



On an Invariance Problem for Parameterized Concurrent Systems

Marius Bozga, Lucas Bueri, Radu Iosif

► To cite this version:

Marius Bozga, Lucas Bueri, Radu Iosif. On an Invariance Problem for Parameterized Concurrent Systems. International Conference on Concurrency Theory 2022, Sep 2022, Varsovie, Poland. hal-03826296

HAL Id: hal-03826296

<https://hal.science/hal-03826296>

Submitted on 24 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On an Invariance Problem for Parameterized Concurrent Systems

Marius Bozga and Lucas Bueri and Radu Iosif

Univ. Grenoble Alpes, CNRS, Grenoble INP, VERIMAG, 38000, France

Abstract

We consider concurrent systems consisting of replicated finite-state processes that synchronize via joint interactions in a network with user-defined topology. The system is specified using a resource logic with a multiplicative connective and inductively defined predicates, reminiscent of Separation Logic [19]. The problem we consider is if a given formula in this logic defines an invariant, namely whether any model of the formula, following an arbitrary firing sequence of interactions, is transformed into another model of the same formula. This property, called *havoc invariance*, is quintessential in proving the correctness of reconfiguration programs that change the structure of the network at runtime. We show that the havoc invariance problem is many-one reducible to the entailment problem $\phi \models \psi$, asking if any model of ϕ is also a model of ψ . Although, in general, havoc invariance is found to be undecidable, this reduction allows to prove that havoc invariance is in 2EXP, for a general fragment of the logic, with a 2EXP entailment problem.

2012 ACM Subject Classification • Software and its engineering → Formal software verification

Keywords and phrases parameterized verification, invariant checking, resource logics, reconfigurable systems, tree automata

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

The parameterized verification problem asks to decide whether a system consisting of an arbitrary number of finite-state processes that communicate via synchronized (joint) actions satisfies a specification, such as deadlock freedom, mutual exclusion or a temporal logic property e.g., every request is eventually answered. The literature in this area has a wealth of decidability and complexity results (see [3] for a survey) classified according to the communication type (e.g., rendez-vous, broadcast) and the network topology e.g., rings where every process interacts with its left/right neighbours, cliques where each two process may interact, stars with a controller interacting with unboundedly many workers, etc.

As modern computing systems are dynamically adaptive, recent effort has been put into designing *reconfigurable systems*, whose network topologies change at runtime (see [12] for a survey) in order to address maintenance (e.g., replacement of faulty and obsolete components by new ones, firmware updates, etc.) and internal traffic issues (e.g., re-routing to avoid congestion in a datacenter [18]). Unfortunately the verification of dynamic reconfigurable systems (i.e., proving the absence of design errors) remains largely unexplored. Consequently, such systems are prone to bugs that may result in e.g., denial of services or data corruption¹.

Proving correctness of parameterized reconfigurable networks is tackled in [1], where a Hoare-style program logic is proposed to write proofs of reconfiguration programs i.e., programs that dynamically add and remove processes and interactions from the network during runtime. The assertion language used by these proofs is a logic that describes sets of configurations defining the network topology and the local states of the processes. The logic views processes and interactions as resources that can be joined via a *separating conjunction*,

¹ Google reports on a cascading cloud failure due to reconfiguration: <https://status.cloud.google.com/incident/appengine/19007>



in the spirit of Separation Logic [19]. The separating conjunction supports *local reasoning*, which is the ability of describing reconfigurations *only with respect to those components and interactions that are involved in the mutation*, while disregarding the rest of the system's configuration. Moreover, the separating conjunction allows to concisely describe networks of unbounded size, that share a similar architectural style (e.g., pipelines, rings, stars, trees) by means of inductively defined predicates.

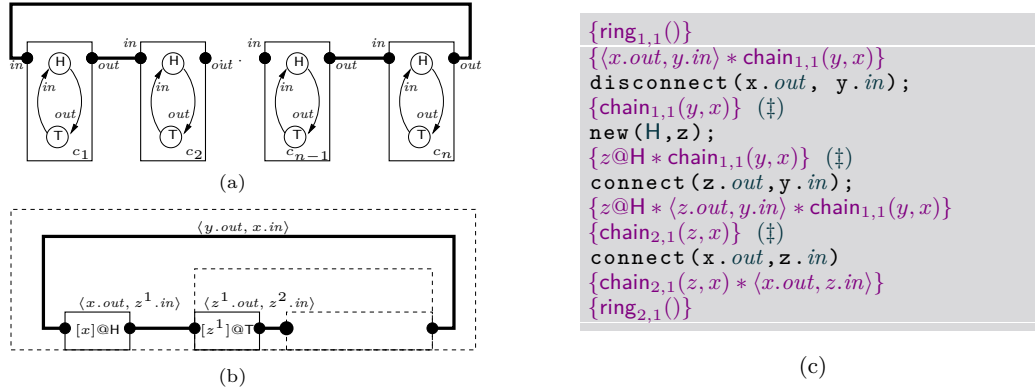
Due to the interleaving of reconfigurations and interactions between components, the annotations of the reconfiguration program form a valid proof under so-called *havoc invariance* assumptions, stating global properties about the configurations, that remain, moreover, *unchanged under the ongoing interactions in the system*. These assumptions are needed to apply the sequential composition rule that infers a Hoare triple $\{\phi\} P; Q \{\psi\}$ from two premisses $\{\phi\} P \{\theta\}$ and $\{\theta\} Q \{\psi\}$, where P and Q are reconfiguration actions that add and/or remove processes and communication channels. Essentially, because the states of the processes described by the intermediate assertion θ might change between the end of P and the beginning of Q , this rule is sound provided that θ is a havoc invariant formula.

This paper contributes to the automated generation of reconfiguration proofs, by a giving a procedure that discharges the havoc invariance side conditions. The challenge is that a formula of the configuration logic (that contains inductively defined predicates) describes an infinite set of configurations of arbitrary sizes. The main result is that the havoc invariance problem is effectively many-one reducible to the entailment problem $\phi \models \psi$, that asks if every model of a formula ϕ is a model of another formula ψ . Here ψ is the formula whose havoc invariance is being checked and ϕ defines the set of configurations γ' obtained from a model γ of ψ , by executing one interaction from γ . The reduction is polynomial if certain parameters are bound by a constant (i.e., the arity of the predicates, the size of interactions and the number of predicate atoms is an inductive rule), providing a 2EXP upper bound for a fragment of the logic with a decidable (2EXP) entailment problem [4, §6]. Having a polynomial reduction motivates, moreover, future work on the definition of fragments of lower (e.g., polynomial) entailment complexity (see e.g., [9] for a fragment of Separation Logic with a polynomial entailment problem), that are likely to yield efficient decision procedures for the havoc invariance problem as well. In addition, we provide a 2EXP-hard lower bound for the havoc invariance problem in this fragment of the logic (i.e., assuming predicates of unbounded arity) and show that havoc invariance is undecidable, when unrestricted formulæ are considered as input. For space reasons, the technical proofs are given in Appendix A.

Related Work Specifying parameterized concurrent systems by inductive definitions is reminiscent of *network grammars* [20, 16, 13], that use inductive rules to describe systems with linear (pipeline, token-ring) architectures obtained by composition of an unbounded number of processes. In contrast, we use predicates of unrestricted arities to describe network topologies that can be, in general, more complex than trees. Moreover, we write inductive definitions using a resource logic, suitable also for writing Hoare logic proofs of reconfiguration programs, based on local reasoning [8].

Verification of network grammars against safety properties (reachability of error configurations) requires the synthesis of *network invariants* [21], computed by rather costly fixpoint iterations [17] or by abstracting (forgetting the particular values of indices in) the composition of a small bounded number of instances [14]. In previous work, we have developed an invariant synthesis method based on *structural invariants*, that are synthesized with little computational effort and prove to be efficient in many practical examples [5, 6].

The havoc invariance problem considered in this paper is, however, different from safety checking and has not been addressed before, to the best of our knowledge. An explanation is that verification of reconfigurable systems has received fairly scant attention, relying mostly



■ **Figure 1** Inductive Specification and Reconfiguration of a Token Ring

on runtime verification [7, 10, 15, 11], instead of deductive verification, reported in [1]. In [1] we addressed havoc invariance with a set of inference rules used to write proofs manually, whereas the goal of this paper is to discharge such conditions automatically.

1.1 A Motivating Example

Consider, for instance, a system consisting of a finite but unbounded number of processes, called *components* in the following. The components execute the same machine with states T and H , denoting whether the component has a token (T) or a hole (H). The components are placed in a ring, each component having exactly one left and one right neighbour, as in Fig. 1 (a). A component without a token may receive one, by executing a transition $H \xrightarrow{in} T$, simultaneously with its left neighbour, that executes the transition $T \xrightarrow{out} H$ simultaneously, as in Fig. 1 (a). Note that there can be more than one token, moving independently in the system, such that no token overtakes another token. The configurations of the token ring system are described by the following inductive rules:

$$\begin{aligned}
 \text{ring}_{h,t}() &\leftarrow \exists x \exists y . \langle x.out, y.in \rangle * \text{chain}_{h,t}(y, x) \\
 \text{chain}_{h,t}(x, y) &\leftarrow \exists z. [x]@q * \langle x.out, z.in \rangle * \text{chain}_{h',t'}(z, y), \text{ for both } q \in \{H, T\} \\
 \text{chain}_{0,1}(x, x) &\leftarrow [x]@T \quad \text{chain}_{1,0}(x, x) \leftarrow [x]@H \quad \text{chain}_{0,0}(x, x) \leftarrow [x] \\
 \text{where } h' &\stackrel{\text{def}}{=} \begin{cases} \max(h-1, 0) & , \text{ if } q = H \\ h & , \text{ if } q = T \end{cases} \quad \text{and } t' \stackrel{\text{def}}{=} \begin{cases} \max(t-1, 0) & , \text{ if } q = T \\ t & , \text{ if } q = H \end{cases}
 \end{aligned}$$

The predicate $\text{ring}_{h,t}()$ describes a ring with at least h (t) components in state H (T). The ring consists of an interaction between the ports *out* and *in* of two components x and y , respectively, described by $\langle x.out, y.in \rangle$ and a separate chain of components between x and y , described by $\text{chain}_{h,t}(y, x)$. Inductively, a chain consists of a component $[x]@q$ in state $q \in \{H, T\}$, an interaction $\langle x.out, z.in \rangle$ and a separate $\text{chain}_{h',t'}(z, y)$, where h' and t' are the least numbers of components in state H and T , respectively, after the removal of the component x . Fig. 1 (b) depicts the unfolding of the inductive definition of $\text{ring}_{h,t}()$ with the existentially quantified variables z from the above rules α -renamed to z^1, z^2 , etc.

A *reconfiguration action* is an atomic creation or deletion of a component or interaction. A *reconfiguration sequence* is a finite sequence of reconfiguration actions that takes as input a mapping of program variables to components and executes the actions from the sequence, interleaving with the interactions in the system. For instance, the reconfiguration sequence from Fig. 1 (c) takes as input the mapping of x and y to two adjacent components in the token ring, removes the interaction $\langle x.out, y.in \rangle$ by executing $\text{disconnect}(x.out, y.in)$ and creates a new component in state H (by executing $\text{new}(x, H)$) that is connected in between x and y

via two new interactions created by executing `connect(z.out, y.in)` and `connect(x.out, z.in)`, respectively. Fig. 1 (c) shows a proof (with annotations in curly braces) of the fact that outcome of the reconfiguration of a ring of components is a ring whose least number of components in state H is increased from one to two. This proof is split into several subgoals:

1. Entailments required to apply the consequence rule of Hoare logic e.g., $\text{ring}_{1,1}() \models \exists x \exists y . \langle x.out, y.in \rangle * \text{chain}_{1,1}(y, x)$. The entailment problem has been addressed in [4, §6], with the definition of a general fragment of the configuration logic, for which the entailment problem is decidable in double exponential time.
2. Hoare triples that describe the effect of the atomic reconfiguration actions e.g., $\{\langle x.out, y.in \rangle * \text{chain}_{1,1}(y, x)\} \text{disconnect}(x.out, y.in) \{\text{chain}_{1,1}(y, x)\}$. These are obtained by applying the frame rule to the local² specifications of the atomic actions. The local specification of reconfiguration actions and the frame rule for local actions are described in [1, §4.2].
3. *Havoc invariance* proofs for the annotations marked with (\dagger) in Fig. 1 (c). For instance, the formula $\text{chain}_{1,1}(y, x)$ is havoc invariant because the interactions in a chain of components will only move tokens to the right without creating more or losing any, hence there will be the same number of components in state H (T) no matter which interactions are fired.

2 Definitions

We denote by \mathbb{N} the set of positive integers, including zero. For a set A , we denote $A^1 \stackrel{\text{def}}{=} A$, $A^{i+1} \stackrel{\text{def}}{=} A^i \times A$, for all $i \geq 0$, where \times denotes the Cartesian product, and $A^+ \stackrel{\text{def}}{=} \bigcup_{i \geq 1} A^i$. The cardinality of a finite set A is denoted by $\|A\|$. By writing $A \subseteq_{fin} B$ we mean that A is a finite subset of B . Given integers i and j , we write $[i, j]$ for the set $\{i, i+1, \dots, j\}$, assumed to be empty if $i > j$. For a function $f : A \rightarrow B$, we denote by $f[a_i \leftarrow b_i]_{i \in [1, n]}$ the function that maps a_i into b_i for each $i \in [1, n]$ and agrees with f everywhere else.

2.1 Configurations

We model a parallel system as a hypergraph, whose vertices are *components* (i.e., the nodes of the network) and hyperedges are *interactions* (i.e., describing the way the components communicate with each other). The components are taken from a countably infinite set \mathbb{C} , called the *universe*. We consider that each component executes its own copy of the same *behavior*, represented as a finite-state machine $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$, where \mathcal{P} is a finite set of *ports*, \mathcal{Q} is a finite set of *states* and $\rightarrow \subseteq \mathcal{Q} \times \mathcal{P} \times \mathcal{Q}$ is a transition relation. Intuitively, each transition $q \xrightarrow{p} q'$ of the behavior \mathbb{B} is triggered by a visible event, represented by the port p . The universe \mathbb{C} and the behavior $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$ are considered to be fixed in the following.

A *configuration* is a snapshot of the system, describing the topology of the network (i.e., the set of present components and interactions) together with the local state of each component, formally defined below:

► **Definition 1.** A configuration is a tuple $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$, where:

- $\mathcal{C} \subseteq_{fin} \mathbb{C}$ is a finite set of components, that are present in the configuration,
- $\mathcal{I} \subseteq_{fin} (\mathbb{C} \times \mathcal{P})^+$ is a finite set of interactions, where each interaction is a sequence $(c_i, p_i)_{i \in [1, n]} \in (\mathbb{C} \times \mathcal{P})^n$ that binds together the ports p_1, \dots, p_n of the pairwise distinct components c_1, \dots, c_n , respectively. The ordered sequence of ports (p_1, \dots, p_n) is called an interaction type and we denote by \mathcal{P}^+ the set of interaction types.

² A Hoare triple $\{\phi\} P \{\psi\}$ is *local* if it mentions only those components and interactions added or deleted by P . Local specifications are plugged into a global context by the *frame rule* that infers $\{\phi * F\} P \{\psi * F\}$ from $\{\phi\} P \{\psi\}$ if the variables modified by P are not free in F .

166 \dashv $\varrho : \mathbb{C} \rightarrow \mathcal{Q}$ is a state map associating each (possibly absent) component, a state of the
 167 behavior \mathbb{B} , such that the set $\{c \in \mathbb{C} \mid \varrho(c) = q\}$ is infinite, for each $q \in \mathcal{Q}$.
 168 We denote by Γ the set of configurations.

169 The last condition requires that there is an infinite pool of components in each state $q \in \mathcal{Q}$;
 170 since \mathbb{C} is infinite and \mathcal{Q} is finite, this condition is feasible.

171 **► Example 2.** The configurations of the system from Fig. 1 (a) are $(\{c_1, \dots, c_n\}, \{(c_i, out,$
 172 $c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}, \varrho)$, where $\varrho : \mathbb{C} \rightarrow \{\mathsf{H}, \mathsf{T}\}$ is a state map. The ring topology is
 173 given by components $\{c_1, \dots, c_n\}$ and interactions $\{(c_i, out, c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}$. ■

174 Note that Def. 1 allows configurations with interactions that involve absent components
 175 i.e., not from the set \mathcal{C} of present components in the given configuration. The following
 176 definition distinguishes such configurations:

177 **► Definition 3.** A configuration $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ is said to be tight if and only if for any
 178 interaction $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$ we have $\{c_i \mid i \in [1, n]\} \subseteq \mathcal{C}$ and loose otherwise.

