regular model checking without transducers (On efficient verification of parameterized systems). In Orna Grumberg and Michael Huth, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4424 of *Lecture Notes in Computer Science*, pages 721–736. Springer, 2007. doi:10.1007/978-3-540-71209-1\_56.
- [AJNd02] Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Julien d’Orso. Regular model checking made simple and efficient. In Lubos Brim, Petr Jancar, Mojmir Kretínský, and Antonín Kucera, editors, *CONCUR 2002 - Concurrency Theory, 13th International Conference, Brno, Czech Republic, August 20-23, 2002, Proceedings*, volume 2421 of *Lecture Notes in Computer Science*, pages 116–130. Springer, 2002. doi:10.1007/3-540-45694-5\_9.
- [AJNS04] Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Mayank Saksena. A survey of regular model checking. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 35–48. Springer, 2004. doi:10.1007/978-3-540-28644-8\_3.
- [AST18] Parosh Aziz Abdulla, A. Prasad Sistla, and Muralidhar Talupur. Model checking parameterized systems. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 685–725. Springer, 2018. doi:10.1007/978-3-319-10575-8\_21.

- [BEI<sup>+</sup>20] Marius Bozga, Javier Esparza, Radu Iosif, Joseph Sifakis, and Christoph Welzel. Structural invariants for the verification of systems with parameterized architectures. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 228–246. Springer, 2020. doi:10.1007/978-3-030-45190-5\_13.
- [BHRV12] Ahmed Bouajjani, Peter Habermehl, Adam Rogalewicz, and Tomas Vojnar. Abstract regular (tree) model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):167–191, 2012. doi:10.1007/S10009-011-0205-Y.
- [BHV04] Ahmed Bouajjani, Peter Habermehl, and Tomas Vojnar. Abstract regular model checking. In Rajeev Alur and Doron A. Peled, editors, *Computer Aided Verification, 16th International Conference, CAV 2004, Boston, MA, USA, July 13-17, 2004, Proceedings*, volume 3114 of *Lecture Notes in Computer Science*, pages 372–386. Springer, 2004. doi:10.1007/978-3-540-27813-9\_29.
- [BIS21] Marius Bozga, Radu Iosif, and Joseph Sifakis. Checking deadlock-freedom of parametric component-based systems. *J. Log. Algebraic Methods Program.*, 119:100621, 2021. doi:10.1016/J.JLAMP.2020.100621.
- [BJNT00] Ahmed Bouajjani, Bengt Jonsson, Marcus Nilsson, and Tayssir Touili. Regular model checking. In E. Allen Emerson and A. Prasad Sistla, editors, *Computer Aided Verification, 12th International Conference, CAV 2000, Chicago, IL, USA, July 15-19, 2000, Proceedings*, volume 1855 of *Lecture Notes in Computer Science*, pages 403–418. Springer, 2000. doi:10.1007/10722167\_31.
- [BLW03] Bernard Boigelot, Axel Legay, and Pierre Wolper. Iterating transducers in the large (extended abstract). In Warren A. Hunt Jr. and Fabio Somenzi, editors, *Computer Aided Verification, 15th International Conference, CAV 2003, Boulder, CO, USA, July 8-12, 2003, Proceedings*, volume 2725 of *Lecture Notes in Computer Science*, pages 223–235. Springer, 2003. doi:10.1007/978-3-540-45069-6\_24.
- [BT12] Ahmed Bouajjani and Tayssir Touili. Widening techniques for regular tree model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):145–165, 2012. doi:10.1007/S10009-011-0208-8.
- [CHLR17] Yu-Fang Chen, Chih-Duo Hong, Anthony W. Lin, and Philipp Rummer. Learning to prove safety over parameterised concurrent systems. In Daryl Stewart and Georg Weissenbacher, editors, *2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, October 2-6, 2017*, pages 76–83. IEEE, 2017. doi:10.23919/FMCAD.2017.8102244.
- [Del03] Giorgio Delzanno. Constraint-based verification of parameterized cache coherence protocols. *Formal Methods Syst. Des.*, 23(3):257–301, 2003. doi:10.1023/A:1026276129010.
- [DLS01] Dennis Dams, Yassine Lakhnech, and Martin Steffen. Iterating transducers. In Gerard Berry, Hubert Comon, and Alain Finkel, editors, *Computer Aided Verification, 13th International Conference, CAV 2001, Paris, France, July 18-22, 2001, Proceedings*, volume 2102 of *Lecture Notes in Computer Science*, pages 286–297. Springer, 2001. doi:10.1007/3-540-44585-4\_27.
- [EKM98] Jacob Elgaard, Nils Klarlund, and Anders Moller. MONA 1.x: New techniques for WS1S and WS2S. In Alan J. Hu and Moshe Y. Vardi, editors, *Computer Aided Verification, 10th International Conference, CAV ’98, Vancouver, BC, Canada, June 28 - July 2, 1998, Proceedings*, volume 1427 of *Lecture Notes in Computer Science*, pages 516–520. Springer, 1998. doi:10.1007/BFB0028773.
- [ERW21a] Javier Esparza, Mikhail A. Raskin, and Christoph Welzel. Abduction of trap invariants in parameterized systems. In *GandALF*, volume 346 of *EPTCS*, pages 1–17, 2021.
- [ERW21b] Javier Esparza, Mikhail A. Raskin, and Christoph Welzel. Computing parameterized invariants of parameterized Petri nets. In Didier Buchs and Josep Carmona, editors, *Application and Theory of Petri Nets and Concurrency - 42nd International Conference, PETRI NETS 2021, Virtual Event, June 23-25, 2021, Proceedings*, volume 12734 of *Lecture Notes in Computer Science*, pages 141–163. Springer, 2021. doi:10.1007/978-3-030-76983-3\_8.
- [ERW22a] Javier Esparza, Mikhail Raskin, and Christoph Welzel. Repository of examples. <https://doi.org/10.5281/zenodo.6483615>, April 2022. doi:10.5281/zenodo.6483615.
- [ERW22b] Javier Esparza, Mikhail A. Raskin, and Christoph Welzel. Regular model checking upside-down: An invariant-based approach. In Bartek Klin, Slawomir Lasota, and Anca Muscholl, editors, *33rd International Conference on Concurrency Theory, CONCUR 2022, September 12-16, 2022*,

