Computing Inductive Invariants of Regular Abstraction Frameworks

Philipp Czerner
Technical University of Munich, Germany

Javier Esparza
Technical University of Munich, Germany

Valentin Krasotin
Technical University of Munich, Germany

Christoph Welzel-Mohr
Technical University of Munich, Germany

Abstract

Regular transition systems (RTS) are a popular formalism for modeling infinite-state systems in general, and parameterised systems in particular. In a CONCUR 22 paper, Esparza et al. introduce a novel approach to the verification of RTS, based on inductive invariants. The approach computes the intersection of all inductive invariants of a given RTS that can be expressed as CNF formulas with a bounded number of clauses, and uses it to construct an automaton recognising an overapproximation of the reachable configurations. The paper shows that the problem of deciding if the language of this automaton intersects a given regular set of unsafe configurations is in EXPSPACE and PSPACE-hard.

We introduce *regular abstraction frameworks*, a generalisation of the approach of Esparza et al., very similar to the regular abstractions of Hong and Lin. A framework consists of a regular language of *constraints*, and a transducer, called the *interpretation*, that assigns to each constraint the set of configurations of the RTS satisfying it. Examples of regular abstraction frameworks include the formulas of Esparza et al., octagons, bounded difference matrices, and views. We show that the generalisation of the decision problem above to regular abstraction frameworks remains in EXPSPACE, and prove a matching (non-trivial) EXPSPACE-hardness bound.

EXPSPACE-hardness implies that, in the worst case, the automaton recognising the overapproximation of the reachable configurations has a double-exponential number of states. We introduce a learning algorithm that computes this automaton in a lazy manner, stopping whenever the current hypothesis is already strong enough to prove safety. We report on an implementation and show that our experimental results improve on those of Esparza et al.

2012 ACM Subject Classification Theory of computation → Regular languages; Theory of computation → Problems, reductions and completeness; Theory of computation → Verification by model checking

Keywords and phrases Regular model checking, abstraction, inductive invariants

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2024.19

Related Version *Full Version*: `https://arxiv.org/abs/2404.10752` [16]

Supplementary Material *Software*: `https://doi.org/10.5281/zenodo.12734991` [33]
archived at `swh:1:snp:9789ae3f53de50f43367d1c0f108665ea96b619e`

1 Introduction

Regular transition systems (RTS) are a popular formalism for modelling infinite-state systems satisfying the following conditions: configurations can be encoded as words, the set of initial configurations is recognised by a finite automaton, and the transition relation is recognised by a transducer. Model checking RTS has been intensely studied under the name of *regular* *model checking* (see [23, 13, 24, 10] and the surveys [5, 1, 6, 2]). Most regular model checking algorithms address the *safety problem*: given a regular set of unsafe configurations, decide if its intersection with the set of reachable configurations is empty or not. They combine algorithms for the computation of increasingly larger regular subsets of the reachable configurations with acceleration, abstraction, and widening techniques [13, 23, 17, 4, 10, 12, 14, 11, 26, 15].

Recently, Esparza et al. have introduced a novel approach that, starting with the set of all configurations of the RTS, computes increasingly smaller inductive invariants, that is, inductive supersets of the reachable configurations. More precisely, [19] considers invariants given by Boolean formulas in conjunctive normal form with at most b clauses. The paper proves that, for every bound $b \geq 0$, the intersection of *all* inductive b-invariants of the system is recognised by a DFA of double exponential size in the RTS. As a corollary, they obtain that, for every $b \geq 0$, deciding if this intersection contains some unsafe configuration is in EXPSPACE. They also show that the problem is PSPACE-hard, and leave the question of closing the gap open.

In [20] (a revised version of [19]), the EXPSPACE proof is conducted in a more general setting than in [19]. Inspired by this, in our first contribution we show that the approach of [19] can be vastly generalised to arbitrary *regular abstraction frameworks*, consisting of a regular language of *constraints*, and an *interpretation*. Interpretations are functions, represented by transducers, that assign to each constraint a set of configurations, viewed as the set of configurations that *satisfy* the constraint. Examples of regular abstraction frameworks include the formulas of [19] for every $b \geq 0$, views [3], and families of Presburger arithmetic formulas like octagons [27] or bounded difference matrices [25, 8]. A framework induces an abstract interpretation, in which, loosely speaking, the word encoding a constraint is the abstraction of the set of configurations satisfying the constraint. Just as regular model checking started with the observation that different classes of *systems* could be uniformly modeled as RTSs [5, 1, 6, 2], we add the observation, also made in [21], that different classes of *abstractions* can be uniformly modeled as regular abstraction frameworks. We show that the generalisation of the verification problem of [19, 20] to arbitrary regular abstraction frameworks remains in EXPSPACE.

In our second contribution we show that our problem is also EXPSPACE-hard. The reduction (from the acceptance problem for exponentially bounded Turing machines) is surprisingly involved. Loosely speaking, it requires to characterise the set of prefixes of the run of a Turing machine on a given word as an intersection of inductive invariants of a very restrictive kind. We think that this construction can be of independent interest.

Our third and final contribution is motivated by the EXPSPACE-hardness result. A consequence of this lower bound is that the automaton recognising the overapproximation of the reachable configurations must necessarily have a double-exponential number of states in the worst case. We present an approach, based on automata learning, that constructs increasingly larger automata that recognise increasingly smaller overapproximations, and checks whether they are precise enough to prove safety. A key to the approach is solving the separability problem: given a pair (c, c') of configurations, is there an inductive constraint that *separates* c and c', i.e. is satisfied by c but not by c'? We show that the problem is PSPACE-complete and NP-complete for interpretations captured by length-preserving transducers. We provide an implementation on top of a SAT solver for the latter case (this is the only case considered in [19, 20]). An experimental comparison shows that this approach beats the one of [19, 20].

Related work. As mentioned above, our first contribution is a reformulation of results of [20] into a more ambitious formalism; it is a conceptual but not a technical novelty. The second and third contributions are new technical results.

Our regular abstraction frameworks are in the same spirit as the regular abstractions of Hong and Lin [21], which use regular languages as abstract objects. In this paper we concentrate on the inductive invariant approach of [19], and in particular on its complexity. This is unlike the approach of [21], which on the one hand is more general, since it also considers liveness properties, but on the other hand does not contain complexity results.

Automata learning has been explored for the verification of regular transition systems multiple times [28, 31, 15, 32, 29]. Roughly speaking, all these approaches formulate a learning process to obtain a *regular* inductive invariant of the system that proves a safety property. Since it is impossible to algorithmically identify the cases where such regular inductive invariant exists, timeouts [15] and resource limits [28] are used as heuristics. In contrast, our approach is designed to always terminate. In particular, we either provide a regular set of constraints that suffices to establish the safety property or a pair of configurations that cannot be separated by inductive constraints of the considered framework. This information can be used to design a more precise framework by adding a new type of constraints.

2 Preliminaries and regular transition systems

Automata. Let Σ be an alphabet. A *nondeterministic finite automaton (NFA)* over Σ is a tuple $A = (Q, \Sigma, \delta, Q_0, F)$ where Q is a finite set of *states*, $\delta : Q \times \Sigma \to \mathcal{P}(Q)$ is the *transition function*, $Q_0 \subseteq Q$ is the set of *initial states*, and $F \subseteq Q$ is the set of *final states*. A *run* of A on a word $w = w_1 \cdots w_l \in \Sigma^l$ is a sequence $q_0 q_1 \cdots q_l$ of states where $q_0 \in Q_0$ and $\forall i \in [l] : q_i \in \delta(q_{i-1}, w_i)$. A run on w is *accepting* if $q_l \in F$, and A *accepts* w if there exists an accepting run of A on w. The language *recognised* by A, denoted $L(A)$ or L_A, is the set of words accepted by A. If $|Q_0| = 1$ and $|\delta(q, a)| = 1$ for every $q \in Q, a \in \Sigma$ $|Q_0| = 1$, then A is a deterministic finite automaton *(DFA)*. In this case, we write $\delta(q, a) = q'$ instead of $\delta(q, a) = \{q'\}$ and have a single initial state q_0 instead of a set Q_0.

Relations. Let $R \subseteq X \times Y$ be a relation. The *complement* of R is the relation $\overline{R} := \{(x, y) \in X \times Y \mid (u, w) \notin R\}$. The *inverse* of R is the relation $R^{-1} := \{(y, x) \in Y \times X \mid (x, y) \in R\}$. The *projections* of R onto its first and second components are the sets $R|_1 := \{x \in X \mid \exists y \in Y : (x, y) \in R\}$ and $R|_2 := \{y \in Y \mid \exists x \in X : (x, y) \in R\}$. The *join* of two relations $R \subseteq X \times Y$ and $S \subseteq Y \times Z$ is the relation $R \circ S := \{(x, z) \in X \times Z \mid \exists y \in Y : (x, y) \in R, (y, z) \in S\}$. The *post-image* of a set $X' \subseteq X$ under a relation $R \subseteq X \times Y$, denoted $X' \circ R$ or $R(X')$, is the set $\{y \in Y \mid \exists x \in X' : (x, y) \in R\}$; the *pre-image*, denoted $R \circ Y$ or $R^{-1}(Y)$, is defined analogously. Throughout this paper, we only consider relations where $X = \Sigma^*$ and $Y = \Gamma^*$ for some alphabets Σ, Γ. We just call them relations. A relation $R \subseteq \Sigma^* \times \Gamma^*$ is *length-preserving* if $(u, w) \in R$ implies $|u| = |w|$.