179 For instance, every configuration of the system from Fig. 1 (a) is tight and becomes loose if
 180 a component is deleted.

181 2.2 Configuration Logic

182 Let \mathbb{V} and \mathbb{A} be countably infinite sets of *variables* and *predicates*, respectively. For each
 183 predicate $A \in \mathbb{A}$, we denote its arity by $\#A$. The formulæ of the *Configuration Logic* (CL)
 184 are described inductively by the following syntax:

185 $\phi := \text{emp} \mid [x] \mid \langle x_1.p_1, \dots, x_n.p_n \rangle \mid x@q \mid x = y \mid x \neq y \mid A(x_1, \dots, x_{\#A}) \mid \phi * \phi \mid \exists x. \phi$

186 where $x, y, x_1, \dots \in \mathbb{V}$, $q \in \mathcal{Q}$ and $A \in \mathbb{A}$. A formula $[x]$, $\langle x_1.p_1, \dots, x_n.p_n \rangle$, $x@q$ and
 187 $A(x_1, \dots, x_{\#A})$ is called a *component*, *interaction*, *state* and *predicate* atom, respectively. We
 188 use the shorthand $[x]@q \stackrel{\text{def}}{=} [x] * x@q$. Intuitively, a formula $[x]@q * [y]@q' * \langle x.out, y.in \rangle * \langle x.in, y.out \rangle$
 189 describes a configuration consisting of two distinct components, denoted by the
 190 values of x and y , in states q and q' , respectively, and two interactions binding the *out* port
 191 of one to the *in* port of the other component.

192 A formula with no occurrences of predicate atoms (resp. existential quantifiers) is called
 193 *predicate-free* (resp. *quantifier-free*). A qpf formula is both predicate- and quantifier-free. A
 194 variable is *free* if it does not occur in the scope of a quantifier and $\text{fv}(\phi)$ is the set of free
 195 variables of ϕ . A *substitution* $\phi[x_i/y_i]_{i \in [1, n]}$ replaces simultaneously every free occurrence of
 196 x_i by y_i in ϕ , for all $i \in [1, n]$. The size of a formula ϕ is the total number of occurrences of
 197 symbols needed to write it down, denoted by $\text{size}(\phi)$.

198 The only connective of the logic is the *separating conjunction* $*$. Intuitively, $\phi_1 * \phi_2$
 199 means that ϕ_1 and ϕ_2 hold separately, on disjoint parts of the same configuration. Its formal
 200 meaning is coined by the following definition of *composition* of configurations:

201 **► Definition 4.** The composition of two configurations $\gamma_i = (\mathcal{C}_i, \mathcal{I}_i, \varrho)$, for $i = 1, 2$, such that
 202 $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ and $\mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset$, is defined as $\gamma_1 \bullet \gamma_2 \stackrel{\text{def}}{=} (\mathcal{C}_1 \cup \mathcal{C}_2, \mathcal{I}_1 \cup \mathcal{I}_2, \varrho)$. The composition
 203 $\gamma_1 \bullet \gamma_2$ is undefined if $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$ or $\mathcal{I}_1 \cap \mathcal{I}_2 \neq \emptyset$.

204 **► Example 5.** Let $\gamma_i = (\{c_i\}, \{(c_i, out, c_{3-i}, in)\}, \varrho)$ be configurations, for $i = 1, 2$. Then
 205 $\gamma_1 \bullet \gamma_2 = (\{c_1, c_2\}, \{(c_1, out, c_2, in), (c_2, out, c_1, in)\}, \varrho)$. ■

206 The meaning of the predicates is given by a set of inductive definitions:

207 **► Definition 6.** A set of inductive definitions (SID) Δ consists of rules $A(x_1, \dots, x_{\#A}) \leftarrow \phi$.

Note that having distinct parameters in a rule is without loss of generality, as e.g., a rule $A(x_1, x_1) \leftarrow \phi$ can be equivalently written as $A(x_1, x_2) \leftarrow x_1 = x_2 * \phi$. As a convention, we shall always use the names $x_1, \dots, x_{\#A}$ for the parameters of a rule that defines A . An example of a SID is given in §1.1.

The size of a SID is $\text{size}(\Delta) \stackrel{\text{def}}{=} \sum_{A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta} \text{size}(\phi) + \#A + 1$. Other parameters, relevant for complexity evaluation, are the maximal (1) arity $\#(\Delta) \stackrel{\text{def}}{=} \max\{\#A \mid A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta\}$ of a defined predicate, (2) size of an interaction type $N(\Delta) \stackrel{\text{def}}{=} \max\{n \mid \langle y_1.p_1, \dots, y_n.p_n \rangle \text{ occurs in } \Delta\}$, and (3) number of predicate atoms $H(\Delta) \stackrel{\text{def}}{=} \max\{h \mid A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=1}^h B_\ell(z_\ell), \phi \text{ is a qpf formula}\}$.

The semantics of CL formulæ is defined by a satisfaction relation $\gamma \models_\Delta^\nu \phi$ between configurations and formulæ. This relation is parameterized by a *store* $\nu : \mathbb{V} \rightarrow \mathbb{C}$ mapping the free variables of a formula into components from the universe (possibly absent from γ) and an SID Δ . The definition of the satisfaction relation is by induction on the structure of formulæ, where $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ is a configuration (Def. 1):

$$\begin{array}{lll}
 \gamma \models_\Delta^\nu \text{emp} & \iff & \mathcal{C} = \emptyset \text{ and } \mathcal{I} = \emptyset \\
 \gamma \models_\Delta^\nu [x] & \iff & \mathcal{C} = \{\nu(x)\} \text{ and } \mathcal{I} = \emptyset \\
 \gamma \models_\Delta^\nu \langle x_1.p_1, \dots, x_n.p_n \rangle & \iff & \mathcal{C} = \emptyset \text{ and } \mathcal{I} = \{(\nu(x_1), p_1, \dots, \nu(x_n), p_n)\} \\
 \gamma \models_\Delta^\nu x @ q & \iff & \gamma \models_\Delta^\nu \text{emp} \text{ and } \varrho(\nu(x)) = q \\
 \gamma \models_\Delta^\nu x \sim y & \iff & \gamma \models_\Delta^\nu \text{emp} \text{ and } \nu(x) \sim \nu(y), \text{ for all } \sim \in \{=, \neq\} \\
 \gamma \models_\Delta^\nu A(y_1, \dots, y_{\#A}) & \iff & \gamma \models_\Delta^\nu \phi[x_1/y_1, \dots, x_{\#A}/y_{\#A}], \text{ for some rule } \\
 & & A(x_1, \dots, x_{\#A}) \leftarrow \phi \text{ from } \Delta \\
 \gamma \models_\Delta^\nu \phi_1 * \phi_2 & \iff & \text{there exist } \gamma_1 \text{ and } \gamma_2, \text{ such that } \gamma = \gamma_1 \bullet \gamma_2 \text{ and } \\
 & & \gamma_i \models_\Delta^\nu \phi_i, \text{ for all } i = 1, 2 \\
 \gamma \models_\Delta^\nu \exists x . \phi & \iff & \gamma \models_\Delta^{\nu[x \leftarrow c]} \phi, \text{ for some } c \in \mathbb{C}
 \end{array}$$

If $\gamma \models_\Delta^\nu \phi$, we say that the pair (γ, ν) is a Δ -model of ϕ . If ϕ is a predicate-free formula, the satisfaction relation does not depend on the SID, written $\gamma \models^\nu \phi$. A formula ϕ is *satisfiable* if and only if it has a model. A formula ϕ Δ -*entails* a formula ψ , written $\phi \models_\Delta \psi$, if and only if any Δ -model of ϕ is a Δ -model of ψ . Two formulæ are Δ -*equivalent*, written $\phi \equiv_\Delta \psi$ if and only if $\phi \models_\Delta \psi$ and $\psi \models_\Delta \phi$. A formula ϕ is Δ -*tight* if γ is tight (Def. 3), for any Δ -model (γ, ν) of ϕ . We omit mentioning Δ whenever it is clear from the context or not needed.

2.3 The Havoc Invariance Problem

This paper is concerned with the *havoc invariance* problem i.e., the problem of deciding whether the set of models of a given CL formula is closed under the execution of a sequence of interactions. The execution of an interaction $(c_i, p_i)_{i \in [1, n]}$ synchronizes transitions labeled by the ports p_1, \dots, p_n from the behaviors (i.e., replicas of the state machine \mathbb{B}) of c_1, \dots, c_n , respectively. This joint execution of several transitions in different components of the system is formally described by the *step* relation below:

► **Definition 7.** The step relation $\Rightarrow \subseteq \Gamma \times (\mathbb{C} \times \mathcal{P})^+ \times \Gamma$ is defined as:

$$\begin{array}{l}
 (\mathcal{C}, \mathcal{I}, \varrho) \xRightarrow{(c_i, p_i)_{i \in [1, n]}} (\mathcal{C}', \mathcal{I}', \varrho') \text{ if and only if } (c_i, p_i)_{i \in [1, n]} \in \mathcal{I} \text{ and } \varrho' = \varrho[c_i \leftarrow q'_i]_{i \in [1, n]} \\
 \text{where } \varrho(c_i) = q_i \text{ and } q_i \xrightarrow{p_i} q'_i \text{ is a transition of } \mathbb{B}, \text{ for all } i \in [1, n]
 \end{array}$$

The havoc relation \rightsquigarrow^* is the reflexive and transitive closure of the relation $\rightsquigarrow \subseteq \Gamma^2$: $(\mathcal{C}, \mathcal{I}, \varrho) \rightsquigarrow (\mathcal{C}', \mathcal{I}', \varrho')$ if and only if $(\mathcal{C}, \mathcal{I}, \varrho) \xRightarrow{(c_i, p_i)_{i \in [1, n]}} (\mathcal{C}', \mathcal{I}', \varrho')$, for some interaction $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$.

240 ► **Example 8.** Let $\gamma_i = (\{c_1, c_2, c_3\}, \{(c_i, out, c_{i \bmod 3+1}, in) \mid i \in [1, 3]\}, \varrho_i)$, for $i \in [1, 3]$ be
 241 configurations, where $\varrho_1(c_1) = \varrho_1(c_2) = H$, $\varrho_1(c_3) = T$, $\varrho_2(c_1) = T$, $\varrho_2(c_2) = H$, $\varrho_2(c_3) = H$,
 242 $\varrho_3(c_1) = \varrho_3(c_3) = H$, $\varrho_3(c_2) = T$. Then we have $\gamma_i \rightsquigarrow^* \gamma_j$, for all $i, j \in [1, 3]$. ■

243 Two interactions $(c_1, p_1, \dots, c_n, p_n)$ and $(c_{i_1}, p_{i_1}, \dots, c_{i_n}, p_{i_n})$ such that $\{i_1, \dots, i_n\} =$
 244 $[1, n]$, are equivalent from the point of view of the step relation, since the set of executed
 245 transitions is the same; nevertheless, we chose to distinguish them in the following, for reasons
 246 of simplicity. Note, moreover, that the havoc relation does not change the component or the
 247 interaction set of a configuration, only its state map.

248 ► **Definition 9.** Given an SID Δ and a predicate A , the problem $\text{HavocInv}[\Delta, A]$ asks whether
 249 for all $\gamma, \gamma' \in \Gamma$ and each store ν , such that $\gamma \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ and $\gamma \rightsquigarrow^* \gamma'$, it is the case
 250 that $\gamma' \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$?

251 ► **Example 10.** Consider a model $\gamma = (\{c_1, \dots, c_n\}, \{(c_i, out, c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}, \varrho)$
 252 of the formula $\text{ring}_{1,1}()$ i.e., having the property that $\varrho(c_i) = H$ and $\varrho(c_j) = T$ for at least
 253 two indices $i \neq j \in [1, n]$, where the SID that defines $\text{ring}_{1,1}()$ is given in §1.1. Similar to
 254 Example 8, in any configuration $\gamma' = (\{c_1, \dots, c_n\}, \{(c_i, out, c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}, \varrho')$
 255 such that $\gamma \rightsquigarrow^* \gamma'$, we have $\varrho'(c_k) = H$ and $\varrho'(c_\ell) = T$, for some $k \neq \ell \in [1, n]$, hence γ' is a
 256 model of $\text{ring}_{1,1}()$, meaning that $\text{ring}_{1,1}()$ is havoc invariant. Examples of formulæ that are
 257 not havoc invariant include e.g., $[x]@T * \langle x.out, y.in \rangle * [y]@H$. ■

258 Without loss of generality, we consider the havoc invariance problem only for single
 259 predicate atoms. This is because, for any formula ϕ , such that $\text{fv}(\phi) = \{x_1, \dots, x_n\}$, one may
 260 consider a fresh predicate symbol (i.e., not in the SID) A_ϕ and add the rule $A_\phi(x_1, \dots, x_n) \leftarrow \phi$
 261 to the SID. Then ϕ is havoc invariant if and only if $A_\phi(x_1, \dots, x_n)$ is havoc invariant.

262 3 From Havoc Invariance to Entailment

263 We describe a many-one reduction of the havoc invariance (Def. 9) to the entailment problem,
 264 following three steps. Given an instance $\text{HavocInv}[\Delta, A]$ of the havoc invariance problem, the
 265 SID Δ is first translated into a tree automaton recognizing trees labeled with predicate-free
 266 formulæ, that symbolically encode the set of Δ -models of the predicate atom $A(x_1, \dots, x_{\#A})$.
 267 Second, we define a structure-preserving tree transducer that simulates the effect of executing
 268 exactly one interaction from such a model. Third, we compute the image of the language
 269 recognized by the first tree automaton via the transducer, as a second tree automaton, which
 270 is translated back into another SID $\bar{\Delta}$ defining one or more predicates $\bar{A}_1, \dots, \bar{A}_p$, among
 271 other. Finally, we prove that $\text{HavocInv}[\Delta, A]$ has a positive answer if and only if each of the
 272 entailments $\{\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\Delta \cup \bar{\Delta}} A(x_1, \dots, x_{\#A})\}_{i=1}^p$ produced by the reduction, hold.

273 For the sake of self-containment, we recall below the definitions of trees, tree automata
 274 and (structure-preserving) tree transducers. Let $(\Sigma, \#)$ be a ranked alphabet, where each
 275 symbol $\alpha \in \Sigma$ has an associated arity $\#\alpha \geq 0$. A tree over Σ is a finite partial function
 276 $t : \mathbb{N}^* \rightarrow_{fin} \Sigma$, whose domain $\text{dom}(t) \subseteq_{fin} \mathbb{N}^*$ is both *prefix-closed* i.e., $u \in \text{dom}(t)$, for all
 277 $u, v \in \mathbb{N}^*$, such that $u \cdot v \in \text{dom}(t)$, and *complete* i.e., $\{n \in \mathbb{N} \mid u \cdot n \in \text{dom}(t)\} = [1, \#t(u)]$,
 278 for all $u \in \text{dom}(t)$. Given $u \in \text{dom}(t)$, the *subtree* of t rooted at u is the tree $t|_u$, such that
 279 $\text{dom}(t|_u) \stackrel{\text{def}}{=} \{w \mid u \cdot w \in \text{dom}(t)\}$ and $t|_u(w) \stackrel{\text{def}}{=} t(u \cdot w)$. We denote by $\mathbb{T}(\Sigma)$ the set of trees
 280 over a ranked alphabet Σ .

281 A tree automaton (TA) is a tuple $\mathcal{A} = (\Sigma, \mathcal{S}, \mathcal{F}, \delta)$, where Σ is a ranked alphabet,
 282 \mathcal{S} is a finite set of states, $\mathcal{F} \subseteq \mathcal{S}$ is a set of final states and δ is a set of transitions
 283 $\alpha(s_1, \dots, s_{\#a}) \rightarrow s$; when $\#\alpha = 0$, we write $\alpha \rightarrow s$ instead of $\alpha() \rightarrow s$. A run of \mathcal{A} over

a tree t is a function $\pi : \text{dom}(t) \rightarrow \mathcal{S}$, such that, for all $u \in \text{dom}(t)$, we have $\pi(u) = s$ if $(t(u))(\pi(u \cdot 1), \dots, \pi(u \cdot \#t(u))) \rightarrow s \in \delta$. Given a state $q \in \mathcal{S}$, a run π of is q -accepting if and only if $\pi(\epsilon) = q$, in which case \mathcal{A} is said to q -accept t . We denote by $\mathcal{L}_q(\mathcal{A})$ the set of trees q -accepted by \mathcal{A} and let $\mathcal{L}(\mathcal{A}) \stackrel{\text{def}}{=} \bigcup_{q \in \mathcal{F}} \mathcal{L}_q(\mathcal{A})$. A language L is *recognizable* if and only if there exists a TA \mathcal{A} , such that $L = \mathcal{L}(\mathcal{A})$.