- Warsaw, Poland, volume 243 of *LIPICs*, pages 23:1–23:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.CONCUR.2022.23.
- [JN00] Bengt Jonsson and Marcus Nilsson. Transitive closures of regular relations for verifying infinite-state systems. In Susanne Graf and Michael I. Schwartzbach, editors, *Tools and Algorithms for Construction and Analysis of Systems, 6th International Conference, TACAS 2000, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS 2000, Berlin, Germany, March 25 - April 2, 2000, Proceedings*, volume 1785 of *Lecture Notes in Computer Science*, pages 220–234. Springer, 2000. doi:10.1007/3-540-46419-0\_16.
- [Leg12] Axel Legay. Extrapolating (omega-)regular model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):119–143, 2012. doi:10.1007/S10009-011-0209-7.
- [MP20] Kenneth L. McMillan and Oded Padon. Ivy: A multi-modal verification tool for distributed algorithms. In Shuvendu K. Lahiri and Chao Wang, editors, *Computer Aided Verification, 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020, Proceedings, Part II*, volume 12225 of *Lecture Notes in Computer Science*, pages 190–202. Springer, 2020. doi:10.1007/978-3-030-53291-8\_12.
- [PMP<sup>+</sup>16] Oded Padon, Kenneth L. McMillan, Aurojit Panda, Mooly Sagiv, and Sharon Shoham. Ivy: safety verification by interactive generalization. In Chandra Krintz and Emery D. Berger, editors, *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, pages 614–630. ACM, 2016. doi:10.1145/2908080.2908118.
- [PRZ01] Amir Pnueli, Sitvanit Ruah, and Lenore D. Zuck. Automatic deductive verification with invisible invariants. In Tiziana Margaria and Wang Yi, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001 Genova, Italy, April 2-6, 2001, Proceedings*, volume 2031 of *Lecture Notes in Computer Science*, pages 82–97. Springer, 2001. doi:10.1007/3-540-45319-9\_7.

