



Decision Problems in a Logic for Reasoning about Reconfigurable Distributed Systems

Marius Bozga, Lucas Bueri, Radu Iosif

► To cite this version:

Marius Bozga, Lucas Bueri, Radu Iosif. Decision Problems in a Logic for Reasoning about Reconfigurable Distributed Systems. International Joint Conference on Automated Reasoning, Aug 2022, Haifa, Israel. hal-03750586

HAL Id: hal-03750586

<https://hal.science/hal-03750586>

Submitted on 12 Aug 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.

Decision Problems in a Logic for Reasoning about Reconfigurable Distributed Systems

Marius Bozga¹, Lucas Bueri¹, and Radu Iosif¹

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

Abstract. We consider a logic used to describe sets of configurations of distributed systems, whose network topologies can be changed at runtime, by reconfiguration programs. The logic uses inductive definitions to describe networks with an unbounded number of components and interactions, written using a multiplicative conjunction, reminiscent of Bunched Implications [38] and Separation Logic [40]. We study the complexity of the satisfiability and entailment problems for the configuration logic under consideration. Additionally, we consider robustness properties, such as tightness (are all interactions entirely connected to components?) and degree boundedness (is every component involved in a bounded number of interactions?), the latter being an ingredient for decidability of entailments.

1 Introduction

Distributed systems are increasingly used as critical parts of the infrastructure of our digital society, as in e.g., datacenters, e-banking and social networking. In order to address maintenance (e.g., replacement of faulty and obsolete network nodes by new ones) and data traffic issues (e.g., managing the traffic inside a datacenter [36]), the distributed systems community has recently put massive effort in designing algorithms for *reconfigurable systems*, whose network topologies change at runtime [25]. However, dynamic reconfiguration is an important source of bugs that may result in e.g., denial of services or even data corruption¹.

This paper contributes to a logical framework that addresses the timely problems of formal *modeling* and *verification* of reconfigurable distributed systems. The basic building blocks of this framework are (i) a Hoare-style program proof calculus [1] used to write formal proofs of correctness of reconfiguration programs, and (ii) an invariant synthesis method [6] that proves the safety (i.e., absence of reachable error configurations) of the configurations defined by the assertions that annotate a reconfiguration program. These methods are combined to prove that an initially correct distributed system cannot reach an error state, following the execution of a given reconfiguration sequence.

The assertions of the proof calculus are written in a logic that defines infinite sets of configurations, consisting of *components* (i.e., processes running on different nodes of the network) connected by *interactions* (i.e., multi-party channels alongside which messages between components are transferred). Systems that share the same architectural style (e.g., pipeline, ring, star, tree, etc.) and differ by the number of components

¹ <https://status.cloud.google.com/incident/appengine/19007>

and interactions are described using inductively defined predicates. Such configurations can be modified either by (a) adding or removing components and interactions (reconfiguration), or (b) changing the local states of components, by firing interactions.

The assertion logic views components and interactions as *resources*, that can be created or deleted, in the spirit of resource logics *à la* Bunched Implications [38], or Separation Logic [40]. The main advantage of using resource logics is their support for *local reasoning* [12]: reconfiguration actions are specified by pre- and postconditions mentioning only the resources involved, while framing out the rest of the configuration.

The price to pay for this expressive power is the difficulty of automating the reasoning in these logics. This paper makes several contributions in the direction of proof automation, by studying the complexity of the *satisfiability* and *entailment* problems, for the configuration logic under consideration. Additionally, we study the complexity of two robustness properties [28], namely *tightness* (are all interactions entirely connected to components?) and *degree boundedness* (is every component involved in a bounded number of interactions?). In particular, the latter problem is used as a prerequisite for defining a fragment with a decidable entailment problem.

1.1 Motivating Example

The logic studied in this paper is motivated by the need for an assertion language that supports reasoning about dynamic reconfigurations in a distributed system. For instance, consider a distributed system consisting of a finite (but unknown) number of *components* (processes) placed in a ring, executing the same finite-state program and communicating via *interactions* that connect the *out* port of a component to the *in* port of its right neighbour, in a round-robin fashion, as in Fig. 1 (a). The behavior of a component is a machine with two states, T and H, denoting whether the component has a token (T) or not (H). A component c_i without a token may receive one, by executing a transition $H \xrightarrow{in} T$, simultaneously with its left neighbour c_j , that executes the transition $T \xrightarrow{out} H$. Then, we say that the interaction (c_j, out, c_i, in) has fired, moving a token one position to the right in the ring. Note that there can be more than one token, moving independently in the system, as long as no token overtakes another token.

The token ring system is formally specified by the following inductive rules:

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

The predicate $\text{ring}_{h,t}(x)$ describes a ring with at least two components, such that at least h (resp. t) components are in state H (resp. T). The ring consists of a component x in state q , described by the formula $[x]@q$, an interaction from the *out* port of x to the *in* port of another component z , described as $\langle x.out, z.in \rangle$, a separate chain of components stretching from z to y ($\text{chain}_{h',t'}(z, y)$), and an interaction connecting the *out* port of component y to the *in* port of component x ($\langle y.out, x.in \rangle$). Inductively, a chain consists of




Fig. 1: Inductive Specification and Reconfiguration of a Token Ring

a component $[x]@q$, an interaction $\langle x.out, z.in \rangle$ and a separate $chain_{h',t'}(z,y)$. Fig. 1 (b) depicts the unfolding of the inductive definition of the token ring, with the existentially quantified variables z from the above rules α -renamed to z^1, z^2, \dots to avoid confusion.

A *reconfiguration program* takes as input a mapping of program variables to components and executes a sequence of *basic operations* i.e., component/interaction creation/deletion, involving the components and interactions denoted by these variables. For instance, the reconfiguration program in Fig. 1 (c) takes as input three adjacent components, mapped to the variables x , y and z , respectively, removes the component y together with its left and right interactions and reconnects x directly with z . Programming reconfigurations is error-prone, because the interleaving between reconfiguration actions and interactions in a distributed system may lead to bugs that are hard to trace. For instance, if a reconfiguration program removes the last component in state T (resp. H) from the system, no token transfer interaction may fire and the system deadlocks.

We prove absence of such errors using a Hoare-style proof system [1], based on the above logic as assertion language. For instance, the proof from Fig. 1 (c) shows that the reconfiguration sequence applied to a component y in state H (i.e., $[y]@H$) in a ring with at least $h \geq 2$ components in state H and at least $t \geq 1$ components in state T leads to a ring with at least $h - 1$ components in state H and at least t in state T; note that the states of the components may change during the execution of the reconfiguration program, as tokens are moved by interactions.

For reasons of proof scalability, a basic operation is specified only with regard to the components and interactions required to avoid faulting. For instance $\{[x]@q\} \text{ delete}(x) \{\text{emp}\}$ (resp. $\{\langle x.out, y.in \rangle\} \text{ disconnect}(x.out, y.in) \{\text{emp}\}$) means that delete (resp. disconnect) requires a component (resp. interaction) and returns an empty configuration, whereas $\{\text{emp}\} \text{ connect}(x.out, y.in) \{\langle x.out, y.in \rangle\}$ means that connect requires nothing and creates an interaction between the given ports of the given components. These local specifications are plugged into a context described by a frame formula F , using the *frame rule* $\{\phi\} P \{\psi\} \Rightarrow \{\phi * [F]\} P \{\psi * F\}$; for readability, the frame formulæ (from the preconditions of the conclusion of the frame rule applications) are enclosed in boxes.

The proof also uses the *consequence rule* $\{\phi\} P \{\psi\} \Rightarrow \{\phi'\} P \{\psi'\}$ that applies if ϕ' is stronger than ϕ and ψ' is weaker than ψ . The side conditions of the consequence rule require checking the validity of the entailments $\text{ring}_{h,t}(y) \models \exists x \exists z . \langle x.out, y.in \rangle * [y]@H * \langle y.out, z.in \rangle * \text{chain}_{h-1,t}(z, x)$ and $\text{chain}_{h-1,t}(z, x) * \langle x.out, z.in \rangle \models \text{ring}_{h-1,t}(z)$, for all $h \geq 2$ and $t \geq 1$. These side conditions can be automatically discharged using the results on the decidability of entailments given in this paper. Additionally, checking the satisfiability of a precondition is used to detect trivially valid Hoare triples.

1.2 Related Work

Formal modeling coordinating architectures of component-based systems has received lots of attention, with the development of architecture description languages (ADL), such as BIP [4] or REO [2]. Many such ADLs have extensions that describe programmed reconfiguration, e.g., [21,31], classified according to the underlying formalism used to define their operational semantics: *process algebras* [13,34], *graph rewriting* [42,45,33], *chemical reactions* [44] (see the surveys [7,11]). Unfortunately, only few ADLs support formal verification, mainly in the flavour of runtime verification [10,17,32,22] or finite-state model checking [14].

Parameterized verification of unbounded networks of distributed processes uses mostly hard-coded coordinating architectures (see [5] for a survey). A first attempt at specifying architectures by logic is the *interaction logic* of Konnov et al. [30], a combination of Presburger arithmetic with monadic uninterpreted function symbols, that can describe cliques, stars and rings. More structured architectures (pipelines and trees) can be described using a second-order extension [35]. However, these interaction logics are undecidable and lack support for automated reasoning.

Specifying parameterized component-based systems by inductive definitions is not new. *Network grammars* [41,33,27] use context-free grammar 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 architectural styles that are, 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 [12].

Local reasoning about concurrent programs has been traditionally the focus of Concurrent Separation Logic (CSL), based on a parallel composition rule [37], initially with a non-interfering (race-free) semantics [8] and later combining ideas of assume-and-rely-guarantee [39,29] with local reasoning [24,43] and abstract notions of framing [16,15,23]. However, the body of work on CSL deals almost entirely with shared-memory multithreading programs, instead of distributed systems, which is the aim of our work. In contrast, we develop a resource logic in which the processes do not just share and own resources, but become mutable resources themselves.