A *tree transducer* (TT) is a tree automaton over an alphabet of pairs $\mathcal{T} = (\Sigma^2, \mathcal{S}, \mathcal{F}, \delta)$, such that $\#\alpha = \#\beta = n$, for each transition $(\alpha, \beta)(s_1, \dots, s_n) \rightarrow s \in \delta$. Intuitively, a transition of the transducer reads a symbol α from the input tree and writes another symbol β to the output tree, at the same position. Clearly, any tree $t : \mathbb{N}^* \rightarrow_{\text{fin}} \Sigma^2$ with labels from the set of pairs $\{(\alpha, \beta) \in \Sigma^2 \mid \#\alpha = \#\beta\}$ can be viewed as a pair of trees (t_1, t_2) over Σ , such that $\text{dom}(t_1) = \text{dom}(t_2) = \text{dom}(t)$. In order to define the image of a tree language via a transducer, we define (i) *projection* $L \downarrow_i \stackrel{\text{def}}{=} \{t_i \mid (t_1, t_2) \in L\}$, for all $i = 1, 2$, where $L \subseteq \mathbb{T}(\Sigma^2)$, and (ii) *cylindrification* $L \uparrow_i \stackrel{\text{def}}{=} \{(t_1, t_2) \mid t_i \in L\}$, for all $i = 1, 2$, where $L \subseteq \mathbb{T}(\Sigma)$. The *image* of a language $L \subseteq \mathbb{T}(\Sigma)$ via a transducer \mathcal{T} is the language $\mathcal{T}(L) \stackrel{\text{def}}{=} (L \uparrow^1 \cap \mathcal{L}(\mathcal{T})) \downarrow_2$. It is manifest that $\mathcal{T}(L)$ is recognizable whenever L is recognizable.

3.1 From SID to Tree Automata and Back

We define a two-way connection between SIDs and TAs, as follows:

1. Given a finite SID Δ we define a TA \mathcal{A}_Δ , whose states q_A are named after the predicates A that occur in Δ and whose alphabet consists of the predicate-free formulæ from the rules of Δ , with variables mapped to canonical names, together with a tuple of arities, needed for later bookkeeping. Each tree $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$ defines a unique predicate-free formula $\Phi(t)$, such that the Δ -models of a predicate atom $A(x_1, \dots, x_{\#A})$ are exactly the models of some $\Phi(t)$, for $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$.
2. Conversely, given a TA \mathcal{A} over an alphabet of formulæ annotated with arities, the tuple of arities associated with each alphabet symbol allows to define a SID $\Delta_{\mathcal{A}}$, whose predicates A_q are named after the states q of the TA, such that the models of the formulæ $\Phi(t)$, such that $t \in \mathcal{L}_q(\mathcal{A})$ are exactly the $\Delta_{\mathcal{A}}$ -models of the predicate atom $A_q(x_1, \dots, x_{\#A_q})$.

Let us fix a countably infinite set of variables $\widetilde{\text{Var}} \stackrel{\text{def}}{=} \{\tilde{x}_i \mid i \geq 1\} \cup \{\tilde{z}_i^{(\ell)} \mid i, \ell \geq 1\}$, with the understanding that \tilde{x}_i are canonical names for the variables from the left-hand side and $\tilde{z}_i^{(\ell)}$ are canonical names for the variables occurring in the ℓ -th predicate atom from the right-hand side of a rule. An *alphabet symbol* $\alpha = \langle \psi, a_0, \dots, a_h \rangle$ consists of a predicate-free formula ψ and a tuple of positive integers $a_0, \dots, a_h \in \mathbb{N}$, such that $\text{fv}(\psi) = \{\tilde{x}_i \mid i \in [1, a_0]\} \cup \{\tilde{z}_i^{(\ell)} \mid \ell \in [1, h], i \in [1, a_\ell]\}$. We take the arity of such a symbol to be $\#\alpha \stackrel{\text{def}}{=} h$ and denote by $\widetilde{\Sigma}$ the (infinite) set of alphabet symbols. Trees labeled with symbols from $\widetilde{\Sigma}$ define predicate-free *characteristic formulæ*, as follows:

► **Definition 11.** Given a tree $t \in \mathbb{T}(\widetilde{\Sigma})$, where $t(\epsilon) = \langle \exists y_1 \dots \exists y_m \cdot \phi, a_0, \dots, a_h \rangle$ with ϕ a qpf formula, and a node $u \in \mathbb{N}^*$, we define the qpf characteristic formula:

$$\Psi^u(t) \stackrel{\text{def}}{=} \phi[\tilde{x}_j/x_j^u]_{j \in [1, a_0]} [\tilde{z}_j^{(\ell)}/x_j^{u \cdot \ell}]_{\ell \in [1, h], j \in [1, a_\ell]} [y_j/y_j^u]_{j \in [1, m]} * \bigstar_{\ell \in [1, h]} \Psi^{u \cdot \ell}(t|_\ell)$$

Assuming that $t(v) = \langle \exists y_1 \dots \exists y_{m^v} \cdot \phi^v, a_0^v, \dots, a_h^v \rangle$, for all $v \in \text{dom}(t)$, we consider also the predicate-free formula $\Phi^u(t) = (\exists x_j^{u \cdot v})_{v \in \text{dom}(t) \setminus \{\epsilon\}, j \in [1, a_0^v]} (\exists y_j^{u \cdot v})_{v \in \text{dom}(t), j \in [1, m^v]} \cdot \Psi^u(t)$.

► **Example 12.** We consider a system whose components form a tree, in which each parent sends a request (*req*) to and receives replies (*reply*) from both its children. In addition, the leaves of the tree form a ring, with the *out* port of each leaf connected to the *in* port of its



$$325 \quad \text{Root}() \leftarrow \exists n \exists \ell \exists r . \langle r.out, \ell.in \rangle * \text{Node}(n, \ell, r) \quad (1)$$

$$326 \quad \text{Node}(n, \ell, r) \leftarrow \exists n_1 \exists r_1 \exists n_2 \exists \ell_2 . [n] * \langle n.\text{req}, n_1.\text{reply}, n_2.\text{reply} \rangle * \langle r_1.\text{in}, \ell_2.\text{out} \rangle *$$

$$Node(n_1, \ell, r_1) * Node(n_2, \ell_2, r) \quad (2)$$

$$\text{Node}(n, \ell, r) \leftarrow [n]@q_0 \quad \text{Node}(n, \ell, r) \leftarrow [n]@q_1 \quad (3)$$

$$\alpha \stackrel{\text{def}}{=} \langle \exists n \exists \ell \exists r . \langle r.out, \ell.in \rangle * \tilde{z}_1^{(1)} = n * \tilde{z}_2^{(1)} = \ell * \tilde{z}_3^{(1)} = r, 0, 3 \rangle$$

$$\beta \stackrel{\text{def}}{=} \langle \exists n_1 \exists r_1 \exists n_2 \exists l_2 . [n] * \langle n.\text{req}, l.\text{reply}, r.\text{reply} \rangle * \langle r_1.\text{in}, l_2.\text{out} \rangle *$$

$$333 \quad \tilde{z}_1^{(1)} = n_1 * \tilde{z}_2^{(1)} = \ell * \tilde{z}_3^{(1)} = r_1 * \tilde{z}_1^{(2)} = n_2 * \tilde{z}_2^{(2)} = \ell * \tilde{z}_3^{(2)} = r, 3, 3, 3)$$

$$334 \quad \gamma_0 \stackrel{\text{def}}{=} \langle [\tilde{x}_1] @ q_0, 3 \rangle \qquad \gamma_1 \stackrel{\text{def}}{=} \langle [\tilde{x}_1] @ q_1, 3 \rangle$$

336 For simplicity, Fig. 2 shows only the formulæ, not the arity lists of the alphabet symbols. ■

The models of the characteristic formula $\Psi^\epsilon(\mathbf{t})$ of a tree $\mathbf{t} \in \mathbb{T}(\tilde{\Sigma})$ define walks in the tree that correspond to chains of equalities between variables. Formally, a *walk* in \mathbf{t} is a sequence of nodes $u_1, \dots, u_n \in \text{dom}(\mathbf{t})$, such that u_i is either the parent or a child of u_{i+1} , for all $i \in [1, n-1]$. Note that a walk can visit the same node of the tree several times. In particular, if the characteristic formula $\Psi^\epsilon(\mathbf{t})$ is tight (i.e., has only tight models in the sense of Def. 3) there exist equality walks between the node containing an interaction atom and the nodes where these variables are instantiated by component atoms. For instance, walks between the root containing $\langle r.out, \ell.in \rangle$ and the left- and right-most leafs, labeled with component atoms that associate elements of \mathbb{C} to the variables ℓ and r are shown in Fig. 2.

► **Lemma 13.** *Let $\mathfrak{t} \in \mathbb{T}(\widetilde{\Sigma})$ be a tree, such that $\Psi^\epsilon(\mathfrak{t})$ is tight, (γ, ν) be a model of $\Psi^\epsilon(\mathfrak{t})$ and y^v, z^w be two variables that occur in a component and interaction atom of $\Psi^\epsilon(\mathfrak{t})$, respectively. Then $\nu(y^v) = \nu(z^w)$ if and only if there exists a walk u_1, \dots, u_n in \mathfrak{t} and variables $y^v = x_{i_1}^{u_1}, \dots, x_{i_n}^{u_n} = z^w$, such that either $x_{i_j}^{u_j}$ and $x_{i_{j+1}}^{u_{j+1}}$ are the same variable, or the equality atom $x_{i_j}^{u_j} = x_{i_{j+1}}^{u_{j+1}}$ occurs in $\Psi^\epsilon(\mathfrak{t})$, for all $j \in [1, n-1]$.*

Let Δ be a fixed and finite SID in the following. We build a TA \mathcal{A}_Δ that recognizes the Δ -models of each predicate atom defined by Δ , in the sense of Lemma 16 below.

► **Definition 14.** We associate each rule $r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * \bigstar_{\ell \in [1, h]} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$, where ϕ is a qpf formula, with the alphabet symbol:

$$\alpha_r \stackrel{\text{def}}{=} \left\langle \exists y_1 \dots \exists y_m . \left(\phi * \bigstar_{\ell \in [1, h], i \in [1, \#B_\ell]} \tilde{z}_i^{(\ell)} = z_i^\ell \right) [x_j / \tilde{x}_j]_{j \in [1, \#A]}, \#A, \#B_1, \dots, \#B_h \right\rangle$$

Let $\mathcal{A}_\Delta \stackrel{\text{def}}{=} (\Sigma_\Delta, \mathcal{S}_\Delta, \delta_\Delta)$ be a TA, where $\Sigma_\Delta \stackrel{\text{def}}{=} \{\alpha_r \mid r \in \Delta\}$, $\mathcal{S}_\Delta \stackrel{\text{def}}{=} \{q_A \mid A \in \text{Def}(\Delta)\}$ and $\delta_\Delta \stackrel{\text{def}}{=} \{\alpha_r(q_{B_1}, \dots, q_{B_h}) \rightarrow q_A \mid r \in \Delta\}$.

► **Example 15** (contd. from Example 12). The TA corresponding to the SID in Example 12 is $\mathcal{A}_\Delta = (\tilde{\Sigma}, \mathcal{S}_\Delta, \delta_\Delta)$, where $\tilde{\Sigma} = \{\alpha, \beta, \gamma_0, \gamma_1\}$, $\mathcal{S}_\Delta = \{q_{\text{Root}}, q_{\text{Node}}\}$ and $\delta_\Delta = \{\alpha(q_{\text{Node}}) \rightarrow q_{\text{Root}}, \beta(q_{\text{Node}}, q_{\text{Node}}) \rightarrow q_{\text{Node}}, \gamma_0 \rightarrow q_{\text{Node}}, \gamma_1 \rightarrow q_{\text{Node}}\}$. ■

The following lemma proves that the predicate-free formulæ corresponding (in the sense of Def. 11) to the trees recognized by \mathcal{A}_Δ in a state q_A define the Δ -models of the predicate atom $A(x_1, \dots, x_{\#A})$:

► **Lemma 16.** *For any predicate $A \in \text{Def}(\Delta)$, configuration γ , store ν and node $u \in \mathbb{N}^*$, we have $\gamma \models_\Delta^\nu A(x_1^u, \dots, x_{\#A}^u)$ if and only if $\gamma \models^\nu \Phi^u(t)$, for some tree $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$.*

Conversely, given a tree automaton $\mathcal{A} = (\Sigma, \mathcal{S}, \delta)$, we construct a SID $\Delta_\mathcal{A}$ that defines the models of the predicate-free formulæ corresponding (Def. 11) to the trees recognized by \mathcal{A} (Lemma 19). We assume that the alphabet Σ consists of symbols $\langle \psi, a_0, \dots, a_h \rangle$ of arity h , where ψ is a predicate-free formula with free variables $\text{fv}(\psi) = \{\tilde{x}_i \mid i \in [1, a_0]\} \cup \{\tilde{z}_i^{(\ell)} \mid \ell \in [1, h], i \in [1, a_\ell]\}$ and that the transitions of the TA meet the requirement:

► **Definition 17.** *A TA \mathcal{A} is SID-compatible iff for any transitions $\langle \psi, a_0, \dots, a_h \rangle(q_1, \dots, q_h) \rightarrow q_0$ and $\langle \psi', a'_0, \dots, a'_h \rangle(q'_1, \dots, q'_h) \rightarrow q'_0$ of \mathcal{A} , we have $q_i = q'_i$ only if $a_i = a'_i$, for all $i \in [0, h]$.*

Let us fix a SID-compatible TA $\mathcal{A} = (\Sigma, \mathcal{S}, \delta)$ for the rest of this section.

► **Definition 18.** *The SID $\Delta_\mathcal{A}$ has a rule:*

$$A_{q_0}(x_1, \dots, x_{a_0}) \leftarrow \exists y_1^1 \dots \exists y_{a_h}^h . \phi[\tilde{x}_i/x_i]_{i \in [1, a_0]} [\tilde{z}_i^{(\ell)}/y_i^\ell]_{\ell \in [1, h], i \in [1, a_\ell]} * \bigstar_{\ell \in [1, h]} A_{q_\ell}(y_1^\ell, \dots, y_{a_\ell}^\ell)$$

for each transition $\langle \phi, a_0, \dots, a_h \rangle(q_1, \dots, q_h) \rightarrow q_0$ of \mathcal{A} and those rules only.

The following lemma states that $\Delta_\mathcal{A}$ defines the set of models of the characteristic formulæ (Def. 11) of the trees recognized by \mathcal{A} .