## APPENDIX A. DINING PHILOSOPHERS WITH ONE LEFT-HANDED PHILOSOPHER

We sketch the formalization of the case in which the states of the forks are only “free” and “busy”. Consider an RTS with  $\Sigma = \{f, b, t, h, e\}$ . The state  $h$  represents philosophers who already grabbed the first fork and wait for the second one. All other states are used as before. The philosopher at index 1 takes first the fork at index 2 and then the fork at index  $n$ , while any other philosopher  $i > 1$  first takes the fork at index  $i - 1$  and then the fork at index  $i + 1$  (modulo  $n$ ). We modelled right-handed philosophers as taking the lower-index fork first as this looks like we are facing the philosopher. Naturally, the opposite choice of notation would not lead to substantial changes. For example, both reversing the direction and single-position cyclic shifts preserve regularity of languages.

In [ERW21b] the absence of deadlocks in this example is shown via only a few inductive assertions. These assertions can be equivalently expressed as 3-invariants. Moreover, these assertions are actually enough to completely characterize *Reach* in this example. To this end, observe that, analogously to Example 2.4, *Reach* is completely characterized by the absence of a few *invalid* patterns. These patterns separate into three cases: First, a philosopher should use some fork, but this fork is still considered free. Second, two philosophers are in states that require the same fork. Third, no adjacent philosopher currently uses some fork, yet this fork is busy. As long as neither of these violations happens, the configuration can be reached from the initial configuration with the philosophers picking up the forks they should be using in an arbitrary order, absence of the second forbidden pattern guarantees the lack of conflicts. Afterwards, the correct set of forks will be in use.

More formally, we get

- $\Sigma (\Sigma \Sigma)^* f (h | e) \Sigma (\Sigma \Sigma)^*, (\Sigma \Sigma)^+ e f (\Sigma \Sigma)^*, (e | h) f (\Sigma \Sigma)^*, e (\Sigma \Sigma)^* f,$
- $(\Sigma \Sigma)^+ e \Sigma (h | e) \Sigma (\Sigma \Sigma)^*, (e | h) \Sigma (e | h) \Sigma (\Sigma \Sigma)^*, e \Sigma (\Sigma \Sigma)^* e \Sigma,$
- $t b t \Sigma (\Sigma \Sigma)^*, (t | h) \Sigma (\Sigma \Sigma)^* (t | h) b, (\Sigma \Sigma)^+ (t | h) b t \Sigma (\Sigma \Sigma)^*.$

The absence of these patterns can be established with the following languages of inductive 1-invariants and inductive 3-invariants:

$$\begin{aligned}
& [\{e\}][\emptyset]([\emptyset][\emptyset])^*[\{e\}][\{f\}] && \begin{bmatrix} \{t, h\} \\ \{t, h\} \\ \emptyset \end{bmatrix} \begin{bmatrix} \emptyset \\ \emptyset \\ \emptyset \end{bmatrix} \left( \begin{bmatrix} \emptyset & \emptyset \\ \emptyset & \emptyset \end{bmatrix} \right)^* \begin{bmatrix} \emptyset & \{b\} \\ \{t, h\} & \emptyset \\ \{t, h\} & \{b\} \end{bmatrix} \\
& [\{e, h\}][\{f\}][\{e, h\}][\emptyset]([\emptyset][\emptyset])^* && \begin{bmatrix} \{t\} \\ \{t\} \\ \emptyset \end{bmatrix} \begin{bmatrix} \{b\} \\ \emptyset \\ \{b\} \end{bmatrix} \begin{bmatrix} \emptyset \\ \{t\} \\ \{t\} \end{bmatrix} \begin{bmatrix} \emptyset \\ \emptyset \\ \emptyset \end{bmatrix} \left( \begin{bmatrix} \emptyset & \emptyset \\ \emptyset & \emptyset \end{bmatrix} \right)^* \\
& [\emptyset][\emptyset]([\emptyset][\emptyset])^*[\{e\}][\{f\}][\{e, h\}][\emptyset]([\emptyset][\emptyset])^* && \begin{bmatrix} \emptyset & \emptyset \\ \emptyset & \emptyset \\ \emptyset & \emptyset \end{bmatrix} \left( \begin{bmatrix} \emptyset & \emptyset \\ \emptyset & \emptyset \end{bmatrix} \right)^* \begin{bmatrix} \{t, h\} \\ \{t, h\} \\ \emptyset \end{bmatrix} \begin{bmatrix} \emptyset \\ \{b\} \\ \{t\} \end{bmatrix} \begin{bmatrix} \{t\} \\ \emptyset \\ \emptyset \end{bmatrix} \left( \begin{bmatrix} \emptyset & \emptyset \\ \emptyset & \emptyset \end{bmatrix} \right)^*
\end{aligned}$$