Convolutions and transducers. Let Σ, Γ be alphabets, let $\# \notin \Sigma \cup \Gamma$ be a padding symbol, and let $\Sigma_\# := \Sigma \cup \{\#\}$ and $\Gamma_\# := \Gamma \cup \{\#\}$. The *convolution* of two words $u = a_1 \ldots a_k \in \Sigma^*$ and $w = b_1 \ldots b_l \in \Gamma^*$, denoted $\left[\begin{smallmatrix} u \\ w \end{smallmatrix}\right]$, is the word over the alphabet $\Sigma_\# \times \Gamma_\#$ defined as follows. Intuitively, $\left[\begin{smallmatrix} u \\ w \end{smallmatrix}\right]$ is the result of putting u on top of w, aligned left, and padding the shorter of u and w with $\#$. Formally, if $k \leq l$, then $\left[\begin{smallmatrix} u \\ w \end{smallmatrix}\right] = \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right] \cdots \left[\begin{smallmatrix} a_k \\ b_k \end{smallmatrix}\right] \left[\begin{smallmatrix} \# \\ b_{k+1} \end{smallmatrix}\right] \cdots \left[\begin{smallmatrix} \# \\ b_l \end{smallmatrix}\right]$, and otherwise $\left[\begin{smallmatrix} u \\ w \end{smallmatrix}\right] = \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right] \cdots \left[\begin{smallmatrix} a_l \\ b_l \end{smallmatrix}\right] \left[\begin{smallmatrix} a_{l+1} \\ \# \end{smallmatrix}\right] \cdots \left[\begin{smallmatrix} a_k \\ \# \end{smallmatrix}\right]$. The convolution of a tuple of words $u_1 \in \Sigma_1^*, \ldots, u_k \in \Sigma_k^*$ is defined analogously, putting all k words on top of each other, aligned left, and padding the shorter words with $\#$.

A *transducer* over $\Sigma \times \Gamma$ is an NFA over $\Sigma_\# \times \Gamma_\#$. The binary relation recognised by a transducer T over $\Sigma \times \Gamma$, denoted $\mathsf{R}(T)$, is the set of pairs $(u, w) \in \Sigma^* \times \Gamma^*$ such that T accepts $\left[\begin{smallmatrix} u \\ w \end{smallmatrix}\right]$. The definition is generalised to relations of higher arity in the obvious way. In the paper transducers recognise binary relations unless mentioned otherwise. A relation is *regular* if it is recognised by some transducer. A transducer is *length-preserving* if it recognises a length-preserving relation.

Complexity of operations on automata and transducers. Given NFAs A_1, A_2 over Σ with n_1 and n_2 states, DFAs B_1, B_2 over Σ with m_1 and m_2 states, and transducers T_1 over $\Sigma \times \Gamma$ and T_2 over $\Gamma \times \Sigma$ with l_1 and l_2 states, the following facts are well known (see e.g. chapters 3 and 5 of [18]):

- there exist NFAs for $\mathsf{L}(A_1) \cup \mathsf{L}(A_2)$, $\mathsf{L}(A_1) \cap \mathsf{L}(A_2)$, and $\overline{\mathsf{L}(A_1)}$ with at most $n_1 + n_2$, $n_1 n_2$, and 2^{n_1} states, respectively;
- there exist DFAs for $\mathsf{L}(B_1) \cup \mathsf{L}(B_2)$, $\mathsf{L}(B_1) \cap \mathsf{L}(B_2)$, and $\overline{\mathsf{L}(B_1)}$ with at most $m_1 m_2$, $m_1 m_2$, and m_1 states, respectively;
- there exist NFAs for $\mathsf{R}(T_1)|_1$ and $\mathsf{R}(T_1)|_2$ and a transducer for $\mathsf{R}(T_1)^{-1}$ with at most l_1 states;
- there exists a transducer for $\mathsf{R}(T_1) \circ \mathsf{R}(T_2)$ with at most $l_1 l_2$ states; and
- there exist NFAs for $\mathsf{L}(A_1) \circ \mathsf{R}(T_1)$ and $\mathsf{R}(T_1) \circ \mathsf{L}(A_2)$ with at most $n_1 l_1$ and $l_1 n_2$ states, respectively.

Regular transition systems

We recall standard notions about regular transition systems and fix some notations. A *transition system* is a pair $\mathcal{S} = (\mathcal{C}, \Delta)$ where $\mathcal{C}$ is the set of all possible *configurations* of the system, and $\Delta \subseteq \mathcal{C} \times \mathcal{C}$ is a *transition relation.* The *reachability relation Reach* is the reflexive and transitive closure of Δ. Observe that, by our definition of post-set, $\Delta(C)$ and $\mathit{Reach}(C)$ are the sets of configurations reachable in one step and in arbitrarily many steps from C, respectively.

Regular transition systems are transition systems where Δ can be finitely represented by a transducer. Formally:

▶ **Definition 1.** *A transition system* $\mathcal{S} = (\mathcal{C}, \Delta)$ *is* regular *if* $\mathcal{C}$ *is a regular language over some alphabet* Σ*, and* Δ *is a regular relation. We abbreviate regular transition system to RTS.*

RTSs are often used to model parameterised systems [5, 1, 6, 2]. In this case, Σ is the set of possible *states* of a process, the set of configurations is $\mathcal{C} = \Sigma^* \setminus \{\varepsilon\}$, and a configuration $a_1 \cdots a_n \in \Sigma^*$ describes the global state of an *array* consisting of n identical copies of the process, with the i-th process in state a_i for every $1 \leq i \leq n$. The transition relation Δ describes the possible transitions of all arrays, of any length.

▶ **Example 2** (Token passing [5])**.** We use a version of the well-known token passing algorithm as running example. We have an array of processes of arbitrary length. At each moment in time, a process either has a token (t) or not (n). Initially, only the first process has a token. A process that has a token can pass it to the process to the right if that process does not have one. We set $\Sigma = \{t, n\}$, and so $\mathcal{C} = \{t, n\}^* \setminus \{\varepsilon\}$. We have $c_2 \in \Delta(c_1)$ iff the word $\left[\begin{smallmatrix} c_1 \\ c_2 \end{smallmatrix}\right]$ belongs to the regular expression $\left(\left[\begin{smallmatrix} n \\ n \end{smallmatrix}\right] + \left[\begin{smallmatrix} t \\ t \end{smallmatrix}\right]\right)^* \left(\left[\begin{smallmatrix} t \\ n \end{smallmatrix}\right]\left[\begin{smallmatrix} n \\ t \end{smallmatrix}\right]\right) \left(\left[\begin{smallmatrix} n \\ n \end{smallmatrix}\right] + \left[\begin{smallmatrix} t \\ t \end{smallmatrix}\right]\right)^*$. For the set of initial configurations $C_I := tn^*$ where only the first process has a token, the set of reachable configurations is $\mathit{Reach}(C_I) = n^* \, t \, n^*$.

3 Regular abstraction frameworks

In the same way that RTSs can model multiple classes of *systems* (e.g. parameterised systems with synchronous/asynchronous, binary/multiway/broadcast communication), regular abstraction frameworks are a formalism to model a wide range of *abstractions*.

▶ **Definition 3.** *An* abstraction framework *is a triple* $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$*, where* $\mathcal{C}$ *is a set of* configurations*,* $\mathcal{A}$ *is a set of* constraints*, and* $\mathcal{V} \subseteq \mathcal{A} \times \mathcal{C}$ *is an* interpretation*.* $\mathcal{F}$ *is* regular *if* $\mathcal{C}$ *and* $\mathcal{A}$ *are regular languages over alphabets* Σ *and* Γ*, respectively, and the interpretation* $\mathcal{V}$ *is a regular relation over* $\mathcal{A} \times \mathcal{C}$*.*

Intuitively, the constraints of an abstraction framework are the abstract objects of the abstraction, and $\mathcal{V}(A)$ is the set of configurations abstracted by A. The following remark formalises this.

▶ Remark 4. An abstraction framework $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$ induces an abstract interpretation as follows. The concrete and abstract domains are $(2^{\mathcal{C}}, \leq_{\mathcal{C}})$ and $(2^{\mathcal{A}}, \leq_{\mathcal{A}})$, respectively, where $\leq_{\mathcal{C}} := \subseteq$ and $\leq_{\mathcal{A}} := \supseteq$. Both are complete lattices. The concretisation function $\gamma\colon 2^{\mathcal{A}} \to 2^{\mathcal{C}}$ and the abstraction function $\alpha\colon 2^{\mathcal{C}} \to 2^{\mathcal{A}}$ are given by:

- $\gamma(\mathcal{A}') := \bigcap_{A \in \mathcal{A}'} \mathcal{V}(A)$. Intuitively, $\gamma(\mathcal{A}')$ is the set of configurations that satisfy all constraints of $\mathcal{A}'$. In particular, $\gamma(\emptyset) = \mathcal{C}$.
- $\alpha(\mathcal{C}') := \{A \in \mathcal{A} \mid \mathcal{C}' \subseteq \mathcal{V}(A)\}$. Intuitively, $\alpha(\mathcal{C}')$ is the set of constraints satisfied by all configurations in $\mathcal{C}'$. In particular, $\alpha(\emptyset) = \mathcal{A}$.