► **Lemma 19.** *For any state $q \in \mathcal{S}$, configuration γ , store ν and node $u \in \mathbb{N}^*$, we have $\gamma \models_{\Delta_\mathcal{A}}^\nu A_q(x_1^u, \dots, x_{\#A_q}^u)$ if and only if $\gamma \models^\nu \Phi^u(t)$, for some tree $t \in \mathcal{L}_q(\mathcal{A})$.*

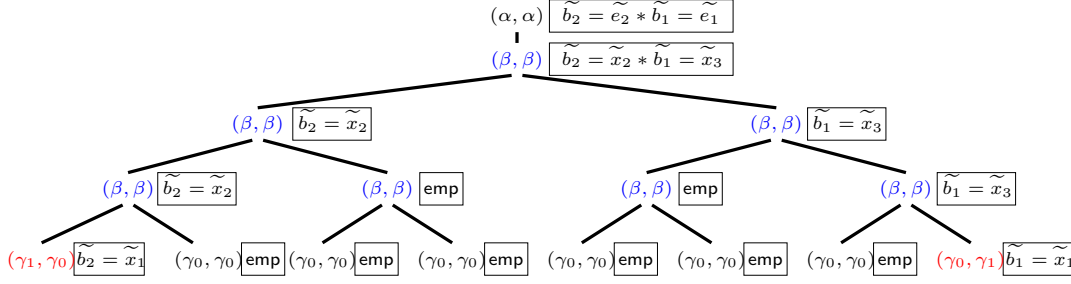
3.2 Encoding Havoc Steps by Tree Transducers

The purpose of this section is the definition of a transducer that simulates one havoc step. Before giving its definition, we note that the havoc invariance problem can be equivalently defined by considering the transformation induced by a single havoc step, instead of an arbitrary sequence of steps. The following lemma can be taken as an equivalent definition:

► **Lemma 20.** *HavocInv $[\Delta, A]$ has a positive answer if and only if, for all $\gamma, \gamma' \in \Gamma$ and each store ν , such that $\gamma \models_\Delta^\nu A(x_1, \dots, x_{\#A})$ and $\gamma \rightsquigarrow \gamma'$, it is the case that $\gamma' \models_\Delta^\nu A(x_1, \dots, x_{\#A})$.*

We fix a SID Δ for the rest of this section and recall the existence of a fixed finite-state behavior $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$ with ports \mathcal{P} , states \mathcal{Q} and transitions $q \xrightarrow{p} q' \in \mathcal{Q} \times \mathcal{P} \times \mathcal{Q}$. We define a transducer \mathcal{T}_τ parameterized by a given interaction type $\tau = (p_1, \dots, p_n) \in \mathcal{P}^+$. The havoc step transducer is the automata-theoretic union of the typed transducers over the set of interaction types that occurs in Δ .

Given a tree $t \in \mathbb{T}(\tilde{\Sigma})$, an interaction-typed transducer \mathcal{T}_τ (1) guesses an interaction atom $\langle z_1.p_1, \dots, z_n.p_n \rangle$ that occurs in some label of t , (2) tracks the equality walks (Lemma 13) between each variable z_i and the component atom $[x_i]@q_i$ that defines the store value of z_i



■ **Figure 3** Tree Transducer for the Interactions of Type (out, in) in the System from Fig. 2

and its current state, and (3) replaces each state component atom $[x_i]@q_i$ by $[x_i]@q'_i$, where $q_i \xrightarrow{p_i} q'_i$ is a transition from \mathbb{B} , for each $i \in [1, n]$. The output of the transducer is a tree $t' \in \mathbb{T}(\Sigma)$, that symbolically encodes the effect of executing some interaction of type τ over t . The main challenge in defining \mathcal{T}_τ is that the equality walks between an interaction atom $\langle z_1.p_1, \dots, z_n.p_n \rangle$ and the component atoms instatiating the variables z_1, \dots, z_n may visit a tree node more than once. To capture this, the transducer will guess at once the equalities summarizing the different fragments of the walk that lie in the currently processed subtree of t . Accordingly, the states of \mathcal{T}_τ are conjunctions of equalities, with special variables \tilde{b}_i (resp. \tilde{e}_i) indicating whether a component (resp. interaction) atom has already been encountered in the current subtree, intuitively marking the *beginning* (resp. *end*) of the walk.

For an interaction type $\tau = (p_1, \dots, p_n)$, let $\widetilde{\text{TVar}}_\tau \stackrel{\text{def}}{=} \{\tilde{x}_i \mid i \in [1, \#(\Delta)]\} \cup \{\tilde{b}_i, \tilde{e}_i \mid i \in [1, n]\}$ and let $Eq(\widetilde{\text{TVar}}_\tau)$ be the set of separating conjunctions of equality atoms i.e., $\varphi \stackrel{\text{def}}{=} *_{i \in I} x_i = y_i$, such that $\text{fv}(\varphi) \subseteq \widetilde{\text{TVar}}_\tau$. Note that $\exists x. \varphi$, for $\varphi \in Eq(\widetilde{\text{TVar}}_\tau)$, is equivalent to a formula from $Eq(\widetilde{\text{TVar}}_\tau)$ obtained by eliminating the quantifier: either x occurs in an atom $x = y$ for a variable y distinct from x then $(\exists x. \varphi) \equiv \varphi[x/y]$, or $x \notin \text{fv}(\varphi)$, in which case $(\exists x. \varphi) \equiv \varphi$.

► **Definition 21.** The transducer $\mathcal{T}_\tau \stackrel{\text{def}}{=} (\Sigma_\Delta^2, \mathcal{S}_\tau, \mathcal{F}_\tau, \delta_\tau)$, where $\tau = (p_1, \dots, p_n)$, is as follows:

- $\mathcal{S}_\tau = \{\varphi \in Eq(\widetilde{\text{TVar}}_\tau) \mid \varphi \not\models (\tilde{b}_i = \tilde{b}_j), \varphi \not\models (\tilde{e}_i = \tilde{e}_j), \varphi \not\models (\tilde{b}_i = \tilde{e}_j), \text{ for any } i \neq j\}$,
- $\mathcal{F}_\tau = \{\varphi \in \mathcal{S}_\tau \mid \varphi \models *_{i \in [1, n]} (\tilde{b}_i = \tilde{e}_i)\}$, and
- δ_τ contains transitions of the form $(\alpha, \alpha')(\varphi_1, \dots, \varphi_h) \rightarrow_\tau \varphi$ where:
 - $\alpha = (\exists y_1 \dots \exists y_m. \psi, a_0, \dots, a_h)$ and $\alpha' = (\exists y_1 \dots \exists y_m. \psi', a_0, \dots, a_h)$, where ψ and ψ' are qpf formulæ such that $\text{fv}(\psi) = \text{fv}(\psi') \subseteq \widetilde{\text{TVar}}_\tau \cup \{y_1, \dots, y_m\}$,
 - there exists a set $I = \{i_1, \dots, i_r\} \subseteq [1, n]$, variables $\xi_1, \dots, \xi_r \in \text{fv}(\psi)$ and transitions $q_1 \xrightarrow{p_{i_1}} q'_1, \dots, q_r \xrightarrow{p_{i_r}} q'_r$ in \mathbb{B} , such that $\psi = (*_{k \in [1, r]} [\xi_k]@q_k) * \eta$ and $\psi' = (*_{k \in [1, r]} [\xi_k]@q'_k) * \eta$, for some qpf formula η ,
 - there exists a set $J \in \{\emptyset, [1, n]\}$, such that ψ contains an interaction atom $\langle \zeta_1.p_1, \dots, \zeta_n.p_n \rangle$ if $J = [1, n]$,
 - the sets I and $\{i \in [1, n] \mid \tilde{b}_i \in \text{fv}(\varphi_\ell)\}_{\ell \in [1, h]}$ are pairwise disjoint,
 - at most one of the sets $J, \{i \in [1, n] \mid \tilde{e}_i \in \text{fv}(\varphi_\ell)\}_{\ell \in [1, h]}$ is not empty,
 - φ is the result of eliminating the quantifiers from the separating conjunction of equalities:

$$\exists \tilde{z}_1^{(1)} \dots \exists \tilde{z}_{a_h}^{(h)} \exists y_1 \dots \exists y_m. *_{\ell \in [1, h]} \varphi_\ell[\tilde{x}_j/\tilde{z}_j^{(\ell)}]_{j \in [1, a_\ell]} * *_{k \in [1, r]} \tilde{b}_{i_k} = \xi_k * *_{\ell \in J} \tilde{e}_\ell = \zeta_\ell * \psi_{eq}$$

where ψ_{eq} is the separating conjunction of the equality atoms from ψ .

► **Example 22.** (contd. from Examples 12 and 15) Fig. 3 shows a run of the transducer $\mathcal{T}_{(out, in)}$, that describes the symbolic execution of the interaction corresponding to the $\langle r.out, \ell.in \rangle$ interaction atom from the root of the tree in Fig. 2. The states of the transducer are separating conjunctions of equality atoms, enclosed within square boxes. The transducer replaces the component atoms $\gamma_1 = \langle [\tilde{x}_1]@q_1, 3 \rangle$ with $\gamma_0 = \langle [\tilde{x}_1]@q_0, 3 \rangle$ (resp. γ_0 with γ_1) in the left-most (resp. right-most) leaf of the tree. ■

Let $L \subseteq \mathbb{T}(\tilde{\Sigma})$ be an arbitrary language. The following lemmas prove that the transducer \mathcal{T}_τ from Def. 21 correctly simulates a havoc step produced by an interaction of type τ .

► **Lemma 23.** *For each tree $\mathbf{t} \in L$, such that $\Phi^\epsilon(\mathbf{t})$ is tight, configurations $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$, $\gamma' \in \Gamma$ and store ν , such that $\gamma \models^\nu \Phi^\epsilon(\mathbf{t})$ and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some $c_1, \dots, c_n \in \mathcal{C}$ and $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$, there exists a tree $\mathbf{t}' \in \mathcal{T}_{(p_1, \dots, p_n)}(L)$, such that $\gamma' \models^\nu \Phi^\epsilon(\mathbf{t}')$.*

Note that the condition of $\Phi^\epsilon(\mathbf{t})$ having only tight models is necessary to avoid interactions $(c_i, p_i)_{i \in [1, n]}$ that fire by “accident” i.e., when the interaction is created by an atom $\langle \zeta_1.p_1, \dots, \zeta_n.p_n \rangle$, with the components c_1, \dots, c_n created by component atoms $[\xi_1]@q_1, \dots, [\xi_n]@q_n$, such that the equality $\xi_i = \zeta_i$ is not the consequence of $\Phi^\epsilon(\mathbf{t})$, for some $i \in [1, n]$. The effect of such interactions is not captured by the transducer introduced by Def. 21. Tightness is, moreover, a necessary condition of Lemma 13, that ensures the existence of equality walks between the variables occurring in an interaction atom and those of the atoms creating the components to which these variables are mapped, in a model of $\Phi^\epsilon(\mathbf{t})$.

► **Lemma 24.** *For each tree $\mathbf{t}' \in \mathcal{T}_{(p_1, \dots, p_n)}(L)$, configuration $\gamma' \in \Gamma$ and store ν , such that $\gamma' \models^\nu \Phi^\epsilon(\mathbf{t}')$, there exists a configuration $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ and a tree $\mathbf{t} \in L$, such that $\gamma \models^\nu \Phi^\epsilon(\mathbf{t})$ and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some $c_1, \dots, c_n \in \mathcal{C}$ and $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$.*

3.3 The Main Result

We establish the main result of this section, which is a many-one reduction of the havoc invariance to the entailment problem. The result is sharpened by proving that the reduction (i) preserves the class of the SID (see Def. 25 below), and (ii) is polynomial when several parameters of the SID are bounded by constants and simply exponential otherwise. In particular, a class-preserving polynomial reduction ensures that the decidability and complexity upper bounds of the entailment problem carry over to the havoc invariance problem.

► **Definition 25.** *For two predicate-free formulæ ϕ and ψ , we write $\phi \simeq \psi$ if and only if they become equivalent when dropping the state atoms from both. For an arity-preserving equivalence relation $\sim \subseteq \mathbb{A} \times \mathbb{A}$ (i.e., $\#A = \#B$, for all $A \sim B$), for any two rules r_1 and r_2 , we write $r_1 \approx r_2$ if and only if $r_1 = A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * \bigstar_{\ell \in [1, h]} B_\ell(\mathbf{z}_\ell)$, $r_2 = A'(x_1, \dots, x_{\#A'}) \leftarrow \exists y'_1 \dots \exists y'_p \cdot \psi * \bigstar_{\ell \in [1, h]} B'_\ell(\mathbf{u}_\ell)$, $\exists y_1 \dots \exists y_m \cdot \phi \simeq \exists y'_1 \dots \exists y'_p \cdot \psi$, $A \sim A'$ and $B_\ell \sim B'_\ell$, for all $\ell \in [1, h]$. For two SIDs Δ_1 and Δ_2 , we write $\Delta_1 \preceq \Delta_2$ if and only if for each rule $r_1 \in \Delta_1$ there exists a rule $r_2 \in \Delta_2$, such that $r_1 \approx r_2$. We denote by $\Delta_1 \approx \Delta_2$ the conjunction of $\Delta_1 \preceq \Delta_2$ and $\Delta_2 \preceq \Delta_1$.*

If $A_1 \sim A_2$ and $\Delta_1 \approx \Delta_2$ then Δ_1 -models of $A_1(x_1, \dots, x_{\#A_1})$ differ from the Δ_2 -models of $A_2(x_1, \dots, x_{\#A_1})$ only by a renaming of the states occurring within state atoms. This is because any derivation of the satisfaction relation $\gamma \models_{\Delta_1}^\nu A_1(x_1, \dots, x_{\#A_1})$ can be mimicked (modulo the state atoms that may change) by a derivation of $\gamma \models_{\Delta_2}^\nu A_2(x_1, \dots, x_{\#A_2})$, and viceversa. We are now in the position of stating the main result of this section:

► **Theorem 26.** *Assuming that $A(x_1, \dots, x_{\#A})$ is a Δ -tight formula, each instance $\text{HavocInv}[\Delta, A]$ of the havoc invariance problem can be reduced to a set $\{\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\Delta \cup \bar{\Delta}} A(x_1, \dots, x_{\#A})\}_{i=1}^p$ of entailments, where $\Delta \approx \bar{\Delta}$, for an arity-preserving equivalence relation $\sim \subseteq \mathbb{A} \times \mathbb{A}$, such that $\bar{A}_i \sim A$, for all $i \in [1, p]$. The reduction is polynomial, if $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$ are bounded by constants and simply exponential, otherwise.*

4 Decidability and Complexity

We prove the undecidability of the havoc invariance problem (Def. 9) using a reduction from the universality of context-free languages, a textbook undecidable problem [2].

► **Theorem 27.** *The $\text{HavocInv}[\Delta, A]$ problem is undecidable.*

The undecidability proof for the havoc invariance problem uses an argument similar to the one used to prove undecidability of the entailment problem [4, Theorem 4]. We leverage further from this similarity and carve a fragment of CL with a decidable havoc invariance problem, based on the reduction from Theorem 26. For self-containment reasons, we recall the definition of a CL fragment for which the entailment problem is decidable (see [4, §6] for more details and proofs). This definition relies on three, easily checkable, syntactic restrictions on the rules of the SID and a decidable semantic restriction on the models of a predicate atom defined by the SID. The syntactic restrictions use the notion of profile:

► **Definition 28.** *The profile of a SID Δ is the pointwise greatest function $\lambda_\Delta : \mathbb{A} \rightarrow \text{pow}(\mathbb{N})$, mapping each predicate A into a subset of $[1, \#A]$, such that, for each rule $A(x_1, \dots, x_{\#A}) \leftarrow \phi$ from Δ , each atom $B(y_1, \dots, y_{\#B})$ from ϕ and each $i \in \lambda_\Delta(B)$, there exists $j \in \lambda_\Delta(A)$, such that x_j and y_i are the same variable.*

The profile identifies the parameters of a predicate that are always replaced by a variable $x_1, \dots, x_{\#A}$ in each unfolding of $A(x_1, \dots, x_{\#A})$, according to the rules in Δ ; it is computed by a greatest fixpoint iteration, in polynomial time.

► **Definition 29.** *A rule $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, where ϕ is a qpf formula, is said to be:*

1. *progressing (P) if and only if $\phi = [x_1] * \psi$, where ψ consists of interaction atoms involving x_1 and (dis-)equalities, such that $\bigcup_{\ell=1}^h \{z_1^\ell, \dots, z_{\#B_\ell}^\ell\} = \{x_2, \dots, x_{\#A}\} \cup \{y_1, \dots, y_m\}$,*
2. *connected (C) if and only if, for each $\ell \in [1, h]$ there exists an interaction atom in ψ that contains both z_1^ℓ and a variable from $\{x_1\} \cup \{x_i \mid i \in \lambda_\Delta(A)\}$,*
3. *equationally-restricted (e-restricted or R) if and only if, for every disequality $x \neq y$ from ϕ , we have $\{x, y\} \cap \{x_i \mid i \in \lambda_\Delta(A)\} \neq \emptyset$.*

A SID Δ is progressing (P), connected (C) and e-restricted (R) if and only if each rule in Δ is progressing, connected and e-restricted, respectively.