Consequently, *IndInv<sub>3</sub>* and *Reach* coincide for this example, as the invariants forbid everything that is not reachable. This immediately implies that *IndInv<sub>3</sub>* proves deadlock-freedom since the system actually is deadlock-free.

However, *IndInv<sub>2</sub>* is insufficient to prove deadlock-freedom: assume there exists some inductive 2-invariant  $I$  that invalidates that  $D = h b t f e b$  can be reached. Then,  $I$  must separate all elements from *Reach* and all configurations  $D'$  with  $D' \rightsquigarrow^* D$  because it is inductive. In particular,  $D' = t b e b e b$  and the reachable configuration  $t b h b e b$ . Hence, one clause of  $I$  contains  $h_{3.6}$ . Consider the following pair of configurations:  $D'' = t f h f t f$  and  $C = t b h f t f$ .  $I$  must separate  $D''$  from  $C$  since  $D'' \rightsquigarrow t b e f t f \rightsquigarrow t b e b h f \rightsquigarrow t f t f h f \rightsquigarrow h b t f h f \rightsquigarrow D$  while  $C \in \text{Reach}$ . Since  $D'' \models h_{3.6}$ , this separation is based on the second clause of  $I$  which must contain  $b_{2.6}$ . This means  $t b h f e b \models I$ . Since  $I$  is

inductive and  $t b h f e b \rightsquigarrow t b e b e b \rightsquigarrow t f t f e b \rightsquigarrow h b t f e b = D$ , the assumption that  $I$  exists is wrong. Consequently,  $D$  cannot be excluded via inductive 2-invariants.

#### APPENDIX B. $IndInv_2$ FOR CACHE COHERENCE PROTOCOLS BERKELEY AND DRAGON

For both protocols, we follow the specification of [Del03].

**Berkeley.** In the Berkeley cache coherence protocol, each cell is in one of four different states: invalid ( $i$ ), unowned ( $u$ ), exclusive ( $e$ ), and shared ( $s$ ). Initially, all cells are invalid. Consequently, the language of initial configurations is  $i^*$ . For the transitions, we consider a few different events. The first one is that the memory is read, and the corresponding cell does provide some value of it; i.e., the cell is *not* in the state  $i$ . In this case, nothing changes:

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* .$$

If, on the other hand, a value is read from some cell that is in the state  $i$ , then this memory cell fetches the information without claiming ownership; i.e., moves into the state  $u$ . Every other memory cell observes this process. Thus, cells that previously were in  $e$  move to  $s$  to account for the fact that another memory cell holds the same information.

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [i] \\ [u] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* .$$

If a value is written to a cell that was invalid before, then this cell claims exclusive ownership; that is, all other cells are invalidated.

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [i] \end{array} \middle| \begin{array}{c} [s] \\ [i] \end{array} \right)^* \left( \begin{array}{c} [i] \\ [e] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [i] \end{array} \middle| \begin{array}{c} [s] \\ [i] \end{array} \right)^* .$$

If a cell already has exclusive ownership of this information, there is nothing to be done. If the cell has only shared ownership of the value, all other cells *that claim shared ownership* are invalidated.

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [i] \end{array} \right)^* \left( \begin{array}{c} [u] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [e] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [i] \end{array} \right)^* .$$

Finally, the cache can decide to drop data at any moment in time. Thus, any cell might move into the state  $i$ .