It is easy to see that the functions α and γ form a Galois connection, that is, for all $C \subseteq \mathcal{C}$ and $A \subseteq \mathcal{A}$, we have $\mathcal{B} \subseteq \alpha(C) \Leftrightarrow C \subseteq \gamma(\mathcal{B})$.

Regular abstractions can be combined to yield more precise ones. Given abstraction frameworks $\mathcal{F}_1 = (\mathcal{C}, \mathcal{A}_1, \mathcal{V}_1)$ and $\mathcal{F}_2 = (\mathcal{C}, \mathcal{A}_2, \mathcal{V}_2)$, we can define new frameworks $(\mathcal{C}, \mathcal{A}, \mathcal{V})$ by means of the following operations:

- Union: $\mathcal{A} := \mathcal{A}_1 \cup \mathcal{A}_2$, $\mathcal{V}(A) := \mathcal{V}_1(A)$ if $A \in \mathcal{A}_1$, else $\mathcal{V}_2(A)$.
  A constraint of the union framework is either a constraint of the first framework, or a constraint of the second.
- Convolution: $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$, $\mathcal{V}(A_1, A_2) := \mathcal{V}_1(A_1) \cap \mathcal{V}_2(A_2)$.
  A constraint of the convolution framework is the conjunction of two constraints, one of each framework. This operation is implicitly used in [19]: the constraint for a Boolean formula with b clauses is the convolution, applied b times, of the constraints for formulas with one clause.

The proof of the following lemma is in the full version of the paper [16].

▶ **Lemma 5.** *Regular abstraction frameworks are closed under union and convolution. If the interpretations of the frameworks are recognised by transducers with* n_1 *and* n_2 *states, then the interpretations of the union and convolution frameworks are recognised by transducers with* $O(n_1 + n_2)$ *and* $O(n_1 n_2)$ *states, respectively.*

Many abstractions used in the literature can be modeled as regular abstraction frameworks. We give some examples.

▶ **Example 6.** Consider a transition system where $\mathcal{C} = \mathbb{N}^d$ for some d, and Δ is given by a formula of Presburger arithmetic $\delta(\mathbf{x}, \mathbf{x}')$, that is, $(\mathbf{n}, \mathbf{n}') \in \Delta$ iff $\delta(\mathbf{n}, \mathbf{n}')$ holds. It is well-known that for any Presburger formula there is a transducer recognising the set of its solutions when numbers are encoded in binary, and so with this encoding $(\mathcal{C}, \Delta)$ is an RTS. Any Presburger formula $\varphi(\mathbf{x}, \mathbf{y})$, where $\mathbf{x}$ has dimension d and $\mathbf{y}$ has some arbitrary dimension e, induces a regular abstraction framework as follows. The set of constraints is the set of all tuples $\mathbf{m} \in \mathbb{N}^e$; the interpretation assigns to $\mathbf{m}$ all tuples $\mathbf{n}$ such that $\varphi(\mathbf{n}, \mathbf{m})$ holds. Intuitively, the constraints are the formulas $\varphi_{\mathbf{m}}(\mathbf{x}) := \varphi(\mathbf{x}, \mathbf{m})$, but using $\mathbf{m}$ as encoding of $\varphi_{\mathbf{m}}$.

Special cases of this setting are used in many different areas. For example, bounded difference matrices (see e.g. [25, 8]) and octagons [27] correspond to abstraction frameworks with constraints $\varphi(x_1, x_2, y)$ of the form $x_1 \pm x_2 \leq y$.

▶ **Example 7.** The approach to regular model checking of [19] is another instance of a regular abstraction framework. The paper encodes sets of configurations as positive Boolean formulas in conjunctive normal form with a bounded number b of clauses. We explain this by means of an example. Consider an RTS with $\Sigma = \{a, b, c\}$ and $\mathcal{C} = \Sigma^*$. Consider the formula $\varphi = (a_{1:5} \vee b_{1:5} \vee a_{3:5}) \wedge b_{4:5}$. We interpret φ on configurations. The intended meaning of a literal, say $a_{1:5}$, is "if the configuration has length 5, then its first letter is an a." So the set of configurations satisfying the formula is $\Sigma^{\leq 4} + \Sigma^6\Sigma^* + (a + b)\Sigma^2 b\Sigma + \Sigma^2 ab\Sigma$. In the formulas of [19] all literals have the same length, where the length of a literal $x_{i:j}$ is j.

Formulas with at most b clauses can be encoded as words over the alphabet $\Gamma = (2^\Sigma)^b$. Each clause is encoded as a word over 2^Σ. For example, the encodings of the clauses $(a_{1:5} \vee b_{1:5} \vee a_{3:5})$ and $b_{4:5}$ are $\{a, b\}\emptyset\{a\}\emptyset\emptyset$ and $\emptyset\emptyset\emptyset\{b\}\emptyset$, and the encoding of φ is the convolution of the encodings of the clauses. It is easy to see that the interpretation of [19] that assigns to a formula the set of configurations satisfying it is a regular relation recognised by a transducer with 2^b states [19]. In particular, for the case $b = 1$ we get the two-state transducer on the left of Figure 1.

▶ **Example 8.** In [3] Abdulla et al. introduce *view abstraction* for the verification of parameterised systems. Given a number $k \geq 1$, a *view* of a word $w \in \Sigma^*$ is a scattered subword of w. Loosely speaking, Abdulla et al. abstract a word by its set of views of length up to k. In our setting, a constraint is a set $F \subseteq \Sigma^{\leq k}$ of "forbidden views", and $\mathcal{V}(V)$ is the set of all words that do not contain any view of F. Since k is fixed, this interpretation is regular.

3.1 The abstract safety problem

We apply regular abstraction frameworks to the problem of deciding whether an RTS avoids some regular set of unsafe configurations. For simplicity, we assume w.l.o.g. that the set of configurations of the RTS is Σ^*[1]. Let us first formalise the Safety problem:

Given: a nondeterministic transducer recognising a regular relation $\Delta \subseteq \Sigma^* \times \Sigma^*$, and two NFAs recognising regular sets $C_I, C_U \subseteq \Sigma^*$ of initial and unsafe configurations, respectively.

Decide: does $\mathit{Reach}(C_I) \cap C_U = \emptyset$ hold?

It is a folklore result that Safety is undecidable. Let us sketch the argument. The configurations of a given Turing machine can be encoded as words of the form $w_l\, q\, w_r$, where w_l, w_r encode the contents of the tape to the left and to the right of the head, and q encodes the current state. With this encoding, the successor relation between configurations of the Turing machine is regular, and so is the set of accepting configurations. Taking the latter as set of unsafe configurations, the Turing machine accepts a given initial configuration iff the RTS started at the initial configuration is unsafe.

[1] By interpreting Δ as a relation over $\Sigma \times \Sigma$, any RTS can be transformed into an equivalent one with the same transitions where the set of configurations is Σ^*.

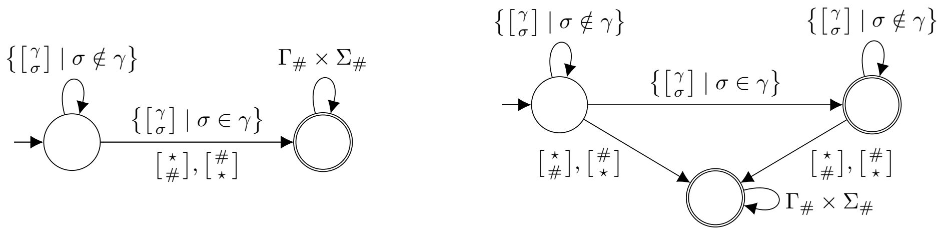

Figure 1 Transducers for the interpretations of Example 7 and 10. We have $\Gamma = 2^\Sigma$, and so the alphabet of the transducer is $(2^\Sigma)_\# \times \Sigma_\#$. The symbols $\left[\begin{smallmatrix}\star\\ \#\end{smallmatrix}\right]$ and $\left[\begin{smallmatrix}\#\\ \star\end{smallmatrix}\right]$ stand for the sets of all letters of the form $\left[\begin{smallmatrix}\gamma\\ \#\end{smallmatrix}\right]$ and $\left[\begin{smallmatrix}\#\\ \sigma\end{smallmatrix}\right]$, respectively.

AbstractSafety. We show that regular abstraction framework induces an "abstract" version of the safety problem, in which we replace the reachability relation by an overapproximation derived from the abstraction framework. Fix an RTS $\mathcal{S} = (\mathcal{C}, \Delta)$ and a regular abstraction framework $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$. We introduce some definitions:

▶ **Definition 9.** *A set $C \subseteq \mathcal{C}$ of configurations is* inductive *if $\Delta(C) \subseteq C$. A constraint A is* inductive *if $\mathcal{V}(A)$ is inductive. We let $Ind \subseteq \mathcal{A}$ denote the set of all inductive constraints of $\mathcal{A}$. Given two configurations c, c' and $A \in Ind$, we say that A* separates *c from c' if $c \in \mathcal{V}(A)$ and $c' \notin \mathcal{V}(A)$.*

It is a folklore result that $Reach(C)$ is the smallest inductive set containing C, and that if some $A \in Ind$ separates c and c', then $(c, c') \notin Reach$. Hence, an abstraction framework $(\mathcal{C}, \mathcal{A}, \mathcal{V})$ induces a *potential reachability* relation $PReach \subseteq \mathcal{C} \times \mathcal{C}$, defined as the set of all pairs of configurations that are not separated by any inductive constraint. Formally:

$$PReach := \{(c, c') \in \mathcal{C} \times \mathcal{C} \mid \forall A \in Ind\colon c \in \mathcal{V}(A) \to c' \in \mathcal{V}(A)\}$$

We have $Reach(C) \subseteq PReach(C)$ for every set of configurations C. In particular, given sets $C_I, C_U \subseteq \mathcal{C}$ of initial and unsafe configurations, if $PReach(C_I) \cap C_U = \emptyset$, then the RTS is safe.