► **Example 30.** For example, the rules for the $\text{chain}_{h,t}(x_1, x_2)$ predicates from the SID in §1.1 are PCR, but not the rules for $\text{ring}_{h,t}()$ predicates, that are neither progressing nor connected. The latter can be replaced with the following PCR rules:

$$\overline{\text{ring}}_{h,t}(x) \leftarrow \exists y \exists z . [x] @ q * \langle x.out, z.in \rangle * \text{chain}_{h',t'}(z, y) * \langle y.out, x.in \rangle, \text{ for all } h, t \in \mathbb{N}$$

Similarly, rule (2) for the *Node* predicate is PCR, but not rules (1) and (3), from Example 12. In order to obtain a SID that is PCR, these rules can be replaced with, respectively:

$$\text{Root}(n) \leftarrow \exists n_1 \exists \ell_1 \exists r_1 \exists n_2 \exists \ell_2 \exists r_2 . [n] * \langle n.req, n_1.reply, n_2.reply \rangle * \langle r_1.in, \ell_2.out \rangle *$$

$$\text{Node}(n_1, \ell_1, r_1) * \text{Node}(n_2, \ell_2, r_2)$$

$$\text{Node}(n, \ell, r) \leftarrow [n] * \langle n.req, \ell.reply, r.reply \rangle * \langle \ell.in, r.out \rangle * \text{Leaf}(\ell) * \text{Leaf}(r) \quad \text{Leaf}(n) \leftarrow [n] \quad \blacksquare$$

A first property is that PCR SIDs define only tight configurations (Def. 3), a prerequisite for the reduction from Theorem 26:

► **Lemma 31.** *Let Δ be a PCR SID and let $A \in \text{Def}(\Delta)$ be a predicate. Then, for any Δ -model (γ, ν) of $A(x_1, \dots, x_{\#A})$, the configuration γ is tight.*

The last restriction for the decidability of entailments relates to the degree of the models of a predicate atom. The degree of a configuration is defined in analogy with the degree of a graph as the maximum number of interactions involving a component:

► **Definition 32.** *The degree of a configuration $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ is defined as $\delta(\gamma) \stackrel{\text{def}}{=} \max_{c \in \mathcal{C}} \delta_c(\gamma)$, where $\delta_c(\gamma) \stackrel{\text{def}}{=} \|\{(c_1, p_1, \dots, c_n, p_n) \in \mathcal{I} \mid c = c_i, i \in [1, n]\}\|$.*

For instance, the configuration of the system from Fig. 1 (a) has degree two. The *degree boundedness* problem $\text{DegreeBound}[\Delta, A]$ asks, given a predicate A and a SID Δ , if the set $\{\delta(\gamma) \mid \gamma \models_{\Delta} \exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})\}$ is finite. This problem is decidable [4, Theorem 3]. The entailment problem $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$ is known to be decidable for PCR SIDs Δ , provided, moreover, that $\text{DegreeBound}[\Delta, A]$ holds:

► **Theorem 33 ([4]).** *The entailment problem $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, where Δ is PCR and $\text{DegreeBound}[\Delta, A]$ has a positive answer, is in 2EXP, if $\#(\Delta)$ and $N(\Delta)$ are bounded by constants and in 4EXP, otherwise.*

Back to the havoc invariance problem, we give first a lower bound using a reduction from the entailment problem $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, where Δ is a PCR SID and $\text{HavocInv}[\Delta, A]$ has a positive answer. To the best of our efforts, we could not prove that the entailment problem is 2EXP-hard under the further assumption that $\#(\Delta)$ is bounded by a constant, which leaves the question of a matching lower bound for the havoc invariance problem open, in this case.

► **Lemma 34.** *The $\text{HavocInv}[\Delta, A]$ problem for PCR SIDs Δ , such that $\text{DegreeBound}[\Delta, A]$ has a positive answer, is 2EXP-hard.*

The main result of this section is a consequence of Theorems 26 and 33. In the absence of a constant bound on the parameters $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$, the entailment resulting from the reduction (Theorem 26) is of simply exponential size in the input and the time complexity of solving the entailments is 4EXP (Theorem 33), yielding a 5EXP upper bound:

► **Theorem 35.** *The $\text{HavocInv}[\Delta, A]$ problem, for PCR SIDs such that $\text{DegreeBound}[\Delta, A]$ has a positive answer is in 2EXP, if $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$ are bounded by constants and in 5EXP, otherwise.*

5 Conclusions

We have considered a logic for describing sets of configurations of parameterized concurrent systems, with user-defined network topology. The havoc invariance problem asks whether a given formula in the logic is invariant under the execution of the system starting from each configuration that is a model of a formula. An algorithm for this problem uses a many-one reduction to the entailment problem, thus leveraging from earlier results on the latter problem. We study the decidability and complexity of the havoc invariance problem and show that a doubly-exponential algorithm exists for a fairly general fragment of the logic, that encompasses all our examples. This result is relevant for automating the generation of correctness proofs for reconfigurable systems, that change the network topology at runtime.

References

- 1 Emma Ahrens, Marius Bozga, Radu Iosif, and Joost-Pieter Katoen. Local reasoning about parameterized reconfigurable distributed systems. *CoRR*, abs/2107.05253, 2021.

- 560 2 Yehoshua Bar-Hillel, Micha Perles, and Eli Shamir. On formal properties of simple phrase
561 structure grammars. *Sprachtypologie und Universalienforschung*, 14:143–172, 1961.
- 562 3 Roderick Bloem, Swen Jacobs, Ayrat Khalimov, Igor Konnov, Sasha Rubin, Helmut Veith, and
563 Josef Widder. *Decidability of Parameterized Verification*. Synthesis Lectures on Distributed
564 Computing Theory. Morgan & Claypool Publishers, 2015.
- 565 4 Marius Bozga, Lucas Bueri, and Radu Iosif. Decision problems in a logic for reasoning about
566 reconfigurable distributed systems. *arXiv 2202.09637, cs.LO*, 2022.
- 567 5 Marius Bozga, Javier Esparza, Radu Iosif, Joseph Sifakis, and Christoph Welzel. Structural
568 invariants for the verification of systems with parameterized architectures. In *Tools and*
569 *Algorithms for the Construction and Analysis of Systems - 26th International Conference,*
570 *TACAS 2020*, volume 12078 of *LNCS*, pages 228–246. Springer, 2020.
- 571 6 Marius Bozga and Radu Iosif. Specification and safety verification of parametric hierarchical
572 distributed systems. In *Formal Aspects of Component Software - 17th International Conference,*
573 *FACS 2021, Virtual Event, October 28-29, 2021, Proceedings*, volume 13077 of *Lecture Notes*
574 *in Computer Science*, pages 95–114. Springer, 2021.
- 575 7 Antonio Bucchiarone and Juan P. Galeotti. Dynamic software architectures verification using
576 dynalloy. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.*, 10, 2008.
- 577 8 Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local action and abstract
578 separation logic. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007),*
579 *10-12 July 2007, Wroclaw, Poland, Proceedings*, pages 366–378. IEEE Computer Society, 2007.
580 doi:10.1109/LICS.2007.30.
- 581 9 Byron Cook, Christoph Haase, Joël Ouaknine, Matthew J. Parkinson, and James Worrell.
582 Tractable reasoning in a fragment of separation logic. In *CONCUR*, volume 6901 of *Lecture*
583 *Notes in Computer Science*, pages 235–249. Springer, 2011.
- 584 10 Julien Dormoy, Olga Kouchnarenko, and Arnaud Lanoix. Using temporal logic for dynamic
585 reconfigurations of components. In Luís Soares Barbosa and Markus Lumpe, editors, *Formal*
586 *Aspects of Component Software - 7th International Workshop, FACS 2010*, volume 6921 of
587 *Lecture Notes in Computer Science*, pages 200–217. Springer, 2010.
- 588 11 Antoine El-Hokayem, Marius Bozga, and Joseph Sifakis. A temporal configuration logic for
589 dynamic reconfigurable systems. In Chih-Cheng Hung, Jiman Hong, Alessio Bechini, and
590 Eunjee Song, editors, *SAC ’21: The 36th ACM/SIGAPP Symposium on Applied Computing,*
591 *Virtual Event, Republic of Korea, March 22-26, 2021*, pages 1419–1428. ACM, 2021. doi:
592 10.1145/3412841.3442017.
- 593 12 Klaus-Tycho Foerster and Stefan Schmid. Survey of reconfigurable data center networks:
594 Enablers, algorithms, complexity. *SIGACT News*, 50(2):62–79, 2019.
- 595 13 Dan Hirsch, Paolo Inverardi, and Ugo Montanari. Graph grammars and constraint solving for
596 software architecture styles. In *Proceedings of the Third International Workshop on Software*
597 *Architecture, ISAW ’98*, page 69–72, New York, NY, USA, 1998. Association for Computing
598 Machinery. doi:10.1145/288408.288426.
- 599 14 Yonit Kesten, Amir Pnueli, Elad Shahar, and Lenore D. Zuck. Network invariants in action. In
600 *CONCUR 2002 - Concurrency Theory, 13th International Conference*, volume 2421 of *LNCS*,
601 pages 101–115. Springer, 2002.
- 602 15 Arnaud Lanoix, Julien Dormoy, and Olga Kouchnarenko. Combining proof and model-checking
603 to validate reconfigurable architectures. *Electron. Notes Theor. Comput. Sci.*, 279(2):43–57,
604 2011.
- 605 16 Daniel Le Metayer. Describing software architecture styles using graph grammars. *IEEE*
606 *Transactions on Software Engineering*, 24(7):521–533, 1998. doi:10.1109/32.708567.
- 607 17 David Lesens, Nicolas Halbwachs, and Pascal Raymond. Automatic verification of param-
608 eterized linear networks of processes. In *The 24th ACM SIGPLAN-SIGACT Symposium on*
609 *Principles of Programming Languages*, pages 346–357. ACM Press, 1997.
- 610 18 Mohammad Noormohammadpour and Cauligi S. Raghavendra. Datacenter traffic control:
611 Understanding techniques and tradeoffs. *IEEE Commun. Surv. Tutorials*, 20(2):1492–1525,
612 2018.

- 19 John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, pages 55–74. IEEE Computer Society, 2002. doi:10.1109/LICS.2002.1029817.
- 20 Ze'ev Shtadler and Orna Grumberg. Network grammars, communication behaviors and automatic verification. In Joseph Sifakis, editor, *Automatic Verification Methods for Finite State Systems, International Workshop*, volume 407 of *LNCS*, pages 151–165. Springer, 1989.
- 21 Pierre Wolper and Vinciane Lovinfosse. Verifying properties of large sets of processes with network invariants. In *Automatic Verification Methods for Finite State Systems, International Workshop*, volume 407 of *LNCS*, pages 68–80. Springer, 1989.

A Proofs

► **Lemma 13.** Let $t \in \mathbb{T}(\tilde{\Sigma})$ be a tree, such that $\Psi^\epsilon(t)$ is tight, (γ, ν) be a model of $\Psi^\epsilon(t)$ and y^v, z^w be two variables that occur in a component and interaction atom of $\Psi^\epsilon(t)$, respectively. Then $\nu(y^v) = \nu(z^w)$ if and only if there exists a walk u_1, \dots, u_n in t and variables $y^v = x_{i_1}^{u_1}, \dots, x_{i_n}^{u_n} = z^w$, such that either $x_{i_j}^{u_j}$ and $x_{i_{j+1}}^{u_{j+1}}$ are the same variable, or the equality atom $x_{i_j}^{u_j} = x_{i_{j+1}}^{u_{j+1}}$ occurs in $\Psi^\epsilon(t)$, for all $j \in [1, n-1]$.

Proof of Lemma 13. By Def. 11, each equality atom within $\Psi^\epsilon(t)$ is of the form $x^u = y^v$, such that either $u = v$, u is the parent of v , or u is a child of v in t . Suppose, for a contradiction, that there is no walk $u_1, \dots, u_n \in \text{dom}(t)$ with the above property. Then we build a loose model (γ', ν') of $\Psi^\epsilon(t)$. Let $\gamma = (\mathcal{C}, \mathcal{I}, \rho)$, $c = \nu(x^u)$ and $\bar{c} \in \mathbb{C} \setminus (\mathcal{C} \cup \{(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}\})$ be a component not occurring in γ . We define (γ', ν') as follows:

- $\gamma' = (\mathcal{C}, \mathcal{I}', \rho)$, where $\mathcal{I}' = \{(c_i, p_i)_{i \in [1, n]} \in \mathcal{I} \mid c \notin \{c_1, \dots, c_n\}\} \cup \{(c'_i, p_i)_{i \in [1, n]} \mid (c_i, p_i)_{i \in [1, n]} \in \mathcal{I}, c \in \{c_1, \dots, c_n\}, c'_i = c_i \text{ if } c_i \neq c \text{ and } c'_i = \bar{c} \text{ otherwise}\}$.
 - for each variable $x \in \mathbb{V}$, we have $\nu'(x) = \bar{c}$ if $\Psi^\epsilon(t) \models z^w = x$ and $\nu'(x) = \nu(x)$, otherwise.
- Clearly, γ' is loose, because \bar{c} occurs in an interaction from \mathcal{I}' but $\bar{c} \notin \mathcal{C}$. We conclude by showing $(\gamma', \nu') \models \Psi^\epsilon(t)$, by induction on the structure of t . However, this contradicts with the assumption that $\Psi^\epsilon(t)$ is tight, which concludes the proof. ◀

► **Lemma 16.** For any predicate $A \in \text{Def}(\Delta)$, configuration γ , store ν and node $u \in \mathbb{N}^*$, we have $\gamma \models_\Delta^\nu A(x_1^u, \dots, x_{\#A}^u)$ if and only if $\gamma \models^\nu \Phi^u(t)$, for some tree $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$.

Proof of Lemma 16. For a store ν , we denote by $\nu[x_i/y_i]_{i \in [1, n]}$ the store that maps x_i to $\nu(y_i)$ and agrees with ν everywhere else.

“ \Rightarrow ” By induction on the definition of the satisfaction relation $\gamma \models_\Delta^\nu A(x_1^u, \dots, x_{\#A}^u)$. Assume that there exists a rule $r \in \Delta$ of the form $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * \bigstar_{\ell \in [1, h]} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, where ϕ is a qpf formula, and a store ν' such that $\nu'(x_j) = \nu(x_j^u)$ for all $j \in [1, \#]$, where $\gamma = \gamma_0 \bullet \dots \bullet \gamma_h$ is such that $\gamma_0 \models^{\nu'} \phi$ and $\gamma_\ell \models_\Delta^{\nu'} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for all $\ell \in [1, h]$. Hence, we obtain $\gamma_\ell \models_\Delta^{\nu'_\ell} B_\ell(x_1^{u \cdot \ell}, \dots, x_{\#B_\ell}^{u \cdot \ell})$, where $\nu'_\ell \stackrel{\text{def}}{=} \nu'[x_i^{u \cdot \ell}/z_i^\ell]_{i \in [1, \#B_\ell]}$, for all $\ell \in [1, h]$. By the induction hypothesis, there exist trees $t_\ell \in \mathcal{L}_{q_{B_\ell}}(\mathcal{A}_\Delta)$, such that $\gamma_\ell \models_\Delta^{\nu'_\ell} \Phi^{u \cdot \ell}(t_\ell)$, for each $\ell \in [1, h]$. We define the tree t as:

- $\text{dom}(t) \stackrel{\text{def}}{=} \{\epsilon\} \cup \bigcup_{\ell \in [1, h]} \ell \cdot \text{dom}(t_\ell)$,
- $t(\epsilon) \stackrel{\text{def}}{=} \langle \psi, \#A, \#B_1, \dots, \#B_h \rangle$, where $\psi \stackrel{\text{def}}{=} \exists y_1 \dots \exists y_m \cdot \left(\phi * \bigstar_{\substack{\ell \in [1, h] \\ i \in [1, \#B_\ell]}} \tilde{z}_i^{(\ell)} = z_i^\ell \right) [x_i/\tilde{x}_i]_{i \in [1, \#A]}$,
- $t_\ell \stackrel{\text{def}}{=} t_\ell$, for all $\ell \in [1, h]$.

By the definition of \mathcal{A}_Δ and t , we obtain $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$ and we are left with proving that $\gamma \models_\Delta^\nu \Phi^u(t)$. To this end, we define the store ν'' , such that:

- 656 $\blacksquare \nu''(x_j^u) = \nu(x_j^u)$, for all $j \in [1, \#A]$,
- 657 $\blacksquare \nu''(y_k^u) = \nu'(y_k^u)$, for all $k \in [1, m]$,
- 658 $\blacksquare \nu''(x_i^{u \cdot \ell}) = \nu'_\ell(x_i^{u \cdot \ell})$, for all $i \in [1, B_\ell]$ and $\ell \in [1, h]$.