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [u] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [i] \end{array} \middle| \begin{array}{c} [s] \\ [i] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [u] \\ [u] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* .$$

We pose now the question whether a configuration can be reached, where two different cells are claiming exclusive access to the same data. The corresponding set  $\mathcal{U}$  corresponds to  $\Sigma^* e \Sigma^* e \Sigma^*$ . As shown in Figure 1,  $IndInv_1$  does not prove this property. Let us see why. Assume there is an inductive 1-invariant  $I$  which invalidates the bad word  $b = e e$ ; that is,  $b \not\models I$ . Observe now that we can reach  $b$  in one step from  $b' = u e$  and  $b'' = e u$ . Consequently,  $I$  cannot be satisfied by  $b'$  or  $b''$  either. Otherwise, since  $I$  is inductive, we already get  $b \models I$ . This means,  $I$  must not contain  $e_{1:2}$ ,  $e_{2:2}$ ,  $u_{1:2}$  nor  $u_{2:2}$ . This, however, makes  $I$  unsatisfiable for the actually reachable configuration  $u u$ .

Using an adapted version of the semi-automatic approach of [ERW21b] and some additional reasoning led us to the following language of inductive 2-invariants which exclude all configurations from  $\mathcal{U}$ :

$$\left[ \begin{array}{c} \emptyset \\ \emptyset \end{array} \right]^* \left[ \begin{array}{c} \{i, s, u\} \\ \{i\} \end{array} \right] \left[ \begin{array}{c} \emptyset \\ \emptyset \end{array} \right]^* \left[ \begin{array}{c} \{i\} \\ \{i, s, u\} \end{array} \right] \left[ \begin{array}{c} \emptyset \\ \emptyset \end{array} \right]^* .$$



Since  $IndInv_2$  is the strongest inductive 2-invariant, and the provided 2-invariant excludes all the configurations from  $\mathcal{U}$ ,  $IndInv_2$  also excludes all the configurations from  $\mathcal{U}$ . Thus,  $IndInv_2$  is strong enough to prove the property.

**Dragon.** The Dragon protocol distinguishes five states. As before, we have states for invalid cells ( $i$ ), cells that maintain an exclusive copy of the data ( $e$ ), and cells that have a (potentially) shared copy of the data ( $s$ ). In contrast to before, the Dragon protocol does not invalidate other copies of some data when it is updated. Instead two new states which mirror  $e$  and  $s$  are introduced but, additionally, indicate that the data might have changed. We refer to these states as  $\hat{e}$  and  $\hat{s}$ , respectively. Regardless, we initialize all cells as invalid; i.e., we have the initial language  $i^*$ .

Assume a read from a “valid” cell; that is, some cell that is not in the state  $i$ . In that case, nothing changes:

$$\left( \begin{array}{c} [i] \\ [e] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* \left( \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right) \left( \begin{array}{c} [i] \\ [e] \end{array} \middle| \begin{array}{c} [s] \\ [\hat{e}] \end{array} \middle| \begin{array}{c} [\hat{s}] \end{array} \right)^* .$$

If a read occurs from an invalid cell – while all cells are invalid – the accessed cell becomes an exclusive reference:

$$\begin{array}{c} [i] \\ [i] \end{array}^* \begin{array}{c} [i] \\ [e] \end{array} \begin{array}{c} [i] \\ [i] \end{array}^* .$$

If not all cells are invalid but a read occurs for an invalid cell, then this cell obtains a copy of the data, having now a shared reference to the data. Moreover, all exclusive references; i.e., cells in the states  $e$  or  $\hat{e}$ , move to their shared counterparts ( $s$  and  $\hat{s}$ , respectively).

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* \begin{array}{c} [i] \\ [s] \end{array} \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* .$$

It is possible here that, although all cells are invalid, the changing cell becomes only shared. Since this configuration can also be reached from the configuration where there are exactly two cells in state  $s$  and all others in  $i$  by one cell moving from its shared state to its invalid state, this does not change the set of actually reachable configurations and, consequently, is immaterial for the correctness analysis.

Writing a cell in the state  $\hat{e}$  does not change anything. On the other hand, writing a cell in the state  $e$  moves that cell into the state  $\hat{e}$ :

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* \begin{array}{c} [\hat{e}] \\ [\hat{e}] \end{array} \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^*$$

and

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* \begin{array}{c} [e] \\ [\hat{e}] \end{array} \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \right)^* .$$

A write operation on a cell that is the only one in state  $s$  or  $\hat{s}$  results in a change to  $\hat{e}$ . If there are other cells in either state, one moves to  $\hat{s}$  while all others move to  $s$ .