▶ **Example 10.** Consider the RTS of the token passing system of Example 2, where $\Sigma = \{t, n\}$. We give two examples of abstraction frameworks. The first one is the abstraction framework of [19], already presented in Example 7, with $b = 1$. We have $\Gamma = 2^\Sigma = \{\emptyset, \{t\}, \{n\}, \Sigma\}$. A constraint like $\varphi = \bigvee_{i=3}^{5} t_{i:5}$ is encoded by the word $\emptyset\emptyset\{t\}\{t\}\{t\} \in \Gamma^*$, and interpreted as the set of all configurations of length 5 that have a token at positions 3, 4, or 5, plus the set of all configurations of length different from 5. The two-state transducer for this interpretation is on the left of Figure 1. For example, the left state has transitions leading to itself for the letters $\left[\begin{smallmatrix}\emptyset\\ n\end{smallmatrix}\right], \left[\begin{smallmatrix}\emptyset\\ t\end{smallmatrix}\right], \left[\begin{smallmatrix}\{t\}\\ n\end{smallmatrix}\right], \left[\begin{smallmatrix}\{n\}\\ t\end{smallmatrix}\right]$. The constraint φ is inductive. In fact, the language of all non-trivial inductive constraints (a constraint is trivial if it is satisfied by all configurations or by none) is $\{n\}^+\emptyset^*\{t\}^* + \{n\}^*\emptyset^*\{t\}^+$. The set of configurations potentially reachable from $C_I = tn^*$ is $PReach(C_I) = (tn + nn^*t)(t + n)^*$. In particular, $PReach(C_I) \cap n^* = \emptyset$, but $tnt \in PReach(C_I)$. So this abstraction framework is strong enough to prove that every reachable configuration has at least one token, but not to prove that it has exactly one.

Consider now the framework in which, instead of a *disjunction* of literals, a constraint is an *exclusive disjunction* of literals, that is, a configuration satisfies the constraint if it satisfies *exactly one* of its literals. So, in particular, the interpretation of $\emptyset\emptyset\{t\}\{t\}\{t\}$ is now that exactly one of the positions 3, 4, and 5 has a token. The interpretation is also

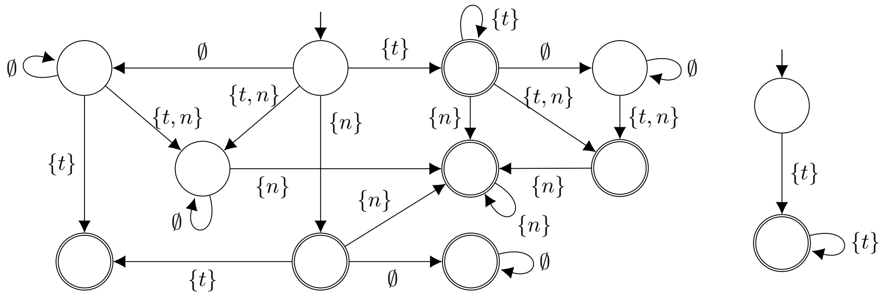

Figure 2 On the left, DFA recognising all non-trivial inductive constraints of Example 17. On the right, fragment with the same interpretation as the DFA on the left.

regular; it is given by the three-state transducer on the right of Figure 1. Examples of inductive constraints are $\{t\}\emptyset\{t,n\}\{n\}$ and all words of $\{t\}^*$. The language of non-trivial inductive constraints is given by the DFA on the left of Figure 2. Observe that the set of words satisfying all constraints of $\{t\}^*$ is the language n^*tn^*. In particular, we have $\mathit{PReach}(C_I) \subseteq n^*tn^* = \mathit{Reach}(C_I)$, and so $\mathit{PReach}(C_I) = \mathit{Reach}(C_I)$.

The AbstractSafety problem is defined exactly as Safety, just replacing the reachability set $\mathit{Reach}(C_I)$ by the potential reachability set $\mathit{PReach}(C_I)$ implicitly defined by the regular abstraction framework:

Given: a nondeterministic transducer recognising a regular relation $\Delta \subseteq \Sigma^* \times \Sigma^*$; two NFAs recognising regular sets $C_I, C_U \subseteq \Sigma^*$ of initial and unsafe configurations, respectively; and a deterministic transducer recognising a regular interpretation $\mathcal{V}$ over $\Gamma \times \Sigma$.

Decide: does $\mathit{PReach}(C_I) \cap C_U = \emptyset$ hold?

Recall that Safety is undecidable. In the rest of this section and in the next one we show that AbstractSafety is EXPSPACE-complete. Membership in EXPSPACE was essentially proved in [20], while EXPSPACE-hardness was left open. We briefly summarise the proof of membership in EXPSPACE presented in [20], for future reference in our paper.

▶ Remark 11. The result we prove in Section 3.2 is slightly more general. In [20], membership in EXPSPACE is only proved for RTSs whose transducers are length-preserving, while we prove it in general. General transducers allow one to model parameterised systems with process creation. For example, we can model a token passing algorithm in which the size of the array can dynamically grow and shrink by adding the transitions $\left(\left[\begin{smallmatrix} n \\ n \end{smallmatrix}\right] + \left[\begin{smallmatrix} t \\ t \end{smallmatrix}\right]\right)^+ \left(\left[\begin{smallmatrix} n \\ \# \end{smallmatrix}\right] + \left[\begin{smallmatrix} \# \\ n \end{smallmatrix}\right]\right)$ to the transition relation of Example 2.

3.2 AbstractSafety is in EXPSPACE

We first show that the set of all inductive constraints of a regular abstraction framework is a regular language. Fix a regular abstraction framework $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$ over an RTS $(\mathcal{C}, \Delta)$. Let n_Δ, $n_\mathcal{V}$, n_I, n_U be the number of states of the transducers and NFAs of a given instance of AbstractSafety.

▶ **Lemma 12** ([20]). *The set $\overline{\mathit{Ind}}$ is regular. Further, one can compute an NFA with at most $n_\Delta \cdot n_\mathcal{V}^2$ states recognising $\overline{\mathit{Ind}}$, and a DFA with at most $2^{n_\Delta \cdot n_\mathcal{V}^2}$ states recognising Ind.*

Proof. By definition, we have

$$\begin{aligned}\overline{Ind} &= \{A \in \Gamma^* \mid \exists c, c' \in \mathcal{C} \colon c \in \mathcal{V}(A), c' \in \Delta(c), \text{ and } c' \notin \mathcal{V}(A)\} \\ &= \{A \in \Gamma^* \mid \exists c, c' \in \mathcal{C} \colon (A, c) \in \mathcal{V}, (c, c') \in \Delta \text{ and } (c', A) \in \overline{\mathcal{V}^{-1}}\}\end{aligned}$$

Let $Id_\Gamma = \{(A, A) \mid A \in \Gamma^*\}$. We obtain $\overline{Ind} = \left(\left(\mathcal{V} \circ \Delta \circ \overline{\mathcal{V}^{-1}}\right) \cap Id_\Gamma \right)|_1$. By the results at the end of Section 2, $\overline{Ind}$ is recognised by a NFA with $n_\mathcal{V} \cdot n_\Delta \cdot n_\mathcal{V} = n_\Delta n_\mathcal{V}^2$ states, and so Ind is recognised by a DFA with $2^{n_\Delta \cdot n_\mathcal{V}^2}$ states. ◀

▶ **Lemma 13** ([20]). *The potential reachability relation PReach is regular. Further, one can compute a nondeterministic transducer with at most $K := n_\mathcal{V}^2 \cdot 2^{n_\Delta \cdot n_\mathcal{V}^2}$ states recognising $\overline{PReach}$, and a deterministic transducer with at most 2^K states recognising PReach.*

Proof. By definition, we have

$$\begin{aligned}\overline{PReach} &= \{(c, c') \in \mathcal{C} \times \mathcal{C} \mid \exists A \in Ind \colon c \in \mathcal{V}(A) \text{ and } c' \notin \mathcal{V}(A)\} \\ &= \left\{(c, c') \in \mathcal{C} \times \mathcal{C} \mid \exists A \in Ind \colon (c, A) \in \mathcal{V}^{-1} \text{ and } (A, c') \in \overline{\mathcal{V}}\right\}\end{aligned}$$

Let $Id_\Gamma = \{(A, A) \mid A \in \Gamma^*\}$. We obtain $\overline{PReach} = \left(\mathcal{V}^{-1} \circ (Id_\Gamma \cap Ind^2) \circ \overline{\mathcal{V}}\right)$. Apply now the results at the end of Section 2 and Lemma 12. ◀

▶ **Theorem 14** ([20]). *AbstractSafety is in* EXPSPACE.