659 By the definition of ν'' , we obtain:

$$660 \quad \gamma \models_{\Delta}^{\nu''} \psi[\tilde{x}_i/x_i]_{i \in [1, \#A]} [\tilde{z}_i^{(\ell)}/x_i^{u \cdot \ell}]_{\substack{i \in [1, \#B_\ell] \\ \ell \in [1, h]}} [y_k/y_k^u]_{k \in [1, m]} * *_{\ell \in [1, h]} \Phi^{u \cdot \ell}(\mathbf{t}_\ell)$$

661 leading to $\gamma \models_{\Delta}^{\nu} \Phi^u(\mathbf{t})$, because ν'' agrees with ν over $x_1^u, \dots, x_{\#A}^u$.

662 " \Leftarrow " Let $\mathbf{t} \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$ be a tree such that $\gamma \models^{\nu} \Phi^u(\mathbf{t})$. We proceed by induction on the
 663 structure of \mathbf{t} . Let be $\alpha_r \stackrel{\text{def}}{=} \mathbf{t}(\epsilon)$ the label of the root of \mathbf{t} , and $\mathbf{t}_1, \dots, \mathbf{t}_h$ be the subtrees
 664 rooted in the children of the root of \mathbf{t} (h can be equal to 0 if \mathbf{t} consists of a leaf). Since
 665 $\mathbf{t} \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$, there exists a rule $r \in \Delta$ and a transition $\alpha_r(q_{B_1}, \dots, q_{B_h}) \rightarrow q_A$ in \mathcal{A}_Δ , where r
 666 is of the form $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * *_{\ell \in [1, h]} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for a qpf formula
 667 ϕ , such that $\mathbf{t}_\ell \in \mathcal{L}_{q_{B_\ell}}(\mathcal{A}_\Delta)$, for all $\ell \in [1, h]$. By Def. 14, we have:

$$668 \quad \alpha_r = \left\langle \exists y_1 \dots \exists y_m \cdot \left(\phi * *_{\substack{\ell \in [1, h] \\ i \in [1, \#B_\ell]}} \tilde{z}_i^{(\ell)} = z_i^\ell \right) [x_i/\tilde{x}_i]_{i \in [1, \#A]}, \#A, \#B_1, \dots, \#B_h \right\rangle$$

669 Since $\gamma \models^{\nu} \Phi^u(\mathbf{t})$, by Def. 11, there exists a store ν' that agrees with ν over $x_1^u, \dots, x_{\#A}^u$ and
 670 configurations $\gamma = \gamma_0 \bullet \dots \bullet \gamma_h$, such that:

$$671 \quad \blacksquare \gamma_0 \models^{\nu'} \left(\phi * *_{\substack{\ell \in [1, h] \\ i \in [1, \#B_\ell]}} x_i^{u \cdot \ell} = z_i^\ell \right) [x_i/x_i^u]_{i \in [1, \#A]} [y_k/y_k^u]_{k \in [1, m]},$$

$$672 \quad \blacksquare \gamma_\ell \models^{\nu'} \Phi^{u \cdot \ell}(\mathbf{t}_\ell), \text{ for all } \ell \in [1, h].$$

673 Since $\mathbf{t}_\ell \in \mathcal{L}_{q_{B_\ell}}(\mathcal{A}_\Delta)$, by the inductive hypothesis we obtain $\gamma_\ell \models_{\Delta}^{\nu'} B_\ell(x_1^{u \cdot \ell}, \dots, x_{\#B_\ell}^{u \cdot \ell})$, for
 674 all $\ell \in [1, h]$. Let us define the store:

$$675 \quad \nu'' \stackrel{\text{def}}{=} \nu' [z_i^\ell/x_i^{u \cdot \ell}]_{\substack{\ell \in [1, h] \\ i \in [1, \#B_\ell]}} [y_k/y_k^u]_{k \in [1, m]}$$

676 We obtain $\gamma \models_{\Delta}^{\nu''} \phi * *_{\ell \in [1, h]} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, hence $\gamma \models_{\Delta}^{\nu} \exists y_1 \dots \exists y_m \cdot \phi * *_{\ell \in [1, h]} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$,
 677 leading to $\gamma \models_{\Delta}^{\nu} A(x_1^u, \dots, x_{\#A}^u)$. \blacktriangleleft

678 **► Lemma 19.** For any state $q \in \mathcal{S}$, configuration γ , store ν and node $u \in \mathbb{N}^*$, we have
 679 $\gamma \models_{\Delta_A}^{\nu} A_q(x_1^u, \dots, x_{\#A_q}^u)$ if and only if $\gamma \models^{\nu} \Phi^u(\mathbf{t})$, for some tree $\mathbf{t} \in \mathcal{L}_q(\mathcal{A})$.

680 **Proof of Lemma 19.** Let $\bar{\mathcal{A}}$ denote the automaton \mathcal{A}_{Δ_A} and \bar{q} denote the state q_{A_q} of $\bar{\mathcal{A}}$
 681 (Def. 14). By Lemma 16, we have $\gamma \models_{\Delta_A}^{\nu} A_q(x_1^u, \dots, x_{\#A_q}^u)$ if and only if $\gamma \models^{\nu} \Phi^u(\bar{\mathbf{t}})$, for
 682 some tree $\bar{\mathbf{t}} \in \mathcal{L}_{\bar{q}}(\bar{\mathcal{A}})$. It is sufficient to prove that for each tree $\bar{\mathbf{t}} \in \mathcal{L}_{\bar{q}}(\bar{\mathcal{A}})$ there exists a
 683 tree $\mathbf{t} \in \mathcal{L}_q(\mathcal{A})$ such that $\Phi^u(\bar{\mathbf{t}})$ is equivalent to $\Phi^u(\mathbf{t})$, and viceversa. The last point is a
 684 consequence of the one-to-one mapping between the transitions of \mathcal{A} and those of $\bar{\mathcal{A}}$:

$$\begin{aligned} & \langle \phi, a_0, \dots, a_h \rangle (q_1, \dots, q_h) \rightarrow q_0 \in \delta \\ & \iff \text{by Def. 18} \\ & A_{q_0}(x_1, \dots, x_{a_0}) \leftarrow \exists y_1^1 \dots \exists y_{a_1}^1 \dots \exists y_1^h \dots \exists y_{a_h}^h \cdot \phi[\tilde{x}_i/x_i]_{i \in [1, a_0]} [\tilde{z}_i^{(\ell)}/y_i^\ell]_{\substack{\ell \in [1, h] \\ i \in [1, a_\ell]}} * *_{\ell \in [1, h]} A_{q_\ell}(y_1^\ell, \dots, y_{a_\ell}^\ell) \in \Delta_A \\ & \iff \text{by Def. 14 and quantifier elimination} \\ & \langle \phi, a_0, \dots, a_h \rangle (\bar{q}_1, \dots, \bar{q}_h) \rightarrow \bar{q}_0 \in \delta_{\Delta_A} \end{aligned}$$

686 \blacktriangleleft

687 ► **Lemma 20.** *HavocInv $[\Delta, A]$ has a positive answer if and only if, for all $\gamma, \gamma' \in \Gamma$ and each*
 688 *store ν , such that $\gamma \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ and $\gamma \rightsquigarrow \gamma'$, it is the case that $\gamma' \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$.*

689 **Proof of Lemma 20.** Let $\text{StepInv}[\Delta, A]$ be the decision problem from the statement of the
 690 Lemma and prove that $\text{StepInv}[\Delta, A] \iff \text{HavocInv}[\Delta, A]$. “ \Rightarrow ” Let $\gamma, \gamma' \in \Gamma$ and ν , such
 691 that $\gamma \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ and $\gamma \rightsquigarrow^* \gamma'$. Then there exist configurations $\gamma = \gamma_0, \dots, \gamma_n = \gamma'$,
 692 such that $\gamma_i \rightsquigarrow \gamma_{i+1}$, for all $i \in [0, n-1]$. Proving $\gamma' \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ goes by induction
 693 over $n \geq 0$. For the base case $n = 0$, we have $\gamma = \gamma'$ and there is nothing to prove. For the
 694 inductive step $n > 0$, we have $\gamma_{n-1} \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ by the inductive hypothesis. Then
 695 $\gamma_{n-1} \rightsquigarrow \gamma'$ and $\gamma' \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ follows, by the hypothesis $\text{StepInv}[\Delta, A]$. “ \Leftarrow ” Trivial,
 696 because $\rightsquigarrow \subseteq \rightsquigarrow^*$. ◀

697 ► **Lemma 23.** *For each tree $t \in L$, such that $\Phi^{\epsilon}(t)$ is tight, configurations $\gamma = (\mathcal{C}, \mathcal{I}, \varrho), \gamma' \in \Gamma$*
 698 *and store ν , such that $\gamma \models^{\nu} \Phi^{\epsilon}(t)$ and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some $c_1, \dots, c_n \in \mathcal{C}$ and*
 699 *$(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$, there exists a tree $t' \in \mathcal{T}_{(p_1, \dots, p_n)}(L)$, such that $\gamma' \models^{\nu} \Phi^{\epsilon}(t')$.*

700 **Proof of Lemma 23.** Let $\gamma \stackrel{\text{def}}{=} (\mathcal{C}, \mathcal{I}, \varrho)$ and $t(u) \stackrel{\text{def}}{=} \langle \phi^u, a_0^u, \dots, a_h^u \rangle$, for each $u \in \text{dom}(t)$.
 701 Because $\gamma \models^{\nu} \Phi^{\epsilon}(t)$, by Def. 11, there exists a store ν' that agrees with ν over $\text{fv}(\Phi^{\epsilon}(t))$,
 702 such that $\gamma \models^{\nu'} \Psi^{\epsilon}(t)$. The proof is split into the following steps:

703 (1) Since there exists an interaction $(c_i, p_i)_{i \in [1, n]}$, such that $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, we define the
 704 tree t' based on this interaction. By Def. 7, it must be the case that $c_i \in \mathcal{C}$, for all $i \in [1, n]$,
 705 $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$ and $\gamma' = (\mathcal{C}, \mathcal{I}, \varrho')$, where $\varrho' = \varrho[c_i \leftarrow q'_i]_{i \in [1, n]}$ and $q_i \xrightarrow{p_i} q'_i$ is a transition
 706 in \mathbb{B} , for each $i \in [1, n]$. For all $i \in [1, n]$, since $c_i \in \mathcal{C}$, there exists a unique node $w_i \in \text{dom}(t)$
 707 and a variable $\xi_i \in \{x_{k_i}^{w_i}, y_{\ell_i}^{w_i}\}$, for some $k_i, \ell_i \geq 1$, such that $\nu'(\xi_i) = c_i$ and $[\xi_i]@q_i$ occurs
 708 in $\Psi^{\epsilon}(t)$, for some $q_i \in \mathcal{Q}$. Then the tree t' is defined as $\text{dom}(t') = \text{dom}(t)$ and:
 709 ■ $t'(w_i) = \langle \psi^{w_i}, a_0^{w_i}, \dots, a_h^{w_i} \rangle$, where ψ^{w_i} is obtained from ϕ^{w_i} by replacing the atom
 710 $[\tilde{x}_{k_i}]@q_i$ (resp. $[y_{\ell_i}]@q_i$) with $[\tilde{x}_{k_i}]@q'_i$ (resp. $[y_{\ell_i}]@q'_i$),
 711 ■ $t'(u) = t(u)$, for all $u \in \text{dom}(t) \setminus \{w_1, \dots, w_n\}$.

712 The check of $\gamma' \models^{\nu'} \Psi^{\epsilon}(t')$ follows from the definition of t' and $\varrho' = \varrho[c_i \leftarrow q'_i]_{i \in [1, n]}$.

713 (2) We prove that $(t, t') \in \mathcal{L}(\mathcal{T}_{(p_1, \dots, p_n)})$ by building an accepting run $\pi : \text{dom}(t) \rightarrow_{\text{fin}} \mathcal{S}_{(p_1, \dots, p_n)}$
 714 of $\mathcal{T}_{(p_1, \dots, p_n)}$ over (t, t') . For each node $u \in \text{dom}(t)$, the formula $\pi(u)$ is the separating
 715 conjunction of the following equality atoms:

- 716 ■ $\tilde{x}_i = \tilde{x}_j$, where $\Psi^u(t_u) \models x_i^u = x_j^u$, for all $i, j \in [1, a_0^u]$,
- 717 ■ $\tilde{x}_j = \tilde{b}_i$, where $\Psi^u(t_u) \models x_j^u = \xi_i$, for all $i \in [1, n]$ and $j \in [1, a_0^u]$,
- 718 ■ $\tilde{x}_j = \tilde{e}_i$, where $\Psi^u(t_u) \models x_j^u = \zeta_i$, for all $i \in [1, n]$ and $j \in [1, a_0^u]$,
- 719 ■ $\tilde{b}_i = \tilde{e}_i$, where $\Psi^u(t_u) \models \xi_i = \zeta_i$, for all $i \in [1, n]$.

720 It is easy to check that π is indeed a run of $\mathcal{T}_{(p_1, \dots, p_n)}$, by Def. 21. We are left with proving
 721 that π is accepting i.e., that $\pi(\epsilon) \in \mathcal{F}_{(p_1, \dots, p_n)}$. Let $\langle \zeta_1.p_1, \dots, \zeta_n.p_n \rangle$, where $\zeta_i \in \{x_{k_i}^{w_0}, y_{\ell_i}^{w_0}\}$,
 722 for some $k_i, \ell_i \geq 1$, be the interaction atom of $\Psi^{\epsilon}(t)$, such that $\nu'(\zeta_i) = c_i$, for all $i \in [1, n]$.
 723 Since $\Psi^{\epsilon}(t)$ is satisfiable, there is exactly one such node in $\text{dom}(t)$. Because $\Phi^{\epsilon}(t)$ is tight, the
 724 formula $\Psi^{\epsilon}(t)$ is tight and, by Lemma 13, there exists an equality walk between the nodes w_i
 725 and w_0 , for all $i \in [1, n]$. By the definition of π , we obtain $\pi(\epsilon) \models \tilde{b}_i = \tilde{e}_i$ for all $i \in [1, n]$,
 726 thus $\pi(\epsilon) \in \mathcal{F}_{(p_1, \dots, p_n)}$. ◀

728 ► **Lemma 24.** *For each tree $t' \in \mathcal{T}_{(p_1, \dots, p_n)}(L)$, configuration $\gamma' \in \Gamma$ and store ν , such that*
 729 *$\gamma' \models^{\nu} \Phi^{\epsilon}(t')$, there exists a configuration $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ and a tree $t \in L$, such that $\gamma \models^{\nu} \Phi^{\epsilon}(t)$*
 730 *and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some $c_1, \dots, c_n \in \mathcal{C}$ and $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$.*

Proof of Lemma 24. Let $\gamma' \stackrel{\text{def}}{=} (\mathcal{C}, \mathcal{I}, \varrho')$. Because $\gamma' \models^\nu \Phi^\epsilon(\mathbf{t}')$, by Def. 11, there exists a store ν' , such that $\gamma' \models^{\nu'} \Psi^\epsilon(\mathbf{t}')$. Moreover, since $\mathbf{t}' \in \mathcal{T}_{(p_1, \dots, p_n)}(L)$, there exists a tree $\mathbf{t} \in L$, such that $\text{dom}(\mathbf{t}) = \text{dom}(\mathbf{t}')$, $(\mathbf{t}, \mathbf{t}') \in \mathcal{L}(\mathcal{T}_{(p_1, \dots, p_n)})$ and let $\pi : \text{dom}(\mathbf{t}) \rightarrow \mathcal{S}_{(p_1, \dots, p_n)}$ be an accepting run of $\mathcal{T}_{(p_1, \dots, p_n)}$ over $(\mathbf{t}, \mathbf{t}')$. By Def. 21, for each $u \in \text{dom}(\mathbf{t})$, the state $\pi(u)$ is the separating conjunction of all equalities entailed by $\Psi^u(\mathbf{t}|_u)$. Since π is accepting, we have $\pi(\epsilon) \models \bigstar_{i \in [1, n]} \tilde{b}_i = \tilde{e}_i$. By Def. 21, there exists nodes $u, v \in \text{dom}(\mathbf{t})$ and variables ξ_i and ζ_i , such that, for all $i \in [1, n]$, the component atom $[\xi_i]@q_i$ (resp. $[\xi_i]@q'_i$) occurs in $\mathbf{t}(u)$ (resp. $\mathbf{t}'(v)$), for a transition $q_i \xrightarrow{P_i} q'_i$ of \mathbb{B} , and the interaction atom $\langle \zeta_1.p_1, \dots, \zeta_n.p_n \rangle$ occurs in both $\mathbf{t}(v)$ and in $\mathbf{t}'(v)$. We define the state map $\varrho = \varrho'[\nu'(\xi_i) \leftarrow q_i]_{i \in [1, n]}$ and let $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$. We are left with proving that $\gamma \models \Phi^\epsilon(\mathbf{t})$ and $\gamma \xrightarrow{(\nu'(\xi_i), p_i)_{i \in [1, n]}} \gamma'$, which are both easy checks.