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{e}] \end{array} \right)^* \begin{array}{c} [\hat{s}] \\ [\hat{e}] \end{array} \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{e}] \end{array} \right)^* ,$$

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{e}] \end{array} \right)^* \begin{array}{c} [s] \\ [\hat{e}] \end{array} \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{e}] \end{array} \right)^*$$

and

$$\left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [s] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{s}] \\ [\hat{s}] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [s] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{s}] \\ [\hat{s}] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \\ \mid \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [s] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [\hat{s}] \\ [\hat{s}] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* \left( \begin{array}{c} [s] \\ [s] \end{array} \middle| \begin{array}{c} [\hat{s}] \\ [\hat{s}] \end{array} \right) \left( \begin{array}{c} [i] \\ [i] \end{array} \middle| \begin{array}{c} [e] \\ [e] \end{array} \middle| \begin{array}{c} [\hat{e}] \\ [\hat{s}] \end{array} \middle| \begin{array}{c} [s] \\ [s] \end{array} \right)^* .$$

If a value is written to a previously invalid cell, then either this cell moves to  $\hat{e}$  (assuming all other cells are  $i$  as well), while an occurrence of another cell with this value causes the written cell to become  $\hat{s}$  and all other cells to move to the state  $s$ .

$$\begin{bmatrix} i \\ i \end{bmatrix}^* \begin{bmatrix} i \\ \hat{e} \end{bmatrix} \begin{bmatrix} i \\ i \end{bmatrix}^*$$

and

$$\begin{aligned} & \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* \begin{bmatrix} i \\ \hat{s} \end{bmatrix} \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* \left( \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \right) \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* \\ & \mid \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* \left( \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \right) \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* \begin{bmatrix} i \\ \hat{s} \end{bmatrix} \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ s \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{s} \\ s \end{bmatrix} \right)^* . \end{aligned}$$

Finally, any cell might drop its content at any point.

$$\left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ e \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ \hat{e} \end{bmatrix} \begin{bmatrix} \hat{s} \\ \hat{s} \end{bmatrix} \right)^* \left( \begin{bmatrix} e \\ i \end{bmatrix} \begin{bmatrix} \hat{e} \\ i \end{bmatrix} \begin{bmatrix} \hat{s} \\ i \end{bmatrix} \begin{bmatrix} s \\ i \end{bmatrix} \right) \left( \begin{bmatrix} i \\ i \end{bmatrix} \begin{bmatrix} e \\ e \end{bmatrix} \begin{bmatrix} s \\ s \end{bmatrix} \begin{bmatrix} \hat{e} \\ \hat{e} \end{bmatrix} \begin{bmatrix} \hat{s} \\ \hat{s} \end{bmatrix} \right)^* .$$

We are interested now to establish that the language  $\Sigma^* \hat{e} \Sigma^* \hat{e} \Sigma^*$  cannot be reached. The proof that  $IndInv_1$  is insufficient to exclude all configurations of  $\Sigma^* \hat{e} \Sigma^* \hat{e} \Sigma^*$  is straightforward: Observe that both  $s \hat{e}$  and  $\hat{e} s$  can reach  $\hat{e} \hat{e}$  in one step. In consequence, analogously to the argument used for the Berkeley protocol, any inductive 1-invariant cannot distinguish between the reachable  $s s$  and the unreachable  $\hat{e} \hat{e}$ .

On the other hand, the language

$$\begin{bmatrix} \emptyset \\ \emptyset \end{bmatrix}^* \begin{bmatrix} \{\hat{s}, i, s\} \\ \{i\} \end{bmatrix} \begin{bmatrix} \emptyset \\ \emptyset \end{bmatrix}^* \begin{bmatrix} \{i\} \\ \{\hat{s}, i, s\} \end{bmatrix} \begin{bmatrix} \emptyset \\ \emptyset \end{bmatrix}^*$$

of inductive 2-invariants (which arose, again, from an adapted version of the semi-automatic approach of [ERW21b]) induces an abstraction disjoint from  $\Sigma^* \hat{e} \Sigma^* \hat{e} \Sigma^*$ . Consequently,  $IndInv_2$  does as well.