The techniques developed in this paper are inspired by existing techniques for similar problems in the context of Separation Logic (SL) [40]. For instance, we use an abstract domain similar to the one defined by Brotherston et al. [9] for checking satisfiability of symbolic heaps in SL and reduce a fragment of the entailment problem in our logic to SL entailment [20]. In particular, the use of existing automated reasoning techniques for SL has pointed out several differences between the expressiveness of our

logic and that of SL. First, the configuration logic describes hypergraph structures, in which edges are ℓ -tuples for $\ell \geq 2$, instead of directed graphs as in SL, where ℓ is a parameter of the problem: considering ℓ to be a constant strictly decreases the complexity of the problem. Second, the degree (number of hyperedges containing a given vertex) is unbounded, unlike in SL, where the degree of heaps is constant. Therefore, we dedicate an entire section (§5) to the problem of deciding the existence of a bound (and computing a cut-off) on the degree of the models of a formula, used as a prerequisite for the encoding of the entailment problems from the configuration logic as SL entailments.

2 Definitions

We denote by \mathbb{N} the set of positive integers. For a set A , we define $A^1 \stackrel{\text{def}}{=} A$, $A^{i+1} \stackrel{\text{def}}{=} A^i \times A$, for all $i \geq 0$, and $A^+ = \bigcup_{i \geq 1} A^i$, where \times denotes the Cartesian product. We denote by $\text{pow}(A)$ the powerset of A and by $\text{mpow}(A)$ the power-multiset (set of multisets) of A . The cardinality of a finite set A is denoted as $\|A\|$. By writing $A \subseteq_{\text{fin}} B$ 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 tuple $\mathbf{t} = \langle t_1, \dots, t_n \rangle$, we define $|\mathbf{t}| \stackrel{\text{def}}{=} n$, $\langle \mathbf{t} \rangle_i \stackrel{\text{def}}{=} t_i$ and $\langle \mathbf{t} \rangle_{[i, j]} \stackrel{\text{def}}{=} \langle t_i, \dots, t_j \rangle$. By writing $x = \text{poly}(y)$, for given $x, y \in \mathbb{N}$, we mean that there exists a polynomial function $f : \mathbb{N} \rightarrow \mathbb{N}$, such that $x \leq f(y)$.

2.1 Configurations

We model distributed systems as hypergraphs, 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 $\mathbb{B} = (\mathcal{P}, Q, \rightarrow)$, where \mathcal{P} is a finite set of *ports*, Q is a finite set of *states* and $\rightarrow \subseteq Q \times \mathcal{P} \times Q$ is a transition relation. Intuitively, each transition $q \xrightarrow{p} q'$ of the behavior is triggered by a visible event, represented by the port p . For instance, the behavior of the components of the token ring system from Fig. 1 (a) is $\mathbb{B} = (\{in, out\}, \{H, T\}, \{H \xrightarrow{in} T, T \xrightarrow{out} H\})$. The universe \mathbb{C} and the behavior $\mathbb{B} = (\mathcal{P}, Q, \rightarrow)$ are considered fixed in the rest of this paper.

We introduce a logic for describing infinite sets of *configurations* of distributed systems with unboundedly many components and interactions. 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:

Definition 1. A configuration is a tuple $\gamma = (C, I, \rho)$, where:

- $C \subseteq_{\text{fin}} \mathbb{C}$ is a finite set of components, that are present in the configuration,
- $I \subseteq_{\text{fin}} (\mathbb{C} \times \mathcal{P})^+$ is a finite set of interactions, where each interaction is a sequence $(c_1, p_1, \dots, c_n, p_n) \in (\mathbb{C} \times \mathcal{P})^n$ that binds together the ports p_1, \dots, p_n of the pairwise distinct components c_1, \dots, c_n , respectively.
- $\rho : \mathbb{C} \rightarrow Q$ is a state map associating each (possibly absent) component, a state of the behavior \mathbb{B} , such that the set $\{c \in \mathbb{C} \mid \rho(c) = q\}$ is infinite, for each $q \in Q$.

The last condition requires that there is an infinite pool of components in each state $q \in Q$; since \mathbb{C} is infinite and Q is finite, this condition is feasible. For example, the configurations of the token ring from Fig. 1 (a) are $(\{c_1, \dots, c_n\}, \{(c_i, out, c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}, \rho)$, where $\rho : \mathbb{C} \rightarrow \{H, T\}$ is a state map. The ring topology is described by the set of components $\{c_1, \dots, c_n\}$ and interactions $\{(c_i, out, c_{(i \bmod n)+1}, in) \mid i \in [1, n]\}$.

Intuitively, an interaction $(c_1, p_1, \dots, c_n, p_n)$ 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. The interactions are classified according to their sequence of ports, called the *interaction type* and let $\text{Inter} \stackrel{\text{def}}{=} \mathcal{P}^+$ be the set of interaction types; an interaction type models, for instance, the passing of a certain kind of message (e.g., request, acknowledgement, etc.). From an operational point of view, two interactions that differ by a permutation of indices e.g., $(c_1, p_1, \dots, c_n, p_n)$ and $(c_{i_1}, p_{i_1}, \dots, c_{i_n}, p_{i_n})$ such that $\{i_1, \dots, i_n\} = [1, n]$, are equivalent, since the set of transitions is the same; nevertheless, we chose to distinguish them in the following, exclusively for reasons of simplicity.

Note that Def. 1 allows configurations with interactions that involve absent components (i.e., not from the set C of present components in the given configuration). The following definition distinguishes such configurations:

Definition 2. Let $\gamma = (C, I, \rho)$ be a configuration. An interaction $(c_1, p_1, \dots, c_n, p_n)$ is loose in γ if and only if $c_i \notin C$, for some $i \in [1, n]$. If I contains at least one interaction that is loose in γ , we say that γ is loose. An interaction (resp. configuration) that is not loose is said to be tight.

For instance, every configuration of the system from Fig. 1 (a) is tight and becomes loose if a component is deleted. Moreover, the reconfiguration program from Fig. 1 (c) manipulates tight configurations only. In particular, loose configurations are useful for the definition of a composition operation, as the union of disjoint sets of components and interactions:

Definition 3. The composition of two configurations $\gamma_i = (C_i, I_i, \rho)$, for $i = 1, 2$, such that $C_1 \cap C_2 = \emptyset$ and $I_1 \cap I_2 = \emptyset$, is defined as $\gamma_1 \bullet \gamma_2 \stackrel{\text{def}}{=} (C_1 \cup C_2, I_1 \cup I_2, \rho)$. The composition $\gamma_1 \bullet \gamma_2$ is undefined if $C_1 \cap C_2 \neq \emptyset$ or $I_1 \cap I_2 \neq \emptyset$.

Note that a tight configuration may be the result of composing two loose configurations, whereas the composition of tight configurations is always tight. The example below shows that, in most cases, a non-trivial decomposition of a tight configuration necessarily involves loose configurations:

Example 1. Let $\gamma_i = (C_i, I_i, \rho)$ be loose configurations, where $C_i = \{c_i\}$, $I_i = \{(c_i, out, c_{3-i}, in)\}$, for all $i = 1, 2$. Then $\gamma \stackrel{\text{def}}{=} \gamma_1 \bullet \gamma_2$ is the tight configuration $\gamma = (\{c_1, c_2\}, \{(c_1, out, c_2, in), (c_2, out, c_1, in)\}, \rho)$. The only way of decomposing γ into two tight subconfigurations γ'_1 and γ'_2 is taking $\gamma'_1 \stackrel{\text{def}}{=} \gamma$ and $\gamma'_2 \stackrel{\text{def}}{=} (\emptyset, \emptyset, \rho)$, or viceversa. ■

In analogy with graphs, the *degree* of a configuration is the maximum number of interactions from the configuration that involve a (possibly absent) component:

Definition 4. The degree of a configuration $\gamma = (C, I, \rho)$ is defined as $\delta(\gamma) \stackrel{\text{def}}{=} \max_{c \in \mathbb{C}} \delta_c(\gamma)$, where $\delta_c(\gamma) \stackrel{\text{def}}{=} \|\{(c_1, p_1, \dots, c_n, p_n) \in I \mid c = c_i, i \in [1, n]\}\|$.

For instance, the configuration of the system from Fig. 1 (a) has degree two.

2.2 Configuration Logic

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

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

where $x, y, x_1, \dots \in \mathbb{V}$, $q \in Q$ and $A \in \mathbb{A}$. A formula $[x]$, $\langle x_1.p_1, \dots, x_n.p_n \rangle$, $x@q$ and $A(x_1, \dots, x_{\#A})$ is called a *component*, *interaction*, *state* and *predicate* atom, respectively. Sometimes, we use the shorthand $[x]@q \stackrel{\text{def}}{=} [x] * x@q$. Intuitively, the formula $[x]@q * [y]@q' * \langle x.out, y.in \rangle * \langle x.in, y.out \rangle$ describes a configuration consisting of two distinct components, denoted by the values of x and y , in states q and q' , respectively, and two interactions binding the *out* port of one to the *in* port of the other component. For instance, $\gamma = \gamma_1 \bullet \gamma_2$ from Example 1 is such a configuration.

A formula is said to be *pure* if and only if it consists of state atoms, equalities and disequalities. A formula with no occurrences of predicate atoms (resp. existential quantifiers) is called *predicate-free* (resp. *quantifier-free*). A variable is *free* if it does not occur within the scope of an existential quantifier and let $\text{fv}(\phi)$ be the set of free variables of ϕ . A *sentence* is a formula with no free variables. A *substitution* $\phi[x_1/y_1 \dots x_n/y_n]$ replaces simultaneously every free occurrence of x_i by y_i in ϕ , for all $i \in [1, n]$. Before defining the semantics of CL formulæ, we introduce the set of inductive definitions that assigns meaning to predicates:

Definition 5. A set of inductive definitions (SID) Δ consists of rules $A(x_1, \dots, x_{\#A}) \leftarrow \phi$, where $x_1, \dots, x_{\#A}$ are pairwise distinct variables, called *parameters*, such that $\text{fv}(\phi) \subseteq \{x_1, \dots, x_{\#A}\}$. The rule $A(x_1, \dots, x_{\#A}) \leftarrow \phi$ defines A and we denote by $\text{def}_\Delta(A)$ the set of rules from Δ that define A .

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 .