Proof. Immediate consequence of Lemma 13, see the full version of the paper [16]. ◀

4 AbstractSafety is EXPSPACE-hard

In [19] it was shown that AbstractSafety was PSPACE-hard, and the paper left the question of closing the gap between the upper and lower bounds open. We first recall and slightly alter the PSPACE-hardness proof of [19], and then present our techniques to extend it to EXPSPACE-hardness.

The proof is by reduction from the problem of deciding whether a Turing machine $\mathcal{M}$ of size n does not accept when started on the empty tape of size n. (For technical reasons, we actually assume that the tape has $n-2$ cells.) Given $\mathcal{M}$, we construct in polynomial time an RTS $\mathcal{S}$ and a set of initial configurations C_I that, loosely speaking, satisfy the following two properties: the execution of $\mathcal{S}$ on an initial configuration simulates the run of $\mathcal{M}$ on the empty tape, and $PReach(C_I) = Reach(C_I)$. We choose C_U as the set of configurations of $\mathcal{S}$ in which $\mathcal{M}$ ends up in the accepting state. Then $\mathcal{S}$ is safe iff $\mathcal{M}$ does not accept.

Turing machine preliminaries. We assume that $\mathcal{M}$ is a deterministic Turing machine with states Q, tape alphabet Γ', initial state q_0 and accepting state q_f.

We represent a configuration of $\mathcal{M}$ as a word $\#\,\beta\, q\, \eta$ of length n, where $\mathcal{M}$ is in state q, the content of the tape is $\beta\,\eta \in \Gamma'^*$, and the head of $\mathcal{M}$ is positioned at the first letter of η. The symbol $\#$ serves as a separator between different configurations. The initial configuration is $\alpha_0 := \# q_0 B^{n-2}$, where B denotes the blank symbol of $\mathcal{M}$; so the tape is initially empty.

We assume w.l.o.g. that the successor of a configuration in state q_f is the configuration itself, so the run of $\mathcal{M}$ can be encoded as an infinite word $\boldsymbol{\alpha} := \alpha_0\alpha_1 \cdots$ where α_i represents the i-th configuration of $\mathcal{M}$. For convenience, we write $\Lambda := Q \cup \Gamma' \cup \{\#\}$ for the set of symbols in $\boldsymbol{\alpha}$. It is easy to see that the symbol at position $i+n$ of $\boldsymbol{\alpha}$ is completely determined by the symbols at positions $i-1$ to $i+2$ and the transition relation of $\mathcal{M}$. We let $\delta(x_1x_2x_3x_4)$ denote the symbol which "should" appear at position $i+n$ when the symbols at positions $i-1$ to $i+2$ are $x_1x_2x_3x_4$; in particular, $\delta(x_1\#x_2x_4) = \#$.

Configurations of $\mathcal{S}$. We choose the set of configurations as $\mathcal{C} := \alpha_0\#(\Lambda \cup \{\square\})^*$, and the initial configurations as $C_I := \alpha_0\#\square^*$. Intuitively, the RTS starts with the representation of the initial configuration of $\mathcal{M}$, followed by some number of blank cells $\square$. During its execution, the RTS will "write" the run of $\mathcal{M}$ into these blanks.

A configuration is unsafe if it contains some occurrence of q_f, the accepting state of $\mathcal{M}$, so $C_U := (\Lambda \cup \{\square\})^*\{q_f\}(\Lambda \cup \{\square\})^*$.

Transitions. For convenience, we will denote the i-th position of a word w as $w(i)$ instead of w_i. Given a configuration c, the set $\Delta(c)$ contains one single configuration c', defined as follows. Let i be some position of c such that $c_{i+n} = \square$. Then c' coincides with c everywhere except at position $i+n$, where instead $c'(i+n) := \delta(c(i-1)c(i)c(i+1)c(i+2))$. It is easy to see that Δ is a regular relation: The transducer nondeterministically guesses the position $i-1$, reads the next four symbols, say $x_1...x_4$, stores $\delta(x_1...x_4)$ in its state, moves to position $i+n$, checks if $c(i+n) = \square$ and writes $c'(i+n) := \delta(x_1...x_4)$. The transducer has $O(n^2)$ states.

It follows from the definitions above that $\mathcal{M}$ accepts the empty word iff $\mathcal{S}$ can reach C_U from C_I, i.e. $\mathit{Reach}(C_I) \cap C_U \neq \emptyset$.

Regular abstraction framework. We define a regular abstraction framework $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$ of polynomial size such that $\mathit{PReach}(C_I) = \mathit{Reach}(C_I)$. Hence, for every configuration $c \notin \mathit{Reach}(C_I)$, we must find an inductive constraint $A \in \mathcal{A}$ which separates C_I and c. (Note that C_I contains exactly one configuration of length $|c|$.)

As the reachable configurations are precisely the prefixes of $\boldsymbol{\alpha}$ with some symbols replaced by $\square$, there is a position i s.t. $c(i) \notin \{\square, \boldsymbol{\alpha}(i)\}$. Let us fix the smallest such i. As we noted above, $\boldsymbol{\alpha}(i)$ is determined entirely by $\boldsymbol{\alpha}(i-n-1)...\boldsymbol{\alpha}(i-n+2)$ via the mapping δ. So the constraint "if $c(i-n-1)...c(i-n+2) = x_1...x_4$, then $c(i) \in \{\square, \delta(x_1...x_4)\}$" is inductive and separates C_I and c.

Therefore, it is sufficient to define an abstraction framework in which every constraint of the above form can be expressed. This is relatively straightforward. We set $\mathcal{A} := \square^*\Lambda^4\square^*\Lambda\square^*$. Given a constraint $A = \square^i x_1...x_4\square^j x\square^k$, define $\mathcal{V}(A)$ as the set of all configurations c s.t. $c(i+1)...c(i+4) = x_1...x_4$ implies $c(i+j+5) \in \{\square, x\}$. Clearly, $\mathcal{V}$ is a regular relation which can be recognised by a transducer with 3 states.

▶ **Theorem 15** ([19]). *The abstract safety problem is* PSPACE*-hard, even for regular abstraction frameworks where the transducer for the interpretation has a constant number of states.*

From PSPACE-hardness to EXPSPACE-hardness

In order to prove EXPSPACE-hardness, we start with a machine $\mathcal{M}$ of size n and run it on a tape with 2^n cells. However, if we proceed exactly as in the PSPACE-hardness proof, we encounter two obstacles: (1) The length of α_0 is 2^n, so our definitions of $\mathcal{C}$ and C_I require automata of exponential size. (2) The transducer for the transition relation Δ needs to "count" to 2^n, as this is the distance between the corresponding symbols of α_i and α_{i+1}. Again, this requires an exponential number of states.

$$
\begin{array}{rlll}
 & 00 & 000 & \#^0\ q_0^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0 \\
\xrightarrow{mark(2,1)} & 01 & 000 & \#^1\ q_0^0\ \Box^1\ \Box^0\ \Box^1\ \Box^0\ \Box^1\ \Box^0\ \Box^1\ \Box^0 \\
\xrightarrow{mark(3,0)} & 01 & 100 & \#^1\ q_0^1\ \Box^1\ \Box^0\ \Box^1\ \Box^1\ \Box^1\ \Box^1\ \Box^1\ \Box^0 \\
\xrightarrow{init} & 00 & 000 & \#^0\ q_0^0\ \Box^0\ B^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0\ \Box^0 \\
\xrightarrow{mark(2,0)} & 10 & 000 & \#^0\ q_0^1\ \Box^0\ B^1\ \Box^0\ \Box^1\ \Box^0\ \Box^1\ \Box^0\ \Box^1 \\
\xrightarrow{mark(3,0)} & 10 & 100 & \#^0\ q_0^1\ \Box^1\ B^1\ \Box^1\ \Box^1\ \Box^0\ \Box^1\ \Box^1\ \Box^1 \\
\xrightarrow{init} & 00 & 000 & \#^0\ q_0^0\ \Box^0\ B^0\ \Box^0\ \Box^0\ \#^0\ \Box^0\ \Box^0\ \Box^0 \\
\xrightarrow{\dots} & 00 & 000 & \#^0\ q_0^0\ B^0\ B^0\ B^0\ B^0\ \#^0\ \Box^0\ \Box^0\ \Box^0 \\
\xrightarrow{\dots} & 10 & 001 & \#^1\ q_0^1\ B^0\ B^1\ B^1\ B^1\ \#^1\ \Box^1\ \Box^0\ \Box^1 \\
\xrightarrow{write} & 00 & 000 & \#^0\ q_0^0\ B^0\ B^0\ B^0\ B^0\ \#^0\ \Box^0\ q_1^1\ \Box^0 \\
\xrightarrow{\dots} & 01 & 010 & \#^1\ q_0^0\ B^1\ B^1\ B^1\ B^1\ \#^1\ \Box^0\ q_1^1\ \Box^1 \\
\xrightarrow{write} & 00 & 000 & \#^0\ q_0^0\ B^0\ B^0\ B^0\ B^0\ \#^0\ x^0\ q_1^0\ \Box^0
\end{array}
$$

Figure 3 A sample run of the regular transition system described in Example 16. Here, $mark(x, y)$ means that the y-th bit of the prime number x is changed to 1, and thus every position not equivalent to $y \pmod{x}$ is unmarked. Note that the first position of the TM part (the one with #) is position 0. We write x^y instead of $\left[\begin{smallmatrix} y \\ x \end{smallmatrix}\right]$. We highlight bits and symbols that were written to in pink (bits which are unmarked by the *mark* transition, but were already unmarked, are drawn in darker pink).