► **Theorem 26.** Assuming that $A(x_1, \dots, x_{\#A})$ is a Δ -tight formula, each instance $\text{HavocInv}[\Delta, A]$ of the havoc invariance problem can be reduced to a set $\{\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\Delta \cup \Delta} A(x_1, \dots, x_{\#A})\}_{i=1}^p$ of entailments, where $\Delta \approx \bar{\Delta}$, for an arity-preserving equivalence relation $\sim \subseteq \mathbb{A} \times \mathbb{A}$, such that $\bar{A}_i \sim A$, for all $i \in [1, p]$. The reduction is polynomial, if $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$ are bounded by constants and simply exponential, otherwise.

Proof of Theorem 26. Let $\text{HavocInv}[\Delta, A]$ be an instance of the havoc invariance problem. We denote by $I(\Delta)$ the set of types $\tau = (p_1, \dots, p_n)$ of an interaction atom $\langle y_1.p_1, \dots, y_n.p_n \rangle$ that occurs in Δ . Let \mathcal{A}_Δ be the tree automaton from Def. 14 and let:

$$\mathcal{T} = \left(\Sigma_\Delta^2, \bigcup_{\tau \in I(\Delta)} \mathcal{S}_\tau, \bigcup_{\tau \in I(\Delta)} \mathcal{F}_\tau, \bigcup_{\tau \in I(\Delta)} \delta_\tau \right)$$

be the automata-theoretic union of the transducers \mathcal{T}_τ , taken over all $\tau \in I(\Delta)$. Let $\bar{\mathcal{A}}$ be the tree automaton that recognizes the language $\mathcal{T}(\mathcal{L}_{q_A}(\mathcal{A}_\Delta))$. The states of $\bar{\mathcal{A}}$ are pairs of the form (q_B, φ) , where $B \in \text{Def}(\Delta)$ and $\varphi \in \mathcal{S}_\tau$, for some $\tau \in I(\Delta)$. Moreover, the final states of $\bar{\mathcal{A}}$ are of the form (q_A, φ) , where $\varphi \in \mathcal{F}_\tau$, for some $\tau \in I(\Delta)$. Let $\bar{\Delta} \stackrel{\text{def}}{=} \Delta_{\bar{\mathcal{A}}}$ be the SID from Def. 18 relative to $\bar{\mathcal{A}}$ and let $\{\bar{A}_i \mid i \in [1, p]\} \stackrel{\text{def}}{=} \{A_{(q_A, \varphi)} \mid \varphi \in \mathcal{F}_\tau, \tau \in I(\Delta)\}$ be the set of predicates corresponding to the final states of $\bar{\mathcal{A}}$. It is easy to check that $\bar{\mathcal{A}}$ is SID-compatible (Def. 17) and that $\# \bar{A}_i = \# A$, for all $i \in [1, p]$. For a store ν , we denote by $\nu[x_i/y_i]_{i \in [1, n]}$ the store that maps x_i to $\nu(y_i)$ and agrees with ν everywhere else. We prove that $\text{HavocInv}[\Delta, A]$ has a positive answer if and only if $\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\bar{\Delta} \cup \Delta} A(x_1, \dots, x_{\#A})$ holds, for all $i \in [1, p]$.

“ \Rightarrow ” Let γ' be a configuration and ν be a store, such that $\gamma' \models_{\bar{\Delta}}^\nu \bar{A}_i(x_1, \dots, x_{\#A})$, for some $i \in [1, p]$. Let $\nu' \stackrel{\text{def}}{=} \nu[x_i^\epsilon/x_i]_{i \in [1, \#A]}$ be a store, such that $\gamma' \models_{\bar{\Delta}}^\nu \bar{A}_i(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$. By Lemma 19, there exists a tree $\mathbf{t}' \in \mathcal{L}_{(q_A, \varphi)}(\bar{\mathcal{A}})$, for a final state (q_A, φ) of $\mathcal{T}_{(p_1, \dots, p_n)}$, for some interaction type $(p_1, \dots, p_n) \in I(\Delta)$, such that $\gamma' \models^{\nu'} \Phi^\epsilon(\mathbf{t}')$. By Lemma 24, there exists a tree $\mathbf{t} \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$ and a configuration γ , such that $\gamma \models^\nu \Phi^\epsilon(\mathbf{t})$ and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some interaction $(c_i, p_i)_{i \in [1, n]}$ from γ . By Lemma 16, we obtain that $\gamma \models_{\bar{\Delta}}^{\nu'} A(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$, leading to $\gamma \models_{\bar{\Delta}}^\nu A(x_1, \dots, x_{\#A})$. By the hypothesis that $\text{HavocInv}[\Delta, A]$ has a positive answer, we obtain that $\gamma' \models_{\bar{\Delta}}^\nu A(x_1, \dots, x_{\#A})$. Since the choices of γ' and i were arbitrary, we obtain $\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\bar{\Delta} \cup \Delta} A(x_1, \dots, x_{\#A})$, for all $i \in [1, p]$.

“ \Leftarrow ” Let γ, γ' be configurations and ν be a store, such that $\gamma \models^\nu A(x_1, \dots, x_{\#A})$ and $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, for some interaction $(c_i, p_i)_{i \in [1, n]}$ from γ . Let $\nu' \stackrel{\text{def}}{=} \nu[x_i^\epsilon/x_i]_{i \in [1, \#A]}$ be

773 a store such that $\gamma \models_{\Delta}^{\nu'} A(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$. By Lemma 16, there exists a tree $t \in \mathcal{L}_{q_A}(\mathcal{A}_\Delta)$,
 774 such that $\gamma \models_{\Delta}^{\nu'} \Phi^\epsilon(t)$. By the assumption in the statement, since $A(x_1, \dots, x_{\#A})$ and hence
 775 $A(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$ is Δ -tight, and every model of $\Phi^\epsilon(t)$ is an Δ -model of $A(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$, we obtain
 776 that $\Phi^\epsilon(t)$ is tight. Since $\gamma \xrightarrow{(c_i, p_i)_{i \in [1, n]}} \gamma'$, there exists a tree $t' \in \mathcal{T}_{(p_1, \dots, p_n)}(\mathcal{L}_{q_A}(\mathcal{A}_\Delta)) =$
 777 $\mathcal{L}_{(q_A, \varphi)}(\overline{\mathcal{A}})$, such that $\gamma' \models_{\Delta}^{\nu'} \Phi^\epsilon(t')$, where φ is a final state of $\mathcal{T}_{(p_1, \dots, p_n)}$, by Lemma 23.
 778 By the definition of $\overline{\Delta}$ and of the predicates $\overline{A}_1, \dots, \overline{A}_p$, by Lemma 16, we obtain $\gamma' \models_{\overline{\Delta}}^{\nu'}$
 779 $\overline{A}_i(x_1^\epsilon, \dots, x_{\#A}^\epsilon)$, for some $i \in [1, p]$. We obtain $\gamma' \models_{\overline{\Delta}}^{\nu'} \overline{A}_i(x_1, \dots, x_{\#A})$ and the hypothesis
 780 $\overline{A}_i(x_1, \dots, x_{\#A}) \models_{\overline{\Delta} \cup \Delta} A(x_1, \dots, x_{\#A})$ yields $\gamma' \models_{\Delta}^{\nu'} A(x_1, \dots, x_{\#A})$. Since the choices of γ
 781 and γ' were arbitrary, we obtain that $\text{HavocInv}[\Delta, A]$ has a positive answer.

782 To prove $\Delta \approx \overline{\Delta}$, consider the relation $B \sim A_{(q_B, \varphi)}$, for all $B \in \mathbb{A}$ and all states
 783 $\varphi \in \bigcup_{\tau \in I(\Delta)} \mathcal{S}_\tau$. Because $\overline{\mathcal{A}}$ is SID-compatible, we have $\#B = \#A_{(q_B, \varphi)}$, for all $B \in \mathbb{A}$,
 784 hence \sim is arity-preserving. By the definition of $\overline{\Delta}$, for each rule $B_0(x_1, \dots, x_{\#B_0}) \leftarrow$
 785 $\exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$ from Δ there exists a rule:

$$786 \quad A_{(q_{B_0}, \varphi_0)} \leftarrow \exists y'_1 \dots \exists y'_p . \psi * \bigstar_{\ell=1}^h A_{(q_{B_\ell}, \varphi_\ell)}(u_1^\ell, \dots, u_{\#B_\ell}^\ell) \quad (4)$$

788 in $\overline{\Delta}$, such that $\exists y_1 \dots \exists y_m . \phi \simeq \exists y'_1 \dots \exists y'_p . \psi$. The last equivalence between predicate-free
 789 formulæ is because $\exists y'_1 \dots \exists y'_p . \psi$ is obtained from $\exists y_1 \dots \exists y_m . \phi$ by changing state atoms
 790 and introducing equality atoms of which one of the variable is existentially quantified (see
 791 Def. 14, Def. 21 and Def. 18). Viceversa, each rule of the form (4) in $\overline{\Delta}$ stems from a rule in
 792 Δ , where the same equivalences hold. We conclude that $\Delta \approx \overline{\Delta}$, by Def. 25.

793 Finally, we compute an upper bound on the time necessary to build $\overline{\Delta}$ from Δ . Note
 794 that the number of interaction atoms, and hence the number of interaction types $\tau \in I(\Delta)$,
 795 that occur in Δ is bounded by $\text{size}(\Delta)$. The number of states in each transducer \mathcal{T}_τ ,
 796 for some $\tau \in I(\Delta)$, is bounded by the number of partitions of $[1, \#(\Delta) + N(\Delta)]$, that is
 797 $2^{\mathcal{O}((\#(\Delta) + N(\Delta)) \cdot \log(\#(\Delta) + N(\Delta)))}$. This is because each state $Eq(\widetilde{\text{TVar}}_\tau)$ corresponds (modulo
 798 logical equivalence) to a partition of the set $\{\tilde{x}_1, \dots, \tilde{x}_{\#A}\} \cup \{\tilde{b}_i, \tilde{c}_i \mid i \in [1, N(\Delta)]\}$ of the
 799 parameters of some predicate $A \in \text{Def}(\Delta)$ to which the variables \tilde{b}_i, \tilde{c}_i , for $i = 1, \dots, n$ are
 800 added and the number of partitions of this set is asymptotically bounded by $(\#A + n)^{(\#A + n)}$.

801 The size of the transducer alphabet is the number of pairs (α, β) , such that $\alpha =$
 802 $\langle \phi, a_0, \dots, a_h \rangle$, $\beta = \langle \psi, a_0, \dots, a_h \rangle$, ϕ stems from a rule in Δ and ψ differs from ϕ by a
 803 renaming of states in at most n state atoms, for $n \leq N(\Delta)$. Let M be the maximum number
 804 of variables that occurs free or bound in a rule in Δ and B be the maximum number of
 805 outgoing transitions $q \xrightarrow{p} q'$ in the behavior $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$. Clearly $M \leq \text{size}(\Delta)$, whereas
 806 B is a constant, because \mathbb{B} is considered to be fixed i.e., not part of the input of the havoc
 807 invariance problem. Then, for each qpf formula that occurs in Δ , the number of alphabet
 808 symbols of the form $(\langle \phi, a_0, \dots, a_h \rangle, \langle \psi, a_0, \dots, a_h \rangle)$ is at most:

$$\begin{aligned} \sum_{i=1}^{N(\Delta)} \binom{M}{i} \cdot B^i &\leq B^{N(\Delta)} \cdot \sum_{i=0}^{N(\Delta)} \binom{M}{i} \\ &\leq B^{N(\Delta)} \cdot \sum_{i=0}^{N(\Delta)} \frac{M^i}{i!} = B^{N(\Delta)} \cdot \sum_{i=0}^{N(\Delta)} \frac{N(\Delta)^i}{i!} \cdot \left(\frac{M}{N(\Delta)}\right)^i \\ 809 &\leq B^{N(\Delta)} \cdot \left(\frac{M}{N(\Delta)}\right)^{N(\Delta)} \cdot \sum_{i=0}^{N(\Delta)} \frac{N(\Delta)^i}{i!} \\ &\leq \left(\frac{B \cdot M \cdot e}{N(\Delta)}\right)^{N(\Delta)} \leq (B \cdot \text{size}(\Delta))^{N(\Delta)} \end{aligned}$$

810 Given alphabet symbols $\alpha = \langle \phi, a_0, \dots, a_h \rangle$, $\beta = \langle \psi, a_0, \dots, a_h \rangle$ and states $\varphi_0, \dots, \varphi_h$,
 811 the time required to decide on the existence of a transition $(\alpha, \beta)(\varphi_1, \dots, \varphi_h) \rightarrow \varphi_0$ is

812 $\mathcal{O}((\text{size}(\phi) + \sum_{i=0}^h \text{size}(\varphi_i))^3) = \mathcal{O}(\text{size}(\Delta)^3 \cdot (\#(\Delta) + \mathbf{N}(\Delta))^3) = \mathcal{O}(\text{size}(\Delta)^3)$, because elim-
 813 inating existential quantifiers from and checking equivalence between separating conjunctions
 814 of equality atoms is cubic in the number of variables, using a standard Floyd-Warshall closure
 815 algorithm. Summing up, the time needed to compute the transducer \mathcal{T} as the union of
 816 $\mathcal{T}_{(p_1, \dots, p_n)}$, for all $(p_1, \dots, p_n) \in \mathbf{I}(\Delta)$ is bounded by:

$$817 \quad \text{size}(\Delta)^{\mathbf{N}(\Delta)+3} \cdot 2^{\mathcal{O}(\mathbf{H}(\Delta) \cdot (\#(\Delta) + \mathbf{N}(\Delta)) \cdot \log(\#(\Delta) + \mathbf{N}(\Delta)))}$$

818 This is because there are at most $(B \cdot \text{size}(\Delta))^{\mathbf{N}(\Delta)} \cdot \|\mathcal{S}_{(p_1, \dots, p_n)}\|^{\mathbf{H}(\Delta)+1}$ transitions of the form
 819 $(\alpha, \beta)(q_1, \dots, q_h) \rightarrow q_0$ in $\mathcal{T}_{(p_1, \dots, p_n)}$ and their enumeration requires time simply exponential
 820 in $\mathbf{H}(\Delta)$. Since the translation between a SID and a tree automaton takes linear time,
 821 the above is an upper bound for the reduction of the $\text{HavocInv}[\Delta, \mathbf{A}]$ instance of the havoc
 822 invariance problem to the set $\{\bar{\mathbf{A}}_i(x_1, \dots, x_{\# \mathbf{A}}) \models_{\Delta \cup \bar{\Delta}} \mathbf{A}(x_1, \dots, x_{\# \mathbf{A}})\}_{i=1}^p$ of instances of the
 823 entailment problem. The reduction is thus polynomial, if $\#(\Delta)$, $\mathbf{N}(\Delta)$ and $\mathbf{H}(\Delta)$ are constant
 824 and simply exponential, otherwise. \blacktriangleleft

825 **► Theorem 27.** *The $\text{HavocInv}[\Delta, \mathbf{A}]$ problem is undecidable.*

826 **Proof of Theorem 27.** A *context-free grammar* $G = (\Sigma, N, T, S, \Delta)$ consists of a finite set
 827 N of *nonterminals*, a finite set T of words over a finite alphabet Σ , called *terminals*, a start
 828 symbol $S \in N$ and a finite set Δ of *productions* of the form $A \rightarrow w$, where $A \in N$ and
 829 $w \in (N \cup T)^*$. Given finite strings $u, v \in (N \cup T)^*$, the relation $u \triangleright v$ replaces a nonterminal
 830 A of u by the right-hand side w of a production $A \rightarrow w$ and \triangleright^* denotes the reflexive and
 831 transitive closure of \triangleright . The *language* of G is the set $\mathcal{L}(G)$ of finite strings $w \in T^*$, such
 832 that $S \triangleright^* w$. The *universality problem* asks, given a grammar G , whether $\Sigma^* \subseteq \mathcal{L}(G)$? Let
 833 $G = (\Sigma, N, T, S, \Delta)$ be a context-free grammar and assume w.l.o.g. the following:

- 834 ■ $\Sigma = \{0, 1\}$, because every symbol can be encoded as a binary string,
- 835 ■ G does not produce the empty word or the single letter words 0 and 1; computing a
 836 grammar G' such that $\mathcal{L}(G') = \mathcal{L}(G) \setminus \{\epsilon, 0, 1\}$ is possible and we can reduce from the
 837 modified universality problem $\Sigma^{\geq 2} \stackrel{\text{def}}{=} \{w \in \Sigma^* \mid |w| \geq 2\} \subseteq \mathcal{L}(G')$ instead of the original
 838 problem $\Sigma^* \subseteq \mathcal{L}(G)$.
- 839 ■ G is in Greibach normal form i.e., contains only production rules of the form $Y_0 \rightarrow$
 840 $aY_1 \dots Y_k$, where $Y_0, \dots, Y_k \in N$, for some $k \geq 0$ and $a \in \Sigma$.