The semantics of CL formulæ is defined by a satisfaction relation $\gamma \models_\Delta^v \phi$ between configurations and formulæ. This relation is parameterized by a *store* $v : \mathbb{V} \rightarrow \mathbb{C}$ mapping the free variables of a formula into components from the universe (possibly absent from γ) and an SID Δ . We write $v[x \leftarrow c]$ for the store that maps x into c and agrees with v on all variables other than x . The definition of the satisfaction relation is by induction on the structure of formulæ, where $\gamma = (C, I, \rho)$ is a configuration (Def. 1):

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

If ϕ is a sentence, the satisfaction relation $\gamma \models_{\Delta}^v \phi$ does not depend on the store, written $\gamma \models_{\Delta} \phi$, in which case we say that γ is a *model* of ϕ . If ϕ is a predicate-free formula, the satisfaction relation does not depend on the SID, written $\gamma \models^v \phi$. A formula ϕ is *satisfiable* if and only if the sentence $\exists x_1 \dots \exists x_n . \phi$ has a model, where $\text{fv}(\phi) = \{x_1, \dots, x_n\}$. A formula ϕ *entails* a formula ψ , written $\phi \models_{\Delta} \psi$ if and only if, for any configuration γ and store v , we have $\gamma \models_{\Delta}^v \phi$ only if $\gamma \models_{\Delta}^v \psi$.

2.3 Separation Logic

Separation Logic (SL) [40] will be used in the following to prove several technical results concerning the decidability and complexity of certain decision problems for CL. For self-containment reasons, we define SL below. The syntax of SL formulæ is described by the following grammar:

$$\phi := \text{emp} \mid x_0 \mapsto (x_1, \dots, x_{\mathfrak{R}}) \mid x = y \mid x \neq y \mid A(x_1, \dots, x_{\#A}) \mid \phi * \phi \mid \exists x . \phi$$

where $x, y, x_0, x_1, \dots \in \mathbb{V}$, $A \in \mathbb{A}$ and $\mathfrak{R} \geq 1$ is an integer constant. Formulæ of SL are interpreted over finite partial functions $h : \mathbb{C} \rightarrow_{\text{fin}} \mathbb{C}^{\mathfrak{R}}$, called *heaps*², by a satisfaction relation $h \Vdash^v \phi$, defined inductively as follows:

$$\begin{aligned} h \Vdash_{\Delta}^v \text{emp} & \iff h = \emptyset \\ h \Vdash_{\Delta}^v x_0 \mapsto (x_1, \dots, x_{\mathfrak{R}}) & \iff \text{dom}(h) = \{v(x_0)\} \text{ and } h(v(x_0)) = \langle v(x_1), \dots, v(x_{\mathfrak{R}}) \rangle \\ h \Vdash_{\Delta}^v \phi_1 * \phi_2 & \iff \text{there exist } h_1, h_2 \text{ such that } \text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset, \\ & h = h_1 \cup h_2 \text{ and } h_i \Vdash_{\Delta}^v \phi_i, \text{ for both } i = 1, 2 \end{aligned}$$

where $\text{dom}(h) \stackrel{\text{def}}{=} \{c \in \mathbb{C} \mid h(c) \text{ is defined}\}$ is the domain of the heap and (dis-)equalities, predicate atoms and existential quantifiers are defined same as for CL.

2.4 Decision Problems

We define the decision problems that are the focus of the upcoming sections. As usual, a decision problem is a class of yes/no queries that differ only in their input. In our case, the input consists of an SID and one or two predicates, written between square brackets.

Definition 6. We consider the following problems, for a SID Δ and predicates $A, B \in \mathbb{A}$:

1. $\text{Sat}[\Delta, A]$: is the sentence $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$ satisfiable for Δ ?
2. $\text{Tight}[\Delta, A]$: is every model γ of the sentence $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$ a tight configuration?
3. $\text{Bnd}[\Delta, A]$: is the set $\{\delta(\gamma) \mid \gamma \models_{\Delta} \exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})\}$ finite?
4. $\text{Entl}[\Delta, A, B]$: does $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#B+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$ hold?

We define the size of a formula ϕ as the total number of occurrences of symbols needed to write it down, denoted by $\text{size}(\phi)$. The size of a SID Δ 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 of a SID Δ are its:

² We use the universe \mathbb{C} here for simplicity, the definition works with any countably infinite set.

- *maximal arity*, denoted as $\text{arity}(\Delta) \stackrel{\text{def}}{=} \max\{\#A \mid A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta\}$,
- *width*, denoted as $\text{width}(\Delta) \stackrel{\text{def}}{=} \max\{\text{size}(\phi) \mid A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta\}$,
- *maximal interaction size*, denoted as $\text{intersize}(\Delta) \stackrel{\text{def}}{=} \max\{n \mid \langle x_1.p_1, \dots, x_n.p_n \rangle \text{ occurs in } \phi, A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta\}$.

For each decision problem $P[\Delta, A, B]$, we consider its (k, ℓ) -bounded versions $P^{(k, \ell)}[\Delta, A, B]$, obtained by restricting the predicates and interaction atoms occurring in Δ to $\text{arity}(\Delta) \leq k$ and $\text{intersize}(\Delta) \leq \ell$, respectively, where k and ℓ are either positive integers or infinity. We consider, for each $P[\Delta, A, B]$, the subproblems $P^{(k, \ell)}[\Delta, A, B]$ corresponding to the three cases (1) $k < \infty$ and $\ell = \infty$, (2) $k = \infty$ and $\ell < \infty$, and (3) $k = \infty$ and $\ell = \infty$. As we explain next, this is because, for the decision problems considered (Def. 6), the complexity for the case $k < \infty, \ell < \infty$ matches the one for the case $k < \infty, \ell = \infty$.

Moreover, for each problem $P[\Delta, A]$ (resp. $P[\Delta, A, B]$), we consider its general version $P[\Delta, \phi]$ (resp. $P[\Delta, \phi, \psi]$), where ϕ and ψ are CL formulæ, whose predicates are interpreted by the rules in Δ . The generalized problems $P[\Delta, \phi]$ involving one predicate atom (points 1 and 3 of Def. 6) can be reduced to their restricted versions $P[\Delta, A]$, by introducing a fresh predicate A_ϕ (not occurring in Δ), of arity $n \geq 0$ and a rule $A_\phi(x_1, \dots, x_n) \leftarrow \phi$, where $\text{fv}(\phi) = \{x_1, \dots, x_n\}$. This reduction is linear in the size of the input and changes none of the following complexity results. Concerning the generalized entailment problem $\text{Entl}[\Delta, \phi, \psi]$, the reduction to the problem $\text{Entl}[\Delta, A, B]$ (Def. 6 4) might affect its decidability status, which is subject to syntactic restrictions on the rules in Δ (details will be given in §6).

Satisfiability (1) and entailment (4) arise naturally during verification of reconfiguration programs. For instance, $\text{Sat}[\Delta, \phi]$ asks whether a specification ϕ of a set configurations (e.g., a pre-, post-condition, or a loop invariant) is empty or not (e.g., an empty precondition typically denotes a vacuous verification condition), whereas $\text{Entl}[\Delta, \phi, \psi]$ is used as a side condition for the Hoare rule of consequence, as in e.g., the proof from Fig. 1 (c). Moreover, entailments must be proved when checking inductiveness of a user-provided loop invariant.

In contrast, the applications of the tightness (2) and boundedness (3) problems are less obvious and require a few explanations. The $\text{Tight}[\Delta, \phi]$ problem is relevant in the context of compositional verification of distributed systems. Suppose we have a distributed system consisting of two interacting subsystems, whose sets of initial configurations are described by ϕ_1 and ϕ_2 , respectively i.e., the initial configurations of the system are described by $\phi_1 * \phi_2$. The compositional verification of a reconfiguration program P reduces checking the validity of a Hoare triple $\{\phi_1 * \phi_2\} P \{\psi_1 * \psi_2\}$ to checking the validity of the simpler $\{\phi_i\} P \{\psi_i\}$, for $i = 1, 2$. Unfortunately, this appealing method faces the problem of *interference* between the subsystems described by ϕ_1 and ϕ_2 , namely the loose interactions of ϕ_i might connect to present components of ϕ_{3-i} and change their states during the execution. In this case, it is sufficient to infer the sets of cross-boundary interactions $\mathcal{F}_{i,3-i}$, describing those interactions from ϕ_i that connect to components from ϕ_{3-i} , and check the validity of the triples $\{\phi_i * \mathcal{F}_{3-i,i}\} P \{\psi_i * \mathcal{F}_{3-i,i}\}$, under a relaxed semantics which considers that the interactions in $\mathcal{F}_{3-i,i}$ can fire anytime, or according to the order described by some regular language. However, if $\text{Tight}[\Delta, \phi_1]$ (resp. $\text{Tight}[\Delta, \phi_2]$) has a negative answer, the set of cross-boundary interactions may be unbounded, hence not representable by a finite sep-

arating conjunction of interaction atoms $\mathcal{F}_{1,2}$ (resp. $\mathcal{F}_{2,1}$). Thus, the tightness problem is important in establishing necessary conditions under which a compositional proof rule can be applied to checking correctness of reconfigurations in a distributed system.

The $\text{Bnd}[\Delta, \phi]$ problem is used to check a necessary condition for the decidability of entailments i.e., $\text{Entl}[\Delta, \phi, \psi]$. If $\text{Bnd}[\Delta, \phi]$ has a positive answer, we can reduce the problem $\text{Entl}[\Delta, \phi, \psi]$ to an entailment problem for SL, which is always interpreted over heaps of bounded degree [20]. Otherwise, the decidability status of the entailment problem is open, for configurations of unbounded degree, such as the one described by the example below.

Example 2. The following SID describes star topologies with a central controller connected to an unbounded number of workers stations:

$$\text{Star}(x) \leftarrow [x] * \text{Worker}(x), \text{Worker}(x) \leftarrow \text{emp} \mid \exists y. \langle x.\text{out}, y.\text{in} \rangle * [y] * \text{Worker}(x) \blacksquare$$

3 Satisfiability

We show that the satisfiability problem (Def. 6, point 1) is decidable, using a method similar to the one pioneered by Brotherston et al. [9], for checking satisfiability of inductively defined symbolic heaps in SL. We recall that a formula π is *pure* if and only if it is a separating conjunction of equalities, disequalities and state atoms.