Obstacle (1) will be easy to overcome. Essentially, instead of starting the RTS with the entire initial configuration α_0 of $\mathcal{M}$ already in place, we set $C_I := \#\, q_0\, \Box^*$ and modify the transitions of $\mathcal{S}$ to also write out α_0.

However, obstacle (2) poses a more fundamental problem. On its face, it is easy to construct an RTS that can count to 2^n by executing multiple transitions in sequence, e.g. by implementing a binary counter. However, we need to balance this with the needs of the abstraction framework: if the RTS is too sophisticated, our constraints can no longer capture its behaviour using only regular languages.

We now sketch an RTS $\mathcal{S}'$ which extends the RTS $\mathcal{S}$ from the PSPACE-hardness proof.

A two-phase system. In order to write the run of $\mathcal{M}$, the RTS $\mathcal{S}'$ uses a "mark and write" approach. In a first phase, it executes n transitions to mark positions with distance m, where $m \geq 2^n$ is some fixed constant. Then, it nondeterministically guesses a marked position, reads and stores 4 symbols from that position, and moves to the next marker to write the symbol according to δ.

Let $p_1, \dots, p_n$ be the first n prime numbers (i.e. $p_1 = 2, p_2 = 3$, etc.). Define $m := \prod_{j=1}^{n} p_j$ and $s := \sum_{j=1}^{n} p_i$. We have $m \geq 2^n$ and, by the Prime Number Theorem, $s \in O(n^2 \log n)$.

The configurations of $\mathcal{S}'$ are of the form $w \left[\begin{smallmatrix} \mathbf{m} \\ c \end{smallmatrix}\right]$, where $w \in \{0,1\}^s$ stores the current state of the mark phase, $\mathbf{m} \in \{0,1\}^*$ are the markers (0 means marked), and $c \in (\Lambda \cup \{\Box\})^*$ is as for $\mathcal{S}$, with the reachable configurations being the prefixes of $\boldsymbol{\alpha}$ with some symbols replaced by $\Box$. We refer to $\left[\begin{smallmatrix} \mathbf{m} \\ c \end{smallmatrix}\right]$ as the *TM part*.

The RTS has three kinds of transitions: $\Delta' := \Delta_{mark} \cup \Delta_{write} \cup \Delta_{init}$.

$$\begin{array}{rcccccccccccccc} u = & 00 & 000 & \#^0\, q_0^0 & \square^0 & \square^0 & \square^0 & \square^0 & \square^0 & \square^0 & \square^0 & \square^0 \\ v = & 00 & 000 & \#^0\, q_0^0 & B^0 & B^0 & B^0 & B^0 & \#^0 & x^0 & q_1^0 & q_f^0 \end{array} \qquad (1)$$

$$\begin{array}{rcccccccccccccc} A_1 = & \square \ldots \square & \square & \square & B & B & B & B & \square & \square & \square & B \end{array} \qquad (2)$$

$$\begin{array}{rcccccccccccccc} A_2 = & 01 & 100 & 0 & 1 & 0 & 2 & 0 & 1 & 0 & 1 & 0 & 2 \end{array} \qquad (3)$$

Figure 4 Constraints in Example 16. (1) Two configurations u, v, where $u \in C_I$, $v \in C_U$. (2) The (not necessarily inductive) constraint A_1, separating u, v. (3) The matching inductive constraint A_2.

In the mark phase, $\mathcal{S}'$ executes a transition of Δ_{mark} for each $j \in [n]$. When executing such a transition, $\mathcal{S}'$ chooses a remainder $r \in [0, p_j - 1]$ and sets the corresponding bit in w. It then unmarks every position in the TM part which is *not* equivalent to r modulo p_j (by replacing the 0 with a 1). Hence, after executing n transitions in Δ_{mark}, the positions of all 0's in the TM part are equivalent modulo every p_j. By the Chinese remainder theorem, these positions must also be equivalent modulo m.

Afterwards, $\mathcal{S}'$ executes either a transition in Δ_{write} or Δ_{init}. To execute Δ_{write}, the RTS nondeterministically guesses a marked position i, reads $x := c(i-1)...c(i+2)$, moves to the next marked position i', and writes $\delta(x)$.

As mentioned in obstacle (1) above, the RTS must write out the initial configuration of $\mathcal{M}$. This is done by Δ_{init}. If the first position of the TM part is not marked, the transducer moves to the first marked position and writes B, otherwise it moves to the second marked position and writes #. By executing this transition multiple times, eventually a configuration $w\left[\begin{smallmatrix}\mathbf{m}\\ c\end{smallmatrix}\right]$ with $c = \#q_0B^{m-2}\#\square^i$ can be reached.

While executing either Δ_{write} or Δ_{init}, the transducer resets the mark phase state and marks all positions, i.e. the resulting configurations have $w = 0^s$ and $\mathbf{m} \in 0^*$.

▶ **Example 16.** Take $n = 2$. Here, we have $p_1 = 2$, $p_2 = 3$, $m = 6$ and $s = 5$. The set of initial and unsafe configurations is thus $C_I := L(0^5[\begin{smallmatrix}0\\ \#\end{smallmatrix}][\begin{smallmatrix}0\\ q_0\end{smallmatrix}][\begin{smallmatrix}0\\ \square\end{smallmatrix}]^*)$ and $C_U := \{0\}^5(\{0\} \times \Lambda)^*\{[\begin{smallmatrix}0\\ q_f\end{smallmatrix}]\}(\{0\} \times \Lambda)^*$, respectively. In Figure 3, we give a possible run of the RTS for a TM with states $\{q_0, q_1, q_f\}$ (q_0 is initial, q_f is final), and one transition from q_0, which reads B, moves the head to the right and goes to state q_1.

The abstraction framework. If $\mathcal{M}$ accepts, no constraint proving safety can exist, as an unsafe configuration is reachable. Consequently, when constructing the abstraction framework we only need to ensure that – provided $\mathcal{M}$ does not accept – for every pair $(u, v) \in C_I \times C_U$ there is an inductive constraint separating u and v.

The abstraction framework $(\mathcal{C}, \mathcal{A}, \mathcal{V})$ is the convolution of two independent parts, i.e. $\mathcal{A} := \mathcal{A}_1 \times \mathcal{A}_2$ and $\mathcal{V}([\begin{smallmatrix}A_1\\ A_2\end{smallmatrix}]) := \mathcal{V}_1(A_1) \cap \mathcal{V}_2(A_2)$.

For every pair $(u, v) \in C_I \times C_U$ there will be a constraint $A_1 \in \mathcal{A}_1$ separating u and v. This is similar to before: v must contain an "error" somewhere, so our constraint will state "if $c(i-1)...c(i+2) = x$, then $c(i+m) \in \{\delta(x), \square\}$", for some i, x. (Depending on v we instead may need A_1 stating just "$c(i) \in \{\boldsymbol{\alpha}(i), \square\}$".) Concretely, we set $\mathcal{A}_1 := \square^s\square^*(\Lambda^4\square^*\Lambda + \Lambda)\square^*$, so the constraint is represented by a word in $\square^*x\square^*\delta(x)\square^*$ (or a word in $\square^*\boldsymbol{\alpha}(i)\square^*$). An example is shown in Figure 4.

This is enough to separate u and v, as v must contain an "error" somewhere (i.e. a deviation from $\boldsymbol{\alpha}$). But it is not inductive: We can take any configuration which has $c(i+m) = \square$, but where the cells have not been marked correctly, s.t. executing Δ_{write} would write to position $i + m$ after reading symbols $c(j-1)...c(j+2)$ with $j \neq i$. So the resulting configuration may have $c(i+m) \neq \delta(x)$, which no longer fulfils A_1.

We solve this issue via $\mathcal{V}_2$. For the constraint A_1 above there is going to be a constraint $A_2 \in \mathcal{A}_2$ s.t. the combination $\mathcal{V}_1(A_1) \cap \mathcal{V}_2(A_2)$ is inductive. Essentially, A_2 will ensure that it is impossible to write to position $i + m$ without reading from position i. (Note that for a particular constraint the value of i is fixed.)

Let $\mathcal{A}_2 := \{0,1\}^s[0,n]^*$. Intuitively, a constraint $x\,y \in \mathcal{A}_2$ (where $|x| = s$) states: "if remainders for the first j primes have been chosen according to x, then exactly the positions k with $y(k) \geq j$ are marked, otherwise positions k with $y(k) = n$ are unmarked", where j is the number of primes that have been chosen.

Again, the constraint A_1 is only concerned with one position i. Moreover, there is only one sequence of remainders $r_1, ..., r_n$ to choose for the Δ_{mark} transitions, s.t. position i is marked (i.e. $r_j \equiv i \pmod{p_j}$). So for each position k we can determine the index in the sequence of Δ_{mark} transitions at which position k will first become unmarked. Concretely, we have $y(k) := \min\{j \mid k \not\equiv i \pmod{p_j}\} - 1$.

This constraint is inductive and, crucially, the intersection of A_1 and A_2 is inductive as well. Essentially, every Δ_{mark} transition either continues the sequence $r_1, ..., r_n$, and then the positions must be marked precisely according to y, or at some point a different remainder has been chosen, and the position i is unmarked and cannot be written to.

To summarise, constraint A_1 is sufficient to exclude any unsafe configuration and, in combination with A_2, does so inductively. Therefore, if $\mathcal{M}$ does not accept, then the RTS can be proven safe using the abstraction framework.