841 We define the behavior $\mathbb{B} = (\{p\}, \{q_0, q_1\}, \rightarrow)$, where $q_0 \xrightarrow{p} q_1$ and $q_1 \xrightarrow{p} q_1$. We encode the
 842 language $\Sigma^{\geq 2}$ by the following set of rules, that define the binary predicate symbols $\mathbf{X}_0(x, y)$
 843 and $\mathbf{X}_1(x, y)$ below:

$$\begin{aligned} \mathbf{X}_0(x, y) &\leftarrow \exists z. [x]@q_0 * \langle x.p, z.p \rangle * \mathbf{X}_1(z, y) \\ \mathbf{X}_0(x, y) &\leftarrow \exists z. [x]@q_1 * \langle x.p, z.p \rangle * \mathbf{X}_0(z, y) \\ 844 \quad \mathbf{X}_1(x, y) &\leftarrow [x]@q_1 * x = y \\ \mathbf{X}_1(x, y) &\leftarrow \exists z. [x]@q_1 * \langle x.p, z.p \rangle * \mathbf{X}_1(z, y) \end{aligned}$$

845 Note that $\mathbf{X}_0(x, y)$ defines those configurations encoding words of length at least two, with
 846 exactly one component in state q_0 and the rest of the components in state q_1 . To encode
 847 the language $\mathcal{L}(G)$, we use a binary predicate symbol $\mathbf{Y}_i(x, y)$, where $i \in [1, n]$, for each
 848 nonterminal from the set $N = \{Y_1, \dots, Y_n\}$ and encode each production rule of G of the
 849 form $Y_{i_0} \rightarrow aY_{i_1} \dots Y_{i_k}$, for some $k \geq 0$, by a rule:

$$850 \quad \mathbf{Y}_{i_0}(x, y) \leftarrow \exists z_1 \dots \exists z_k. [x]@q_1 * \langle x.p, z_1.p \rangle * \bigstar_{\ell=1}^{k-1} \mathbf{Y}_{i_\ell}(z_\ell, z_{\ell+1}) * \mathbf{Y}_{i_k}(z_k, y)$$

851 Note that $\mathbf{Y}_i(x, y)$ encodes the words produced by G starting with nonterminal Y_i , by a chain
 852 configurations, in which all components are in a state q_1 . In particular, the parameter y is

not bound to an present component, allowing to compose Y_i with other predicate atoms Y_j .
Finally, we define the predicate A by the following rules:

$$\begin{aligned} A(x, y) &\leftarrow X_0(x, y) \\ A(x, y) &\leftarrow Y_0(x, y) * [y]@q_1 \end{aligned}$$

assuming w.l.o.g. that Y_0 is the starting nonterminal of G . Let Δ be the SID consisting of the rules above. We prove that $\Sigma^{\geq 2} \subseteq \mathcal{L}(G)$ if and only if for any configurations γ, γ' and each store ν , such that $\gamma \models_{\Delta}^{\nu} A(x, y)$ and $\gamma \rightsquigarrow^* \gamma'$, we have $\gamma' \models_{\Delta}^{\nu} A(x, y)$. “ \Rightarrow ” Let γ, γ' and ν be such that $\gamma \models_{\Delta}^{\nu} A(x, y)$ and $\gamma \rightsquigarrow^* \gamma'$. If $\gamma = \gamma'$ we are done, so we assume $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ and $\gamma' = (\mathcal{C}, \mathcal{I}, \varrho')$, where $\varrho \neq \varrho'$. By the definition of \mathbb{B} , the only possibility is that $\varrho' = \varrho[c \leftarrow q_1]$, for some component $c \in \mathcal{C}$, such that $\varrho(c) = q_0$. But then $\gamma \models_{\Delta}^{\nu} X_0(x, y)$. Since γ encodes a word $w \in \Sigma^{\geq 2}$, we have $w \in \mathcal{L}(G)$, by the hypothesis, hence $\gamma' \models_{\Delta}^{\nu} Y_0(x, y) * [y]@q_1$ and $\gamma' \models_{\Delta}^{\nu} A(x, y)$ follows, by the definition of Δ . “ \Leftarrow ” Let $w \in \Sigma^{\geq 2}$ be a word and let $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ be a configuration such that $\varrho(c) = q_0$, for some $c \in \mathcal{C}$ and $\varrho(c') = q_1$, for all $c' \in \mathcal{C} \setminus \{c\}$. By the definition of Δ , we have $\gamma \models_{\Delta}^{\nu} X_0(x, y)$ for a store ν , hence $\gamma \models_{\Delta}^{\nu} A(x, y)$. Let $\gamma' = (\mathcal{C}, \mathcal{I}, \varrho[c \leftarrow q_1])$. By the hypothesis, we have $\gamma' \models_{\Delta}^{\nu} A(x, y)$, hence $\gamma' \models_{\Delta}^{\nu} Y_0(x, y) * [y]@q_1$ is the only possibility. Then we obtain $w \in \mathcal{L}(G)$, again by the definition of Δ . \blacktriangleleft

► **Lemma 31.** *Let Δ be a PCR SID and let $A \in \text{Def}(\Delta)$ be a predicate. Then, for any Δ -model (γ, ν) of $A(x_1, \dots, x_{\#A})$, the configuration γ is tight.*

Proof of Lemma 31. Let $(\mathcal{C}, \mathcal{I}, \varrho)$ be a configuration and ν be a store, such that $(\mathcal{C}, \mathcal{I}, \varrho) \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$. Let $(c_i, p_i)_{i \in [1, n]} \in \mathcal{I}$ be an interaction for which we shall prove that $c_1, \dots, c_n \in \mathcal{C}$. The proof goes by induction on the definition of the satisfaction relation. Assume that $(\mathcal{C}, \mathcal{I}, \varrho) \models_{\Delta}^{\nu'} [x_1] * \psi * \bigstar_{\ell=1}^h B_{\ell}(z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell})$, for a store ν' that agrees with ν over $x_1, \dots, x_{\#A}$, where ψ is a separating conjunction of interaction atoms. Then there exist configurations $\gamma_0, \dots, \gamma_h$, such that $(\mathcal{C}, \mathcal{I}, \varrho) = \gamma_0 \bullet \dots \bullet \gamma_h$, $\gamma_0 \models_{\nu'}^{\nu'} [x_1] * \psi$ and $\gamma_{\ell} \models_{\Delta}^{\nu'} B_{\ell}(z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell})$. We distinguish two cases:

- There exists an interaction atom $\langle y_1.p_1, \dots, y_n.p_n \rangle$ in ψ such that $\nu'(y_i) = c_i$, for all $i \in [1, n]$. Since Δ is progressing, by Def. 29, we have $y_1, \dots, y_n \in \bigcup_{\ell=1}^h \{z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell}\}$. By [4, 12], we obtain $\nu'(y_j) \in \mathcal{C}_{\ell}$, for each $y_j \in \{z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell}\}$, such that $\gamma_{\ell} = (\mathcal{C}_{\ell}, \mathcal{I}_{\ell}, \varrho_{\ell})$. Consequently, we have $\{c_1, \dots, c_n\} \subseteq \bigcup_{\ell=1}^h \mathcal{C}_{\ell} \subseteq \mathcal{C}$, because $(\mathcal{C}, \mathcal{I}, \varrho) = \gamma_0 \bullet \dots \bullet \gamma_h$.
- Else, the induction hypothesis applies to $\gamma_{\ell} \models_{\Delta}^{\nu'} B_{\ell}(z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell})$, for some $\ell \in [1, h]$.

► **Lemma 34.** *The $\text{HavocInv}[\Delta, A]$ problem for PCR SIDs Δ , such that $\text{DegreeBound}[\Delta, A]$ has a positive answer, is 2EXP-hard.*

Proof of Lemma 34. By reduction from the 2EXP-complete entailment problem $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, where Δ is PCR and the set $\{\delta(\gamma) \mid \gamma \models_{\Delta} \exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})\}$ is finite [4, Thm. 4]. Let us fix the above instance of the entailment problem and assume w.l.o.g. that there are no occurrences of state atoms in Δ — the proof of the lower bound from [4, Thm. 4] is actually independent on the occurrences of state atoms in the rules of the SID. Let $\mathbb{B} = (\mathcal{P}, \{q_0, q_1\}, \{q_0 \xrightarrow{p} q_1, q_1 \xrightarrow{p} q_1 \mid p \in \mathcal{P}\})$ be a behavior, where \mathcal{P} is a finite set of ports that subsumes the ports occurring in an interaction atom from Δ . Moreover, assume w.l.o.g. that:

1. A does not occur on the right-hand side of a rule in Δ — each such occurrence of a predicate atom $A(z_1, \dots, z_{\#A})$ can be replaced by $A_0(z_1, \dots, z_{\#A})$, where A_0 is a fresh predicate with the same definition as A ,

2. each rule that defines A has at least one occurrence of a predicate atom — rules of the form $A(x_1, \dots, x_{\#A}) \leftarrow \phi$, where ϕ has no predicate atoms can be removed without affecting the 2EXP-hardness result from [4, Thm. 4] because only a finite subset of the Δ -models of $A(x_1, \dots, x_{\#A})$ (modulo renaming of components) gets lost.

We define the following copies of Δ :

- Δ_0 is the following set of rules:

$$\begin{aligned} A(x_1, \dots, x_{\#A}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q_0 * \psi, \text{ where} \\ A(x_1, \dots, x_{\#A}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1] * \psi \in \Delta \\ &\text{and } \psi \text{ contains no component atoms} \end{aligned}$$

$$\begin{aligned} B'(x_1, \dots, x_{\#B'}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q_1 * \psi, \text{ where} \\ B'(x_1, \dots, x_{\#B'}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1] * \psi \in \Delta, \\ A \neq B' &\text{ and } \psi \text{ contains no component atoms} \end{aligned}$$

- Δ_1 is the following set of rules:

$$\begin{aligned} B'(x_1, \dots, x_{\#B}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q_1 * \psi, \text{ where} \\ B'(x_1, \dots, x_{\#B}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1] * \psi \in \Delta, \\ &\text{and } \psi \text{ contains no component atoms} \end{aligned}$$

Let $\bar{\Delta}$ be the union of Δ_0 and Δ_1 to which the following rules are added, for a fresh predicate \bar{A} of the same arity as A :

$$\begin{aligned} \bar{A}(x_1, \dots, x_{\#A}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q * \psi, \text{ where} \\ A(x_1, \dots, x_{\#A}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q * \psi \in \Delta_0 \\ \bar{A}(x_1, \dots, x_{\#A}) &\leftarrow \exists x_{\#A+1} \dots \exists x_{\#B} \exists y_1 \dots \exists y_m . [x_1]@q * \psi, \text{ where} \\ B(x_1, \dots, x_{\#B}) &\leftarrow \exists y_1 \dots \exists y_m . [x_1]@q * \psi \in \Delta \end{aligned}$$

We prove that $\text{HavocInv}[\bar{\Delta}, \bar{A}]$ has a positive answer if and only if $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$.

“ \Rightarrow ” Let $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ be a configuration and ν be a store such that $\gamma \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$. Let $\gamma_0 = (\mathcal{C}, \mathcal{I}, \varrho_0)$ be the configuration such that $\varrho_0(\nu(x_1)) = q_0$ and $\varrho_0(c) = q_1$, for all $c \in \mathcal{C} \setminus \{\nu(x_1)\}$. By the definition of Δ_0 , we have $\gamma_0 \models_{\Delta_0}^{\nu} A(x_1, \dots, x_{\#A})$, hence also $\gamma_0 \models_{\bar{\Delta}}^{\nu} \bar{A}(x_1, \dots, x_{\#A})$. Let $\gamma_1 \stackrel{\text{def}}{=} (\mathcal{C}, \mathcal{I}, \varrho_1)$, where $\varrho_1 = \varrho_0[\nu(x_1) \leftarrow q_1]$. Since $\bar{\Delta}$ is progressing and connected, by the assumption (2) above, there exists an interaction involving $\nu(x_1)$, then γ_1 is the result of executing that interaction in γ_0 , hence $\gamma_0 \rightsquigarrow^* \gamma_1$. By the hypothesis $\text{HavocInv}[\bar{\Delta}, \bar{A}]$, we obtain $\gamma_1 \models \bar{A}(x_1, \dots, x_{\#A})$ and, since $\varrho_1(c) = q_1$ for all $c \in \mathcal{C}$, it must be the case that $\gamma_1 \models_{\Delta_1}^{\nu} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$. Since no state atoms occur in Δ , we obtain $\gamma \models_{\Delta} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, by the definition of Δ_1 .

“ \Leftarrow ” Let $\gamma = (\mathcal{C}, \mathcal{I}, \varrho)$ and $\gamma' = (\mathcal{C}, \mathcal{I}, \varrho')$ be configurations and ν be a store such that $\gamma \models_{\bar{\Delta}}^{\nu} \bar{A}(x_1, \dots, x_{\#A})$ and $\gamma \rightsquigarrow^* \gamma'$. If $\gamma = \gamma'$ there is nothing to prove. Otherwise, the only possibility is that $\varrho(\nu(x_1)) = q_0$ and $\varrho'(\nu(x_1)) = q_1$. Then we have $\gamma \models_{\Delta_0}^{\nu} A(x_1, \dots, x_{\#A})$, by the definition of $\bar{\Delta}$ and $\gamma \models_{\Delta}^{\nu} A(x_1, \dots, x_{\#A})$ follows from the assumption that there are no state atoms in Δ . By the hypothesis, we obtain $\gamma \models_{\Delta}^{\nu} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$,

940 hence $\gamma' \models_{\Delta_1}^\nu \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, by the definition of Δ_1 , leading to $\gamma' \models_{\bar{\Delta}}$
 941 $\bar{A}(x_1, \dots, x_{\#A})$.

942 We conclude observing that the construction of $\bar{\Delta}$ takes time linear in the size of Δ . ◀

943 ► **Theorem 35.** *The $\text{HavocInv}[\Delta, A]$ problem, for PCR SIDs such that $\text{DegreeBound}[\Delta, A]$
 944 has a positive answer is in 2EXP, if $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$ are bounded by constants and in
 945 5EXP, otherwise.*

946 **Proof of Theorem 35.** If $\#(\Delta)$, $N(\Delta)$ and $H(\Delta)$ are bounded by constants, each instance
 947 $\text{HavocInv}[\Delta, A]$ of the havoc invariance problem can be converted, in polynomial time, into
 948 several instances $\{\bar{A}_i(x_1, \dots, x_{\#A}) \models_{\Delta \cup \bar{\Delta}} A(x_1, \dots, x_{\#A})\}_{i=1}^p$ of the entailment problem,
 949 such that $\bar{A}_i \sim A$, for all $i \in [1, p]$ and $\Delta \simeq \bar{\Delta}$, by Theorem 26. Hence $\bar{\Delta}$ is PCR and
 950 $\text{DegreeBound}[\bar{\Delta}, \bar{A}_i]$ has a positive answer, for each $i \in [1, p]$. By Theorem 33, the entailment
 951 problems can be answered in 2EXP for each of these instances, hence $\text{HavocInv}[\Delta, A]$ can be
 952 answered in 2EXP, as well. Otherwise, if either $\#(\Delta)$, $N(\Delta)$ or $H(\Delta)$ are unbounded, the
 953 entailments are produced in time simply exponential (Theorem 26) and can be answered in
 954 4EXP in their size (Theorem 33), yielding a 5EXP upper bound. ◀