Definition 7. The closure $\text{cl}(\pi)$ of a pure formula π is the limit of the sequence $\pi^0, \pi^1, \pi^2, \dots$ such that $\pi^0 = \pi$ and, for each $i \geq 0$, π^{i+1} is obtained by joining (with $*$) all of the following formulae to π^i :

- $x = z$, where x and z are the same variable, or $x = y$ and $y = z$ both occur in π^i ,
- $x \neq z$, where $x = y$ and $y \neq z$ both occur in π^i , or
- $y @ q$, where $x @ q$ and $x = y$ both occur in π^i .

Because only finitely many such formulae can be added, the sequence of pure formulae from Def. 7 is bound to stabilize after polynomially many steps. A pure formula is satisfiable if and only if its closure does not contain contradictory literals i.e., $x = y$ and $x \neq y$, or $x @ q$ and $x @ q'$, for $q \neq q' \in Q$. We write $x \approx_\pi y$ (resp. $x \not\approx_\pi y$) if and only if $x = y$ (resp. $x \neq y$) occurs in $\text{cl}(\pi)$ and $\text{not}(x \approx_\pi y)$ (resp. $\text{not}(x \not\approx_\pi y)$) whenever $x \approx_\pi y$ (resp. $x \not\approx_\pi y$) does not hold. Note that e.g., $\text{not}(x \approx_\pi y)$ is not the same as $x \not\approx_\pi y$.

Lemma 1. A pure formula π is satisfiable if and only if the following hold:

1. for all $x, y \in \text{fv}(\pi)$, $x = y$ and $x \neq y$ do not occur both in $\text{cl}(\pi)$,
2. for all $x \in \text{fv}(\pi)$ and $q \neq r \in Q$, $x @ q$ and $x @ r$ do not occur both in $\text{cl}(\pi)$.

Proof. A pure formula π is satisfiable if and only if there exists a store \mathbf{v} and a configuration $(\emptyset, \emptyset, \rho)$, such that $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}} \pi$. “ \Leftarrow ” It is easy to see that \approx_π is an equivalence relation, for each pure formula π . Given any state map ρ , we define \mathbf{v} by assigning each equivalence class of \approx_π a distinct component c , such that $\rho(c) = q$ if $y @ q$ occurs in π , for a variable y in the class. By the conditions of the Lemma, ρ and \mathbf{v} are well defined and we have $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}} \pi$, by definition. “ \Rightarrow ” If $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}} \pi$ then $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}} \text{cl}(\pi)$, because each additional formula in $\text{cl}(\pi)$ is a logical consequence of π . Since $\text{cl}(\pi)$ is satisfiable, the two conditions of the Lemma must hold. \square

Base tuples constitute the abstract domain used by the algorithms for checking satisfiability (point 1 of Def. 6) and boundedness (point 3 of Def. 6), defined as follows:

Definition 8. A base tuple is a triple $t = (C^\sharp, I^\sharp, \pi)$, where:

- $C^\sharp \in \text{mpow}(\mathbb{V})$ is a multiset of variables denoting present components,
- $I^\sharp : \text{Inter} \rightarrow \text{mpow}(\mathbb{V}^+)$ maps each interaction type $\tau \in \text{Inter}$ into a multiset of tuples of variables of length $|\tau|$ each, and
- π is a pure formula.

A base tuple is called *satisfiable* if and only if π is satisfiable and the following hold:

1. for all $x, y \in C^\sharp$, $\text{not}(x \approx_\pi y)$,
2. for all $\tau \in \text{Inter}$, $\langle x_1, \dots, x_{|\tau|} \rangle, \langle y_1, \dots, y_{|\tau|} \rangle \in I^\sharp(\tau)$, there exists $i \in [1, |\tau|]$ such that $\text{not}(x_i \approx_\pi y_i)$,
3. for all $\tau \in \text{Inter}$, $\langle x_1, \dots, x_{|\tau|} \rangle \in I^\sharp(\tau)$ and $1 \leq i < j \leq |\tau|$, we have $\text{not}(x_i \approx_\pi x_j)$.

We denote by SatBase the set of satisfiable base tuples.

Note that a base tuple $(C^\sharp, I^\sharp, \pi)$ is unsatisfiable if C^\sharp (I^\sharp) contains the same variable (tuple of variables) twice (for the same interaction type), hence the use of multisets in the definition of base tuples. It is easy to see that checking the satisfiability of a given base tuple $(C^\sharp, I^\sharp, \pi)$ can be done in time $\text{poly}(\|C^\sharp\| + \sum_{\tau \in \text{Inter}} \|I^\sharp(\tau)\| + \text{size}(\pi))$.

We define a partial *composition* operation on satisfiable base tuples, as follows:

$$(C_1^\sharp, I_1^\sharp, \pi_1) \otimes (C_2^\sharp, I_2^\sharp, \pi_2) \stackrel{\text{def}}{=} (C_1^\sharp \cup C_2^\sharp, I_1^\sharp \cup I_2^\sharp, \pi_1 * \pi_2)$$

where the union of multisets is lifted to functions $\text{Inter} \rightarrow \text{mpow}(\mathbb{V}^+)$ in the usual way. The composition operation \otimes is undefined if $(C_1^\sharp, I_1^\sharp, \pi_1) \otimes (C_2^\sharp, I_2^\sharp, \pi_2)$ is not satisfiable e.g., if $C_1^\sharp \cap C_2^\sharp \neq \emptyset$, $I_1^\sharp(\tau) \cap I_2^\sharp(\tau) \neq \emptyset$, for some $\tau \in \text{Inter}$, or $\pi_1 * \pi_2$ is not satisfiable.

Given a pure formula π and a set of variables X , the projection $\pi \downarrow_X$ removes from π all atomic propositions α , such that $\text{fv}(\alpha) \not\subseteq X$. The *projection* of a base tuple $(C^\sharp, I^\sharp, \pi)$ on a variable set X is formally defined below:

$$(C^\sharp, I^\sharp, \pi) \downarrow_X \stackrel{\text{def}}{=} (C^\sharp \cap X, \lambda \tau. \{ \langle x_1, \dots, x_{|\tau|} \rangle \in I^\sharp(\tau) \mid x_1, \dots, x_{|\tau|} \in X \}, \text{cl}(\text{dist}(I^\sharp) * \pi) \downarrow_X)$$

where $\text{dist}(I^\sharp) \stackrel{\text{def}}{=} *_{\tau \in \text{Inter}} *_{\langle x_1, \dots, x_{|\tau|} \rangle \in I^\sharp(\tau)} *_{1 \leq i < j \leq |\tau|} x_i \neq x_j$

The *substitution* operation $(C^\sharp, I^\sharp, \pi)[x_1/y_1, \dots, x_n/y_n]$ replaces simultaneously each x_i with y_i in C^\sharp , I^\sharp and π , respectively. For a store v , we denote by $v[x_1/y_1, \dots, x_n/y_n]$ the store such that $v[x_1/y_1, \dots, x_n/y_n](x_i) = v(y_i)$ and agrees with v over $\mathbb{V} \setminus \{x_1, \dots, x_n\}$. We lift the composition, projection and substitution operations to sets of satisfiable base tuples, as usual.

Lemma 2. Given a formula ϕ and a substitution $\sigma = [x_1/y_1, \dots, x_n/y_n]$, for any configuration (C, I, ρ) and store v , $(C, I, \rho) \models_\Delta^v \phi \sigma$ only if $(C, I, \rho) \models^{\text{v}\sigma} \phi$.

Proof. By induction on the definition of \models_Δ^v . □

Next, we define the base tuple corresponding to a quantifier- and predicate-free formula $\phi = \Psi * \pi$, where Ψ consists of component and interaction atoms and π is pure. Since, moreover, we are interested in those components and interactions that are visible

input: a SID Δ
output: $\mu\vec{X}.\Delta^\sharp$

- 1: initially $\mu\vec{X}.\Delta^\sharp := \lambda A . \emptyset$
- 2: **for** $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi \in \Delta$, with ϕ quantifier- and predicate-free **do**
- 3: $\mu\vec{X}.\Delta^\sharp(A) := \mu\vec{X}.\Delta^\sharp(A) \cup \text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \downarrow_{x_1, \dots, x_{\#A}}$
- 4: **while** $\mu\vec{X}.\Delta^\sharp$ still change **do**
- 5: **for** $r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$ **do**
- 6: **if** there exist $t_1 \in \mu\vec{X}.\Delta^\sharp(B_1), \dots, t_h \in \mu\vec{X}.\Delta^\sharp(B_h)$ **then**
- 7: $\mu\vec{X}.\Delta^\sharp(A) := \mu\vec{X}.\Delta^\sharp(A) \cup \left(\text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=1}^h t_\ell[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell] \right) \downarrow_{x_1, \dots, x_{\#A}}$

Fig. 2: Algorithm for the Computation of the Least Solution

through a given indexed set of parameters $X = \{x_1, \dots, x_n\}$, for a variable y , we denote by $\{\{y\}\}_\pi^X$ the parameter x_i with the least index, such that $y \approx_\pi x_i$, or y itself, if no such parameter exists. We define the following sets of formulæ:

$$\text{Base}(\phi, X) \stackrel{\text{def}}{=} \begin{cases} \{(C^\sharp, I^\sharp, \pi)\} & , \text{ if } (C^\sharp, I^\sharp, \pi) \text{ is satisfiable} \\ \emptyset & , \text{ otherwise} \end{cases}$$

where $C^\sharp \stackrel{\text{def}}{=} \{\{\{x\}\}_\pi^X \mid [x] \text{ occurs in } \psi\}$
 $I^\sharp \stackrel{\text{def}}{=} \lambda \langle p_1, \dots, p_s \rangle . \{ \langle \{\{y_1\}\}_\pi^X, \dots, \{\{y_s\}\}_\pi^X \rangle \mid \langle y_1.p_1, \dots, y_s.p_s \rangle \text{ occurs in } \psi \}$

We consider a tuple of variables \vec{X} , having a variable $X(A)$ ranging over $\text{pow}(\text{SatBase})$, for each predicate A that occurs in Δ . With these definitions, each rule of Δ :

$$A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(z_1^1, \dots, z_{\#B_1}^1) * \dots * B_h(z_1^h, \dots, z_{\#B_h}^h)$$

where ϕ is a quantifier- and predicate-free formula, induces the constraint:

$$X(A) \supseteq (\text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=1}^h X(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}} \quad (1)$$

Let Δ^\sharp be the set of such constraints, corresponding to the rules in Δ and let $\mu\vec{X}.\Delta^\sharp$ be the tuple of least solutions of the constraint system generated from Δ , indexed by the tuple of predicates that occur in Δ , such that $\mu\vec{X}.\Delta^\sharp(A)$ denotes the entry of $\mu\vec{X}.\Delta^\sharp$ corresponding to A . Since the composition and projection are monotonic operations, such a least solution exists and is unique. Moreover, since SatBase is finite, the least solution can be attained in a finite number of steps, using a standard Kleene iteration (see Fig. 2).