For the full proof, see the full version of the paper [16].

5 Learning regular sets of inductive constraints

Recall the algorithm for AbstractSafety underlying Theorem 14. It computes an automaton recognising the set *Ind* of inductive constraints (Lemma 12); uses this automaton to compute a transducer recognising the potential reachability relation *PReach* (Lemma 13); uses this transducer to compute an automaton recognising $\mathit{PReach}(C_I) \cap C_U$; and finally uses this automaton to check if $\mathit{PReach}(C_I) \cap C_U$ is empty (Theorem 14). The main practical problem of this approach is that, while the automaton for $\overline{\mathit{Ind}}$ has polynomial size in the input, the automaton for *Ind* can be exponential, and, while the automaton for $\overline{\mathit{PReach}}$ has polynomial size in *Ind*, the size of the automaton for *PReach* can be exponential.

In practice one typically does not need all inductive constraints to prove safety. This can be illustrated even on the tiny RTS of Example 2.

▶ **Example 17.** Consider the RTS of the token passing system of Example 2, where $\Sigma = \{t, n\}$, and the second abstraction framework of Example 10, where $\Gamma = 2^\Sigma = \{\emptyset, \{t\}, \{n\}, \{t, n\}\}$. Recall that in this abstraction framework a constraint is an *exclusive disjunction* of literals, that is, a configuration satisfies the constraint if it satisfies *exactly one* of its literals. The minimal DFA recognising all non-trivial inductive constraints was shown on the left of Figure 2. The set of inductive constraints $\{t\}\{t\}^*$ is satisfied by the configurations n^*tn^*, and so the DFA on the right is already strong enough to prove any safety property.

We present a learning algorithm that computes automata recognising increasingly large sets $\mathsf{H} \subseteq \mathit{Ind}$ of inductive constraints until either H is large enough to prove safety, or it becomes clear that even the whole set *Ind* is not large enough. More precisely, recall that, by definition, we have $\mathit{PReach} := \{(c, c') \in \mathcal{C} \times \mathcal{C} \mid \forall A \in \mathit{Ind}\colon c \in \mathcal{V}(A) \to c' \in \mathcal{V}(A)\}$. Given a set $\mathsf{H} \subseteq \mathit{Ind}$, define the relation $\mathit{PReach}_\mathsf{H}$ exactly as *PReach*, just replacing *Ind* by H. Clearly, we have $\mathit{PReach}_\mathsf{H} \supseteq \mathit{PReach}$ and $\mathit{PReach}_{\mathit{Ind}} = \mathit{PReach}$.

5.1 The learning algorithm

Let $\mathcal{S} = (\mathcal{C}, \Delta)$ and $\mathcal{F} = (\mathcal{C}, \mathcal{A}, \mathcal{V})$ be a regular transition system and a regular abstraction framework, respectively. Further, let C_I and C_U be regular sets of initial and unsafe configurations. The algorithm refines Angluin's algorithm L^* for learning a DFA for the full set *Ind* [7, 30]. Recall that Angluin's algorithm involves two agents, usually called Learner and Teacher. Learner sends Teacher membership and equivalence queries, which are answered by Teacher according to the following specification:

Membership Query:
- Input: a constraint $A \in \mathcal{A}$
- Output: ✓ if $A \in \mathit{Ind}$, and $\times$ otherwise.

Equivalence Query:
- Input: a DFA recognising a set $\mathsf{H} \subseteq \mathcal{A}$.
- Output: ✓ if $\mathsf{H} = \mathit{Ind}$, otherwise a constraint $A \in (\mathsf{H} \setminus \mathit{Ind}) \cup (\mathit{Ind} \setminus \mathsf{H})$.

Angluin's algorithm describes a strategy for Learner guaranteeing that Learner eventually learns the minimal DFA recognising *Ind*. The number of equivalence queries asked by Learner is at most the number of states of the DFA.

Answering the queries. We describe the algorithms used by Teacher to answer queries. For membership queries, Teacher constructs an NFA for $\overline{\mathit{Ind}} \cap \{A\}$ with $O(|A| \cdot n_\Delta \cdot n_\mathcal{V}^2)$ states (see Lemma 12), and checks it for emptiness.

For equivalence queries, Teacher proceeds as follows :

1. Teacher first checks whether $\mathsf{H} \setminus \mathit{Ind} \neq \emptyset$ holds by computing an NFA recognising $\mathsf{H} \cap \overline{\mathit{Ind}}$ with $O(n_\mathsf{H} \cdot n_\Delta \cdot n_\mathcal{V}^2)$ states (see Lemma 12), and checking it for emptiness. If $\mathsf{H} \setminus \mathit{Ind}$ is nonempty, then Teacher returns one of its elements.
2. Otherwise, Teacher constructs an automaton for $\mathit{PReach}_\mathsf{H}(C_I) \cap C_U$ of size $O(2^{n_\mathcal{V}^2 \cdot n_\mathsf{H}})$ and checks it for emptiness. There are two cases:
   a. If $\mathit{PReach}_\mathsf{H}(C_I) \cap C_U = \emptyset$, then the system is safe; Teacher reports it and terminates. In this case, the learning algorithm is aborted without having learned a DFA for *Ind*, because it is no longer necessary.
   b. Otherwise, Teacher chooses an element $(c, c') \in \mathit{PReach}_\mathsf{H} \cap (C_I \times C_U)$, and searches for an inductive constraint A such that $c \in \mathcal{V}(A)$ and $c' \notin \mathcal{V}(A)$. We call this problem the *separability problem*, and analyze it further in Section 5.2.

5.2 The separability problem

The Separability problem is formally defined as follows:

Given: a nondeterministic transducer recognising a regular relation $\Delta \subseteq \Sigma^* \times \Sigma^*$; a deterministic transducer recognising a regular interpretation $\mathcal{V}$ over $\Gamma \times \Sigma$; and two configurations $c, c' \in \mathcal{C}$

Decide: is c' separable from c, i.e. does there exist $A \in \mathit{Ind}$ s.t. $c \in \mathcal{V}(A)$ and $c' \notin \mathcal{V}(A)$?

Contrary to AbstractSafety, the complexity of Separability is different for arbitrary transducers, and for length-preserving ones.

▶ **Theorem 18.** *Separability is* PSPACE*-complete, even if Δ is length-preserving. If $\mathcal{V}$ is length-preserving, then Separability is* NP*-complete.*

Proof. See the full version of the paper [16]. ◀

Table 1 Comparison of the sizes of the automata computed by the lazy and direct approaches. In each table, the first three columns contain the name of the RTS and the sizes of the automata for C_I and Δ. The fourth column (Pr.) indicates the checked property, where D, M, and O stand for "deadlock freedom", "mutex" (at most one process in a given state), and "other" (custom properties of the particular RTS). The next two columns give the results for the lazy approach: sizes of the DFAs for H and $PReach_{\mathsf{H}}$ (abbreviated as PR_{H}), and the next two the same results for the direct approach. The last column (Re.) indicates the result of the check: the property could be proved (✓), could not (×), or, in the case of multiple properties, how many of the properties were proved.

				Lazy		Direct		
System	$\|C_I\|$	$\|\Delta\|$	Pr.	$\|\mathsf{H}\|$	$\|PR_{\mathsf{H}}\|$	$\|Ind\|$	$\|PR\|$	Re.
Bakery	3	5	D	1	1	9	8	✓
			M	4	3			✓
Burns	1	6	D	1	1	10	6	✓
			M	5	3			✓
Dijkstra	2	17	D	4	4	218	22	✓
			M	11	8			✓
Dijkstra (ring)	2	12	D	9	9	47	17	✓
			M	9	7			×
Dining crypto.	2	8	D	23	18	86	19	2/2
Herman	2	11	D	3	2	8	7	✓
			O	1	2			✓
Herman (linear)	2	3	D	1	2	7	7	×
			O	1	2			✓
Israeli-Jafon	3	10	D	1	4	21	7	✓
			O	1	4			✓
Token passing	2	3	O	4	4	9	7	✓
Lehmann-Rabin	1	7	D	5	6	29	13	✓
LR phils. 1	1	11	D	13	14	29	15	×
LR phils. 2	1	11	D	25	11	29	9	✓
Atomic phils.	1	8	D	13	9	22	20	✓
Mux array	2	4	D	1	2	7	8	✓
			M	3	5			×
Res. alloc.	1	5	D	5	5	9	8	✓
			M	3	3			×

				Lazy		Direct		
System	$\|C_I\|$	$\|\Delta\|$	Pr.	$\|\mathsf{H}\|$	$\|PR_{\mathsf{H}}\|$	$\|Ind\|$	$\|PR\|$	Re.
Berkeley	1	9	D	1	1	12	9	✓
			O	4	4			2/3
Dragon	1	23	D	1	1	37	11	✓
			O	15	7			6/7
Firefly	1	16	D	1	1	12	7	✓
			O	4	3			0/4
Illinois	1	16	D	1	1	18	14	✓
			O	4	3			0/2
MESI	1	7	D	1	1	8	7	✓
			O	4	4			2/2
MOESI	1	7	D	1	1	15	10	✓
			O	4	4			7/7
Synapse	1	5	D	1	1	8	7	✓
			O	2	3			2/2

				Lazy		Direct		
System	$\|C_I\|$	$\|\Delta\|$	Pr.	$\|\mathsf{H}\|$	$\|PR_{\mathsf{H}}\|$	$\|Ind\|$	$\|PR\|$	Re.
Dijkstra (ring)	2	12	M	9	7			✓
LR phils. 1	1	11	D	34	11			✓
Mux array	2	4	M	5	3			✓
Res. alloc.	1	5	M	5	5			✓

Most applications of regular model checking to the verification of parameterised systems, and in particular all the examples studied in [19, 20], have length-preserving transition functions and length-preserving interpretations. For this reason, in our implementation we only consider this case, and leave an extension for future research. Since SEPARABILITY is NP-complete in the length-preserving case, it is natural to solve it by reduction to SAT. A brief description of the reduction is given in the full version of the paper [16].