Given a base tuple $(C^\sharp, I^\sharp, \pi)$ and a store v , we define the following sets of components and interactions, respectively:

$$v(C^\sharp) \stackrel{\text{def}}{=} \{v(x) \mid x \in C^\sharp\}$$

$$v(I^\sharp) \stackrel{\text{def}}{=} \bigcup_{\langle p_1, \dots, p_n \rangle \in \text{Inter}} \{(v(x_1), p_1, \dots, v(x_n), p_n) \mid (x_1, \dots, x_n) \in I^\sharp(\langle p_1, \dots, p_n \rangle)\}$$

We state below the main result leading to an elementary recursive algorithm for the satisfiability problem (Thm. 1).

Lemma 3. *Given a base tuple $(C^\sharp, I^\sharp, \pi) \in \mu\vec{X}.\Delta^\sharp(A)[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$, a state map ρ and a store v such that $(\emptyset, \emptyset, \rho) \models^v \pi$, a set of components \mathcal{D} disjoint from $v(C^\sharp)$ and a set of interactions \mathcal{J} disjoint from $v(I^\sharp)$, there exists a configuration (C, I, ρ) , such that $(C, I, \rho) \models_\Delta^v A(y_1, \dots, y_{\#A})$, $C \cap \mathcal{D} = \emptyset$ and $I \cap \mathcal{J} = \emptyset$.*

Proof. Let $\sigma \stackrel{\text{def}}{=} [x_1/y_1, \dots, x_{\#A}/y_{\#A}]$ be a substitution and $(C_0^\sharp, I_0^\sharp, \pi_0) \in \mu\vec{X}.\Delta^\sharp(A)$ be a base pair, such that $(C^\sharp, I^\sharp, \pi) = (C_0^\sharp, I_0^\sharp, \pi_0)\sigma$. Since $(\emptyset, \emptyset, \rho) \models^v \pi$, we obtain $(\emptyset, \emptyset, \rho) \models^{v_0} \pi_0$, by Lemma 2, where we define $v_0 \stackrel{\text{def}}{=} v\sigma$. Let

$$\mathcal{K} \stackrel{\text{def}}{=} \mathcal{D} \cup \{c_i \mid (c_1, p_1, \dots, c_n, p_n) \in \mathcal{J}, i \in [1, n]\}$$

The proof is by fixpoint induction on the definition of $(C_0^\sharp, I_0^\sharp, \pi_0)$. Assume that:

$$(C_0^\sharp, I_0^\sharp, \pi_0) \in (\text{Base}(\Psi * \pi', \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=2}^h \mu\vec{X}.\Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}}$$

for a rule $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m. \Psi * \pi' * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$ of Δ , such that $\Psi * \pi'$ is quantifier-free, Ψ consists of component and interaction atoms and π' is the largest pure subformula of $\Psi * \pi'$. Then there exist base tuples $(C_1^\sharp, I_1^\sharp, \pi_1), \dots, (C_h^\sharp, I_h^\sharp, \pi_h)$, such that:

- $(C_1^\sharp, I_1^\sharp, \pi_1) \in \text{Base}(\Psi * \pi', \{x_1, \dots, x_{\#A}\})$,
- $(C_\ell^\sharp, I_\ell^\sharp, \pi_\ell) \in \mu\vec{X}.\Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]$, for all $\ell \in [2, h]$,
- $(C_0^\sharp, I_0^\sharp, \pi_0) = ((C_1^\sharp, I_1^\sharp, \pi_1) \otimes \dots \otimes (C_h^\sharp, I_h^\sharp, \pi_h)) \downarrow_{x_1, \dots, x_{\#A}}$.

From the first and last points, we deduce $\pi_0 = \text{cl}(\pi' * \bigstar_{\ell=2}^h (\text{dist}(I_\ell^\sharp) * \pi_\ell)) \downarrow_{x_1, \dots, x_{\#A}}$. Let $\pi'' \stackrel{\text{def}}{=} \pi' * \bigstar_{\ell=2}^h (\text{dist}(I_\ell^\sharp) * \pi_\ell)$ and define a store v'_0 , by assigning each $\approx_{\pi''}$ -equivalence class the following component:

- $v_0(x_i)$, if x_i belongs to the class, for some $i \in [1, \#A]$,
- else, if the class is disjoint from $\{x_1, \dots, x_{\#A}\}$ and $y@q$ occurs in π' , for a variable y in the class, we assign $c \in \mathbb{C} \setminus \mathcal{K}$, such that $\rho(c) = q$; since π' is satisfiable, there are no two state atoms $y@q$ and $z@r$, such that $y \approx_{\pi'} z$ and $q \neq r$ in π' and, moreover, choosing c is always possible, by the last point of Def. 1,
- otherwise, the class is assigned an arbitrary component $c \in \mathbb{C} \setminus \mathcal{K}$.

Such a store exists, because π'' is satisfiable and, moreover, $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi''$, hence also $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi_\ell$, for all $\ell \in [2, h]$. We define two sequences of sets of components C_1, \dots, C_h and interactions I_1, \dots, I_h , as follows:

- $C_1 \stackrel{\text{def}}{=} \{v'_0(y) \mid [y] \text{ occurs in } \Psi\}$,
- $I_1 \stackrel{\text{def}}{=} \{(v'_0(z_1), p_1, \dots, v'_0(z_t), p_t) \mid \langle z_1.p_1, \dots, z_t.p_t \rangle \text{ occurs in } \Psi\}$,
- for all $\ell \in [2, h]$, assume that $C_1, \dots, C_{\ell-1}$ and $I_1, \dots, I_{\ell-1}$ have been defined and let us define:

$$\begin{aligned} \mathcal{D}_\ell &\stackrel{\text{def}}{=} \mathcal{D} \cup \bigcup_{j=1}^{\ell-1} C_j \cup \bigcup_{j=\ell+1}^h v'_0(C_j^\sharp) \\ \mathcal{J}_\ell &\stackrel{\text{def}}{=} \mathcal{J} \cup \bigcup_{j=1}^{\ell-1} I_j \cup \bigcup_{j=\ell+1}^h v'_0(I_j^\sharp) \end{aligned}$$

First, we prove that $\mathcal{D}_\ell \cap \mathbf{v}'_0(C_\ell^\#) = \emptyset$ and $\mathcal{I}_\ell \cap \mathbf{v}'_0(I_\ell^\#) = \emptyset$ (we prove the first point, the second using a similar argument). Suppose, for a contradiction, that $c \in \mathcal{D}_\ell \cap \mathbf{v}'_0(C_\ell^\#)$. We distinguish the following cases:

- if $c \in \mathcal{D} \cap \mathbf{v}'_0(C_\ell^\#)$, then $c \in \mathcal{D} \cap \mathbf{v}'_0(C_0^\#)$, because $C_\ell^\# \subseteq C_0^\#$, contradiction with $\mathcal{D} \cap \mathbf{v}'_0(C_0^\#) = \mathcal{D} \cap \mathbf{v}(C^\#) = \emptyset$.
- else, if $c \in \mathcal{C}_j \cap \mathbf{v}'_0(C_\ell^\#)$, for some $j \in [1, \ell - 1]$, then $c \in \mathcal{C}_j \cap \mathcal{D}_j$, because $\mathbf{v}'_0(C_\ell^\#) \subseteq \mathcal{D}_j$, contradiction with the inductive hypothesis $\mathcal{C}_j \cap \mathcal{D}_j = \emptyset$.
- otherwise, $c \in \mathbf{v}'_0(C_j^\#) \cap \mathbf{v}'_0(C_\ell^\#)$, hence there exist variables $y_j \in \mathcal{C}_j^\#$ and $y_\ell \in \mathcal{C}_\ell^\#$, such that $y_j \approx_{\pi'} y_\ell$, contradiction with the fact that $(C_j^\#, I_j^\#, \pi_j) \otimes (C_\ell^\#, I_\ell^\#, \pi_\ell)$ is satisfiable.

Second, we apply the inductive hypothesis to obtain configurations (C_ℓ, I_ℓ, ρ) , such that $(C_\ell, I_\ell, \rho) \models_{\Delta}^{\mathbf{v}'_0} \mathbf{B}_\ell(t_1^\ell, \dots, t_{\#B_\ell}^\ell)$, $C_\ell \cap \mathcal{D}_\ell = \emptyset$ and $I_\ell \cap \mathcal{I}_\ell = \emptyset$, for all $\ell \in [2, h]$. By the definitions of \mathcal{D}_ℓ and \mathcal{I}_ℓ , the sets C_ℓ and I_ℓ are pairwise disjoint, respectively, hence the composition $(C, I, \rho) \stackrel{\text{def}}{=} \bullet_{\ell=1}^h (C_\ell, I_\ell, \rho)$ is defined. Moreover, $(C, I, \rho) \models_{\Delta}^{\mathbf{v}'_0} \psi * \pi' * \star_{\ell=2}^h \mathbf{B}_\ell(t_1^\ell, \dots, t_{\#B_\ell}^\ell)$, hence $(C, I, \rho) \models_{\Delta}^{\mathbf{v}'_0} \mathbf{A}(x_1, \dots, x_{\#A})$, leading to $(C, I, \rho) \models^{\mathbf{v}} \mathbf{A}(y_1, \dots, y_{\#A})$.

Finally, we are left with proving that $C \cap \mathcal{D} = \emptyset$ and $I \cap \mathcal{I} = \emptyset$ (we prove the first point only, the second uses a similar reasoning). Since $C = \bigcup_{\ell=1}^h C_\ell$, this is equivalent to proving the following:

- $C_1 \cap \mathcal{D} = \emptyset$: suppose, for a contradiction, that $\mathbf{v}'_0(y) \in \mathcal{D}$, for a variable y , such that $[y]$ occurs in ψ . By the definition of \mathbf{v}'_0 , we have either $\mathbf{v}'_0(y) \in \mathbf{v}'_0(C_0^\#)$ or $\mathbf{v}'_0(y) \in \mathcal{K}$. Since $\mathbf{v}'_0(C_0^\#) \cap \mathcal{D} = \mathcal{K} \cap \mathcal{D} = \emptyset$, both cases lead to a contradiction.
- $C_\ell \cap \mathcal{D} = \emptyset$, for all $\ell \in [2, h]$: because $C_\ell \cap \mathcal{D}_\ell = \emptyset$ and $\mathcal{D} \subseteq \mathcal{D}_\ell$, by definition of \mathcal{D}_ℓ , for all $\ell \in [2, h]$. \square

Lemma 4. *Given a predicate atom $\mathbf{A}(y_1, \dots, y_{\#A})$, a store \mathbf{v} and a configuration (C, I, ρ) , such that $(C, I, \rho) \models_{\Delta}^{\mathbf{v}} \mathbf{A}(y_1, \dots, y_{\#A})$, there exists a base tuple $(C^\#, I^\#, \pi) \in \mu\vec{X}. \Delta^\#(\mathbf{A})[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$, such that $\mathbf{v}(C^\#) \subseteq C$, $\mathbf{v}(I^\#) \subseteq I$ and $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}} \pi$.*

Proof. By fixpoint induction on the definition of the satisfaction relation $\models_{\Delta}^{\mathbf{v}}$. Since $(C, I, \rho) \models_{\Delta}^{\mathbf{v}} \mathbf{A}(y_1, \dots, y_{\#A})$, by Lemma 2, we have $(C, I, \rho) \models_{\Delta}^{\mathbf{v}_0} \mathbf{A}(x_1, \dots, x_{\#A})$, where $\mathbf{v}_0 \stackrel{\text{def}}{=} \mathbf{v}[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$. Hence, Δ has a rule $\mathbf{A}(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m. \psi * \pi' * \star_{\ell=2}^h \mathbf{B}_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, such that $\psi * \pi'$ is quantifier-free, ψ consists of component and interaction atoms and π' is pure and there exists a store \mathbf{v}'_0 , that agrees with \mathbf{v}_0 over $x_1, \dots, x_{\#A}$ and configurations $(C_1, I_1, \rho), \dots, (C_\ell, I_\ell, \rho)$, such that:

- $(C_1, I_1, \rho) \models_{\Delta}^{\mathbf{v}'_0} \psi * \pi'$,
- $(C_\ell, I_\ell, \rho) \models_{\Delta}^{\mathbf{v}'_0} \mathbf{B}_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for all $\ell \in [2, h]$, and
- $(C, I, \rho) = (C_1, I_1, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$.

We consider the following base tuples:

- $(C_1^\#, I_1^\#, \pi_1) \stackrel{\text{def}}{=} \text{Base}(\psi * \pi', \{x_1, \dots, x_{\#A}\})$,

- $(C_\ell^\sharp, I_\ell^\sharp, \pi_\ell) \in \mu\vec{X}.\Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]$, such that $v'_0(C_\ell^\sharp) \subseteq C_\ell$, $v'_0(I_\ell^\sharp) \subseteq I_\ell$ and $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi_\ell$, whose existence is guaranteed by the inductive hypothesis, for all $\ell \in [2, h]$.

By the definition of $\text{Base}(\psi * \pi', \{x_1, \dots, x_{\#A}\})$ and the fact that $(C_1, I_1, \rho) \models^{v'_0} \psi * \pi'$, we obtain $v'_0(C_1^\sharp) = C_1$ and $v'_0(I_1^\sharp) = I_1$. Since the composition $\bullet_{\ell=1}^h (C_\ell, I_\ell, \rho)$ is defined, the sets C_1, \dots, C_h and I_1, \dots, I_h are pairwise disjoint, respectively. Since $v'_0(C_\ell^\sharp) \subseteq C_\ell$ and $v'_0(I_\ell^\sharp) \subseteq I_\ell$, for all $\ell \in [1, h]$, we deduce that $\otimes_{\ell=1}^h (C_\ell^\sharp, I_\ell^\sharp, \pi_\ell)$ is satisfiable, because:

- for all $1 \leq i < j \leq h$, for any two variables $y \in C_i^\sharp$ and $z \in C_j^\sharp$ we have $\text{not}(y \approx_{\pi'} z)$, because $v'_0(C_i^\sharp) \cap v'_0(C_j^\sharp) = \emptyset$,
- for all $1 \leq i < j \leq h$, all $\tau \in \text{Inter}$, for any two tuples $\langle y_1, \dots, y_{|\tau|} \rangle \in I_i^\sharp(\tau)$ and $\langle z_1, \dots, z_{|\tau|} \rangle \in I_j^\sharp(\tau)$, we have $\text{not}(y_k \approx_{\pi'} z_k)$, for at least some $k \in [1, |\tau|]$, because $v'_0(I_i^\sharp) \cap v'_0(I_j^\sharp) = \emptyset$,
- for each tuple $\langle y_1, \dots, y_{|\tau|} \rangle \in I_\ell^\sharp(\langle p_1, \dots, p_n \rangle)$, for $\ell \in [1, h]$, we have $\text{not}(y_i \approx_{\pi'} y_j)$, for all $1 \leq i < j \leq n$, because $(v'_0(y_1), p_1, \dots, v'_0(y_n), p_n) \in I_\ell$, hence $v'_0(y_1), \dots, v'_0(y_n)$ are pairwise distinct,
- $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi' * \ast_{\ell=2}^h \pi_\ell$, hence $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi' * \ast_{\ell=2}^h \text{dist}(I_\ell^\sharp) * \pi_\ell$, by the previous point.

Then we define $(C_0^\sharp, I_0^\sharp, \pi_0) \stackrel{\text{def}}{=} \left(\otimes_{\ell=1}^h (C_\ell^\sharp, I_\ell^\sharp, \pi_\ell) \right) \downarrow_{x_1, \dots, x_{\#A}}$ and $(C^\sharp, I^\sharp, \pi) \stackrel{\text{def}}{=} (C_0^\sharp, I_0^\sharp, \pi_0)[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$.

By the definition of Δ^\sharp , we have:

$$\mu\vec{X}.\Delta^\sharp(A) \supseteq (\text{Base}(\psi * \pi', \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=2}^h \mu\vec{X}.\Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}}$$

and, since, by the construction of $(C_0^\sharp, I_0^\sharp, \pi_0)$,

$$(C_0^\sharp, I_0^\sharp, \pi_0) \in (\text{Base}(\psi * \pi', \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=2}^h \mu\vec{X}.\Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}}$$

we obtain $(C_0^\sharp, I_0^\sharp, \pi_0) \in \mu\vec{X}.\Delta^\sharp(A)$, leading to $(C^\sharp, I^\sharp, \pi) \in \mu\vec{X}.\Delta^\sharp(A)[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$.

Next, we check that $v(C^\sharp) \subseteq \bigcup_{\ell=1}^h v'_0(C_\ell^\sharp) \subseteq \bigcup_{\ell=1}^h C_\ell = C$ and $v(I^\sharp) \subseteq \bigcup_{\ell=1}^h v'_0(I_\ell^\sharp) \subseteq \bigcup_{\ell=1}^h I_\ell = I$. Finally, the requirement $(\emptyset, \emptyset, \rho) \models^v \pi$ follows from the following:

- $\pi = \pi_0[x_1/y_1, \dots, x_{\#A}/y_{\#A}]$, by the definition of $(C^\sharp, I^\sharp, \pi)$,
- $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi'$ and $(\emptyset, \emptyset, \rho) \models^{v'_0} \pi_\ell$, for all $\ell \in [2, h]$,
- $\pi_0 = \text{cl}(\pi' * \ast_{\ell=2}^h \text{dist}(I_\ell^\sharp) * \pi_\ell) \downarrow_{x_1, \dots, x_{\#A}}$, where $(\emptyset, \emptyset, \rho) \models^{v'_0} \text{dist}(I_\ell^\sharp)$ follows from the satisfiability of $(C_\ell^\sharp, I_\ell^\sharp, \pi_\ell)$, for all $\ell \in [2, h]$. \square

Lemma 5. $\text{Sat}[\Delta, A]$ has a positive answer if and only if $\mu\vec{X}.\Delta^\sharp(A) \neq \emptyset$.

Proof. “ \Leftarrow ” follows from Lemma 3 and “ \Rightarrow ” follows from Lemma 4. \square

If the maximal arity of the predicates occurring in Δ is bound by a constant k , no satisfiable base tuple $(C^\sharp, I^\sharp, \pi)$ can have a tuple $\langle y_1, \dots, y_{|\tau|} \rangle \in I^\sharp(\tau)$, for some $\tau \in \text{Inter}$, such that $|\tau| > k$, since all variables $y_1, \dots, y_{|\tau|}$ are parameters denoting distinct components (point 3 of Def. 8). Hence, the upper bound on the size of a satisfiable base tuple is constant, in both the $k < \infty, \ell < \infty$ and $k < \infty, \ell = \infty$ cases, which are, moreover indistinguishable complexity-wise (i.e., both are NP-complete). In contrast, in the cases $k = \infty, \ell < \infty$ and $k = \infty, \ell = \infty$, the upper bound on the size of satisfiable base tuples is polynomial and simply exponential in $\text{size}(\Delta)$, incurring a complexity gap of one and two exponentials, respectively. The theorem below states the main result of this section:

Theorem 1. $\text{Sat}^{(k, \infty)}[\Delta, A]$ is NP-complete for $k \geq 4$, $\text{Sat}^{(\infty, \ell)}[\Delta, A]$ is EXP-complete and $\text{Sat}[\Delta, A]$ is in 2EXP.