5.3 Some experimental results

We have implemented the learning algorithm in a tool prototype, built on top of the libraries `automatalib` and `learnlib` [22] and the SAT solver `sat4j` [9]. We compare our learning approach with the one of [19], which constructs automata for *Ind* and *PReach* using the regular abstraction framework of Example 7. In the rest of this section we call these two approaches the *lazy* and the *direct* approach, respectively. We use the same case studies as [19]. We compare the sizes of the DFA for the final hypothesis H and $PReach_{\mathsf{H}}$ with the sizes of the DFA for *Ind* and *PReach*. The results are available at [33] and are shown in Table 1.

The left table in Table 1 shows results on RTSs modeling mutex and leader election algorithms, and academic examples, like various versions of the dining philosophers. The right top table shows results on models of cache-coherence protocols. Observe that *Ind* and

PReach do not depend on the property, but H and $PReach_{\mathsf{H}}$ do, because the algorithm can finish early. In this case, the sizes given in columns H and $PReach_{\mathsf{H}}$ are the largest ones computed over all properties checked.

The main result is that the automata computed by our tool are significantly smaller than those for [19]. (Note that in all cases we compute minimal DFAs, and so the differences are not due to algorithms for the computation of automata.) Observe that in five cases the deadlock-freedom and the mutex properties could not be proved. In one case (deadlock-freedom of Herman (linear)) this is because the property does not hold. In the other four cases, the problem is that [19] uses only a specific regular abstraction framework, namely the one of Example 7. We can prove the property by refining the abstraction: we take the union of the "disjunctive" and the "exclusive disjunctive" abstractions of Example 10. The bottom-right table gives the results of these four cases.

Both tools take less than three seconds in 54 out of the 59 case studies in the left and top right tables. We do not report the exact times; the implementation of [19] uses MONA, while the experiments of this paper use `automatalib` and `learnlib`, and so small time differences may have any number of reasons. In the other five cases, the implementation of [19] still needs less than one second, while our implementation takes minutes (more than ten minutes in two cases). In these five cases the time performance is dominated by the SAT solver `sat4j`. We have not yet identified a pattern explaining why `sat4j` takes so much time, in particular the number and size of the formulas passed to it is similar to the other cases.

6 Conclusions

We have generalised the technique of [19, 20] for checking safety properties of RTS to arbitrary regular abstraction frameworks. We have shown that the abstract safety problem is EXPSPACE-complete, solving an open problem of [19, 20], by means of a complex reduction of independent interest. For particular abstraction frameworks the complexity can be better.

We have used automata learning to design a lazy algorithm that stops when the inductive constraints computed so far are enough to prove safety. Its combination with other learning techniques, as those proposed in [28, 31, 15, 32, 29], is a question for future research.

References

1 Parosh Aziz Abdulla. Regular model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):109–118, 2012.

2 Parosh Aziz Abdulla. Regular model checking: Evolution and perspectives. In *Model Checking, Synthesis, and Learning*, volume 13030 of *Lecture Notes in Computer Science*, pages 78–96. Springer, 2021.

3 Parosh Aziz Abdulla, Frédéric Haziza, and Lukás Holík. Parameterized verification through view abstraction. *Int. J. Softw. Tools Technol. Transf.*, 18(5):495–516, 2016.

4 Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Julien d'Orso. Regular model checking made simple and efficient. In *CONCUR*, volume 2421 of *Lecture Notes in Computer Science*, pages 116–130. Springer, 2002.

5 Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Mayank Saksena. A survey of regular model checking. In *CONCUR*, volume 3170 of *Lecture Notes in Computer Science*, pages 35–48. Springer, 2004.

6 Parosh Aziz Abdulla, A. Prasad Sistla, and Muralidhar Talupur. Model checking parameterized systems. In *Handbook of Model Checking*, pages 685–725. Springer, 2018.

7 Dana Angluin. Learning regular sets from queries and counterexamples. *Inf. Comput.*, 75(2):87–106, 1987.

8 Clark W. Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Model Checking*, pages 305–343. Springer, 2018.

9 Daniel Le Berre and Anne Parrain. The sat4j library, release 2.2. *J. Satisf. Boolean Model. Comput.*, 7(2-3):59–6, 2010.

10 Bernard Boigelot, Axel Legay, and Pierre Wolper. Iterating transducers in the large (extended abstract). In *CAV*, volume 2725 of *Lecture Notes in Computer Science*, pages 223–235. Springer, 2003.

11 Ahmed Bouajjani, Peter Habermehl, Adam Rogalewicz, and Tomás Vojnar. Abstract regular (tree) model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):167–191, 2012.

12 Ahmed Bouajjani, Peter Habermehl, and Tomás Vojnar. Abstract regular model checking. In *CAV*, volume 3114 of *Lecture Notes in Computer Science*, pages 372–386. Springer, 2004.

13 Ahmed Bouajjani, Bengt Jonsson, Marcus Nilsson, and Tayssir Touili. Regular model checking. In *CAV*, volume 1855 of *Lecture Notes in Computer Science*, pages 403–418. Springer, 2000.

14 Ahmed Bouajjani and Tayssir Touili. Widening techniques for regular tree model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):145–165, 2012.

15 Yu-Fang Chen, Chih-Duo Hong, Anthony W. Lin, and Philipp Rümmer. Learning to prove safety over parameterised concurrent systems. In *FMCAD*, pages 76–83. IEEE, 2017.

16 Philipp Czerner, Javier Esparza, Valentin Krasotin, and Christoph Welzel-Mohr. Computing inductive invariants of regular abstraction frameworks. *CoRR*, abs/2404.10752, 2024. `doi:10.48550/arXiv.2404.10752`.

17 Dennis Dams, Yassine Lakhnech, and Martin Steffen. Iterating transducers. In *CAV*, volume 2102 of *Lecture Notes in Computer Science*, pages 286–297. Springer, 2001.

18 Javier Esparza and Michael Blondin. *Automata Theory – An Algorithmic Approach*. MIT Press, 2023.

19 Javier Esparza, Mikhail A. Raskin, and Christoph Welzel. Regular model checking upside-down: An invariant-based approach. In *CONCUR*, volume 243 of *LIPIcs*, pages 23:1–23:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

20 Javier Esparza, Mikhail A. Raskin, and Christoph Welzel. Regular model checking upside-down: An invariant-based approach. *CoRR*, abs/2205.03060, version 2, 2022.

21 Chih-Duo Hong and Anthony W. Lin. Regular abstractions for array systems. *Proc. ACM Program. Lang.*, 8(POPL):638–666, 2024.

22 Malte Isberner, Falk Howar, and Bernhard Steffen. The open-source learnlib - A framework for active automata learning. In *CAV (1)*, volume 9206 of *Lecture Notes in Computer Science*, pages 487–495. Springer, 2015.

23 Bengt Jonsson and Marcus Nilsson. Transitive closures of regular relations for verifying infinite-state systems. In *TACAS*, volume 1785 of *Lecture Notes in Computer Science*, pages 220–234. Springer, 2000.

24 Yonit Kesten, Oded Maler, Monica Marcus, Amir Pnueli, and Elad Shahar. Symbolic model checking with rich assertional languages. *Theor. Comput. Sci.*, 256(1-2):93–112, 2001.

25 Daniel Kroening and Ofer Strichman. *Decision Procedures - An Algorithmic Point of View, Second Edition*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2016.

26 Axel Legay. Extrapolating (omega-)regular model checking. *Int. J. Softw. Tools Technol. Transf.*, 14(2):119–143, 2012.

27 Antoine Miné. The octagon abstract domain. *High. Order Symb. Comput.*, 19(1):31–100, 2006.

28 Daniel Neider. *Applications of automata learning in verification and synthesis*. PhD thesis, RWTH Aachen University, 2014.

29 Daniel Neider and Nils Jansen. Regular model checking using solver technologies and automata learning. In *NASA Formal Methods*, volume 7871 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2013.

30 Bernhard Steffen, Falk Howar, and Maik Merten. Introduction to active automata learning from a practical perspective. In *SFM*, volume 6659 of *Lecture Notes in Computer Science*, pages 256–296. Springer, 2011.

31 Abhay Vardhan. *Learning To Verify Systems*. PhD thesis, University of Illinois, 2006.

32 Abhay Vardhan, Koushik Sen, Mahesh Viswanathan, and Gul Agha. Learning to verify safety properties. In *ICFEM*, volume 3308 of *Lecture Notes in Computer Science*, pages 274–289. Springer, 2004.

33 Christoph Welzel-Mohr, Valentin Krasotin, Philipp Czerner, and Javier Esparza. dodo, February 2024. `doi:10.5281/zenodo.12734991`.