Proof. Membership (upper bounds). For non-negative integers $m \leq n$ denote by $S_{n,m} \stackrel{\text{def}}{=} \frac{n!}{(n-m)!}$ the number of ordered m -element subsets of a n -element set.

Let $\alpha = \text{arity}(\Delta)$, $\beta = \text{intersize}(\Delta)$ and $p = \|\mathcal{P}\|$. The maximum length of a satisfiable base tuple is $B \stackrel{\text{def}}{=} \alpha + (\sum_{j=1}^{\min(\alpha, \beta)} p^j \cdot S_{\alpha, j}) + (2\alpha^2 + \alpha)$, that is, size of the set of components C^\sharp plus the size of the set of interactions I^\sharp plus the length of the longest pure formula π . In general, for any non-negative integer j there exists at most p^j interaction types of arity j with ports from \mathcal{P} ; moreover, for any such interaction type there exists at most $S_{\alpha, j}$ interactions relating distinct components from an α -element set. Moreover, no such interaction exists neither if $j > \alpha$ nor $j > \beta$.

For any $u \leq \alpha$ it holds that $\sum_{j=1}^u p^j \cdot S_{\alpha, j} \leq p^u \alpha^u$ (an easy check by induction on u). We use the inequality above with $u = \min(\alpha, \beta)$ and obtain that $B \leq 2\alpha + 2\alpha^2 + p^{\min(\alpha, \beta)} \alpha^{\min(\alpha, \beta)} \stackrel{\text{def}}{=} B^*$. We distinguish the three cases:

1. $k < \infty, \ell = \infty$: since $\alpha \leq k$ then α is constant and $B^* = O(1)$,
2. $k = \infty, \ell < \infty$: since $\beta \leq \ell$ and $\alpha = O(\text{size}(\Delta))$ then $B^* = \text{poly}(\text{size}(\Delta))$,
3. $k = \infty, \ell = \infty$: since $\alpha = O(\text{size}(\Delta))$ then $B^* = 2^{\text{poly}(\text{size}(\Delta))}$.

Let $N \stackrel{\text{def}}{=} 2^{B^*}$, that is, (an over-approximation of) the total number of base tuples. Clearly, N is constant in case (1) and respectively $2^{\text{poly}(\text{size}(\Delta))}$ and $2^{2^{\text{poly}(\text{size}(\Delta))}}$ in cases (2), (3). Let L be the number of predicates occurring in Δ and H be the maximum number of predicates used in a term in Δ . Let observe that both L and H are in general $O(\text{size}(\Delta))$. Then the least solution $\mu \vec{X}. \Delta^\sharp$ has at most N base tuples for each predicate, hence at most $L \cdot N$ base tuples. Furthermore, for each rule of Δ the time to check and/or produce the base tuple $(C_0^\sharp, I_0^\sharp, \pi_0)$ with respect to the rule constraint (1) and given arguments $(C_j^\sharp, I_j^\sharp, \pi_j)_{j=1, h}$ is polynomial $\text{poly}(B^*, \text{size}(\Delta))$. That is, both composition and projection take at most $(H+1)B^* + \text{size}(\Delta)^3$ time as they need to process (union or scan) at most $H+1$ base tuples of length B^* each plus the closure of pure formula with at most $\text{size}(\Delta)$ variables.

1. $k < \infty, \ell = \infty$: We define a non-deterministic algorithm as follows. Let (Δ, A) be the input instance. We guess a witness $\langle W_1, \dots, W_K \rangle$ for a least solution, where $1 \leq K \leq L \cdot N$ and each W_i entry is of the form $(T_i, r_i, e_{i,1}, \dots, e_{i,h_i})$, where T_i is a base tuple, r_i an index of a rule of Δ and $e_{i,1}, \dots, e_{i,h_i}$ are index values from $\{i+1, \dots, K\}$ for $0 \leq h_i \leq H$. The length of every witness entry is therefore at

most $B^* + \lceil \log_2(\text{size}(\Delta)) \rceil + H \lceil \log_2(L \cdot N) \rceil$. As N is constant when $k < \infty$, and L and H are $O(\text{size}(\Delta))$ it follows that the number of guesses is polynomial for building $\langle W_1, \dots, W_K \rangle$. We now check that $\langle W_1, \dots, W_K \rangle$ represents indeed a valid computation of a base tuples from the least solution i.e., $(C^\#, I^\#, \pi) \in \mu \vec{X}. \Delta^\#(A)$. For this, we need to check: (a) every entry is well-formed, that is, the rule indexed by r_i instantiates precisely h_i predicates; moreover, for every $1 \leq j \leq h_i$ the index $e_{i,j}$ designates an entry $W_{e_{i,j}}$ whose rule defines the j -th predicate instantiated by the rule r_i ; (b) the base tuple of every entry is satisfiable and correctly computed, that is, T_i is the result of applying the constraint (1) for rule r_i with actual arguments $T_{e_{i,1}}, \dots, T_{e_{i,h_i}}$ from the referred entries; (c) the rule r_1 of the first entry W_1 defines the predicate A . Again, as B^* and N are constant in this case, all these checks are done in polynomial time. Since both the generation and the checking of the witness are polynomial time, this ensures membership in NP.

2. $k = \infty, \ell < \infty$: Consider the computation of the least solution $\mu \vec{X}. \Delta^\#$ using standard Kleene iteration. At every step, a rule of Δ and a tuple of at most H base tuples arguments are selected to produce a new base tuple. Thus, in the worst case, at most $\text{size}(\Delta)$ rules in combination with at most N^H base tuples need to be selected and evaluated. If no new base tuple is generated the fixpoint is reached and the algorithm stops. Since there are at most $L \cdot N$ base tuples in the least solution, the total time will be therefore $L \cdot N \cdot \text{size}(\Delta) \cdot N^H \cdot t(B^*, \text{size}(\Delta))$ where $t(B^*, \text{size}(\Delta))$ is the (polynomial) time to process one selection. It is an easy check that the above is $2^{\text{poly}(\text{size}(\Delta))}$ since $N = 2^{\text{poly}(\text{size}(\Delta))}$ in this case.
3. $k = \infty, \ell = \infty$: Following the same reasoning as in the previous case the complexity is $2^{2^{\text{poly}(\text{size}(\Delta))}}$ as $N = 2^{2^{\text{poly}(\text{size}(\Delta))}}$ in this case.

Hardness (lower bounds). The restricted fragment of CL to $*$, $=$, \neq is equisatisfiable to the restricted fragment of SL restricted to $*$, $=$, \neq . The satisfiability of the above SL fragment has been proven respectively NP-hard, if the arities of predicates are bounded by a constant $k \geq 3$ [9, Theorem 4.9] and EXP-hard, in general [9, Theorem 4.15]. Yet, the reductions considered in these proofs rely on the use of a predefined *nil* constant symbol in the SL logic; this constant can be nevertheless replaced by a variable consistently propagated along the SID, that is, at the price of increasing the arities of all predicates by one. Therefore, it follows immediately that $\text{Sat}^{(k, \infty)}[\Delta, \phi]$ is NP-hard for $k \geq 4$ and $\text{Sat}^{(\infty, \ell)}[\Delta, \phi]$, $\text{Sat}[\Delta, \phi]$ are both EXP-hard. \square

Example 3. The doubly-exponential upper bound for the algorithm computing the least solution of a system of constraints of the form (1) is necessary, in general, as illustrated by the following worst-case example. Let n be a fixed parameter and consider the n -arity predicates A_1, \dots, A_n defined by the following SID:

$$\begin{aligned} A_i(x_1, \dots, x_n) &\rightarrow *_{j=0}^{n-i} A_{i+1}(x_1, \dots, x_{i-1}, [x_i, \dots, x_n]^j), \text{ for all } i \in [1, n-1] \\ A_n(x_1, \dots, x_n) &\rightarrow \langle x_1.p, \dots, x_n.p \rangle \\ A_n(x_1, \dots, x_n) &\rightarrow \text{emp} \end{aligned}$$

where, for a list of variables x_i, \dots, x_n and an integer $j \geq 0$, we write $[x_i, \dots, x_n]^j$ for the list rotated to the left j times (e.g., $[x_1, x_2, x_3, x_4, x_5]^2 = x_3, x_4, x_5, x_1, x_2$). In this example, when starting with $A_1(x_1, \dots, x_n)$ one eventually obtains predicate atoms $A_n(x_{i_1}, \dots, x_{i_n})$,

for any permutation x_{i_1}, \dots, x_{i_n} of x_1, \dots, x_n . Since A_n may choose to create or not an interaction with that permutation of variables, the total number of base tuples generated for A_1 is $2^{n!}$. That is, the fixpoint iteration generates $2^{2^{O(n \log n)}}$ base tuples, whereas the size of the input of $\text{Sat}[\Delta, A]$ is $\text{poly}(n)$. ■

4 Tightness

The tightness problem (Def. 6, point 2) is the complement of a problem slightly stronger than satisfiability (1): given a SID Δ and a formula ϕ , such that $\text{fv}(\phi) = \{x_1, \dots, x_n\}$, the *looseness problem* $\text{Loose}[\Delta, A]$ asks for the existence of a loose configuration γ (Def. 2), such that $\gamma \models_{\Delta} \exists x_1 \dots \exists x_n . \phi$. We establish upper and lower bounds for the complexity of the looseness problem by a reduction to and from the satisfiability problem. The bounds for the tightness problem follow by standard complementation of the complexity classes for the looseness problem.

From Looseness to Satisfiability. Let Δ be a given SID and A be a predicate. For each predicate B that occurs in Δ , we consider a fresh predicate B' , not occurring in Δ , such that $\#B' = \#B + 1$. The SID $\tilde{\Delta}$ consists of Δ and, for each rule of Δ of the form:

$$B_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

where ϕ is a quantifier- and predicate-free formula, $\tilde{\Delta}$ has the following rules:

$$B'_0(x_1, \dots, x_{\#B_0+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * x_{\#B_0+1} = z * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

if z occurs in an interaction atom from ϕ , and:

$$B'_0(x_1, \dots, x_{\#B_0+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B'_i(t_1^i, \dots, t_{\#B_i}^i, x_{\#B_0+1}) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

for some $i \in [1, h]$, if $h \geq 1$. Moreover, for each rule of Δ of the form above, with no predicate atoms (i.e., $h = 0$), $\tilde{\Delta}$ contains the rule:

$$B'(x_1, \dots, x_{\#B+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * [x_{\#B+1}]$$

if and only if ϕ contains no predicate atoms. Finally, there is a fresh predicate \tilde{A} , of arity $\#A$, with a rule:

$$\tilde{A}(x_1, \dots, x_{\#A}) \leftarrow \exists y . A'(x_1, \dots, x_{\#A}, y) * [y]$$

Intuitively, the last parameter of a B' predicate binds to an arbitrary variable of an interaction atom. A configuration γ is loose if and only if the value (component) of some variable occurring in an interaction atom is absent, in which case the component can be added to γ (without clashing with a present component of γ) by the last rule. The reduction is polynomial, since the number of rules in $\tilde{\Delta}$ is linear in the number of rules in Δ and the size of each newly added rule is increased by a constant. The following lemma states the correctness of the reduction:

Lemma 6. *Given a SID Δ and a predicate A , the problem $\text{Loose}[\Delta, A]$ has a positive answer if and only if the problem $\text{Sat}[\tilde{\Delta}, \tilde{A}]$ has a positive answer.*

Proof. “ \Rightarrow ” Let $\gamma \stackrel{\text{def}}{=} (C, I, \rho)$ be a loose configuration, such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$, for some store v . Since γ is loose, there exists an interaction $(c_1, p_1, \dots, c_n, p_n) \in I$, such that $c_i \notin C$, for some $i \in [1, n]$. We prove that $\gamma \models_{\tilde{\Delta}}^{v[y \leftarrow c_i]} A'(x_1, \dots, x_{\#A}, y)$, by fixpoint induction on the definition of $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. This is sufficient, because then we obtain $(C \cup \{c_i\}, I, \rho) \models_{\Delta}^v \tilde{A}(x_1, \dots, x_{\#A})$, thus $\text{Sat}[\tilde{\Delta}, \tilde{A}]$ has a positive answer. Consider the rule of Δ :

$$A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

where ϕ is quantifier- and predicate-free and $c'_1, \dots, c'_m \in \mathbb{C}$ are components, such that $(C, I, \rho) \models_{\Delta}^{v'} \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$, where $v' \stackrel{\text{def}}{=} v[y_1 \leftarrow c'_1, \dots, y_m \leftarrow c'_m]$. We distinguish the following cases:

- if $(c_1, p_1, \dots, c_n, p_n) \in I$ because of an interaction atom $\langle z_1.p_1, \dots, z_n.p_n \rangle$ from ϕ , such that $v'(z_i) = c_i$, for all $i \in [1, n]$, then $(C, I, \rho) \models_{\Delta}^{v[y \leftarrow c_i]} \phi * y = z_i * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$, hence $(C, I, \rho) \models_{\tilde{\Delta}}^{v[y \leftarrow c_i]} A'(x_1, \dots, x_{\#A}, y)$, by the definition of $\tilde{\Delta}$.
- else $(c_1, p_1, \dots, c_n, p_n) \in I$ because of a configuration γ' , such that $\gamma = \gamma' \bullet \gamma''$, for some configuration γ'' , and $\gamma' \models_{\Delta}^{v'} B_i(t_1^i, \dots, t_{\#B_i}^i)$. By the inductive hypothesis, we obtain $\gamma' \models_{\Delta}^{v[y \leftarrow c_i]} B'_i(t_1^i, \dots, t_{\#B_i}^i, y)$, hence $\gamma \models_{\tilde{\Delta}}^{v[y \leftarrow c_i]} A'(x_1, \dots, x_{\#A}, y)$, because $\tilde{\Delta}$ contains the rule:

$$A'(x_1, \dots, x_{\#A+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B'_i(t_1^i, \dots, t_{\#B_i}^i, x_{\#A+1}) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

$$\text{and } \gamma'' \models_{\Delta}^{v'} \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B'_{i-1}(t_1^{i-1}, \dots, t_{\#B_{i-1}}^{i-1}) * B'_{i+1}(t_1^{i+1}, \dots, t_{\#B_{i+1}}^{i+1}) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h) \text{ follows from } \gamma = \gamma' \bullet \gamma''.$$

“ \Leftarrow ” Let $\gamma \stackrel{\text{def}}{=} (C, I, \rho)$ be a configuration and v be a store, such that $\gamma \models_{\tilde{\Delta}}^v \tilde{A}(x_1, \dots, x_{\#A})$. Since the only rule of $\tilde{\Delta}$ that defines \tilde{A} is:

$$\tilde{A}(x_1, \dots, x_{\#A}) \leftarrow \exists y . A'(x_1, \dots, x_{\#A}, y) * [y]$$

there exists a component $c \in C$, such that $(C \setminus \{c\}, I, \rho) \models_{\tilde{\Delta}}^{v[y \leftarrow c]} A'(x_1, \dots, x_{\#A}, y)$. We prove the following:

- there exists an interaction $(c_1, p_1, \dots, c_n, p_n) \in I$, such that $c_i = c$, and
- $(C \setminus \{c\}, I, \rho) \models_{\Delta}^v A(x_1, \dots, x_{\#A})$,

by fixpoint induction on the definition of $(C \setminus \{c\}, I, \rho) \models_{\tilde{\Delta}}^{v[y \leftarrow c]} A'(x_1, \dots, x_{\#A}, y)$. Based on the definition of $\tilde{\Delta}$, we distinguish the following cases, where ϕ is quantifier- and predicate-free, $c'_1, \dots, c'_m \in \mathbb{C}$ are components and $v' \stackrel{\text{def}}{=} v[y_1 \leftarrow c'_1, \dots, y_m \leftarrow c'_m]$:

- $(C \setminus \{c\}, I, \rho) \models_{\Delta}^{v[y \leftarrow c]} \phi * y = z * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$, where z occurs in an interaction atom from ϕ . In this case, there exists an interaction $(c_1, p_1, \dots, c_n, p_n) \in I$, such that $c = c_i$, for some $i \in [1, n]$. Moreover, $(C \setminus \{c\}, I, \rho) \models_{\Delta}^v A(x_1, \dots, x_{\#A})$, because Δ has a rule:

$$A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

such that $(C \setminus \{c\}, I, \rho) \models_{\Delta}^v \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$.

- $(C \setminus \{c\}, I, \rho) \models_{\Delta}^{v[y \leftarrow c]} \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B'_i(t_1^i, \dots, t_{\#B_i}^i, y) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$. In this case, there exists configurations γ' and γ'' , such that $(C \setminus \{c\}, I, \rho) = \gamma' \bullet \gamma''$, $\gamma' \models_{\Delta}^{v[y \leftarrow c]} B'_i(t_1^i, \dots, t_{\#B_i}^i, y)$ and $\gamma'' \models_{\Delta}^v \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$. By the inductive hypothesis, there exists an interaction $(c_1, p_1, \dots, c_n, p_n)$ in γ' , such that $c = c_i$, for some $i \in [1, n]$ and $\gamma' \models_{\Delta}^v B_i(t_1^i, \dots, t_{\#B_i}^i)$. Then $(c_1, p_1, \dots, c_n, p_n) \in I$ and $(C \setminus \{c\}, I, \rho) \models_{\Delta}^v A(x_1, \dots, x_{\#A})$, since Δ has a rule:

$$A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$$

such that $\gamma' \bullet \gamma'' \models_{\Delta}^v \phi * B_1(t_1^1, \dots, t_{\#B_1}^1) * \dots * B_h(t_1^h, \dots, t_{\#B_h}^h)$.

- $(C \setminus \{c\}, I, \rho) \models_{\Delta}^{v[y \leftarrow c]} \phi * [y]$ this case contradicts the semantics of CL. \square

From Satisfiability to Looseness. Given a SID Δ and a predicate A , we build a SID $\tilde{\Delta}$ that defines a predicate \tilde{A} , of equal arity, not occurring in Δ , such that $\text{Sat}[\Delta, A]$ has a positive answer if and only if there exists a loose configuration γ and a store v , such that $\gamma \models_{\Delta}^v \tilde{A}(x_1, \dots, x_{\#A})$. The rules of $\tilde{\Delta}$ are the rules of Δ , to which the following rule is added, for some ports $p_1, p_2 \in \mathcal{P}$:

$$\tilde{A}(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \exists y_2 . A(x_1, \dots, x_{\#A}) * \langle y_1.p_1, y_2.p_2 \rangle$$

This reduction is polynomial, because we add one rule, of size linear on $\#A$. The following lemma states the correctness of the reduction:

Lemma 7. *Given a SID Δ and a predicate A , the problem $\text{Sat}[\Delta, A]$ has a positive answer if and only if the problem $\text{Loose}[\tilde{\Delta}, \tilde{A}]$ has a positive answer.*

Proof. “ \Rightarrow ” If $\text{Sat}[\Delta, A]$ has a positive answer, there exists a configuration $\gamma \stackrel{\text{def}}{=} (C, I, \rho)$ and a store v , such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. Consider the configuration $\gamma' \stackrel{\text{def}}{=} (\emptyset, \{(c_1, p_1, c_2, p_2)\}, \rho)$, for some components $c_1, c_2 \notin C$. Then the composition $\gamma \bullet \gamma'$ is defined and we have $\gamma \bullet \gamma' \models_{\Delta}^{v[y_1 \leftarrow c_1, y_2 \leftarrow c_2]} A(x_1, \dots, x_{\#A}) * \langle y_1.p_1, y_2.p_2 \rangle$, leading to $\gamma \bullet \gamma' \models_{\Delta}^v \tilde{A}(x_1, \dots, x_{\#A})$. Moreover, $\gamma \bullet \gamma'$ is loose, because $c_1, c_2 \notin C$. “ \Leftarrow ” If $\gamma \models_{\Delta}^v \tilde{A}(x_1, \dots, x_{\#A})$, we necessarily have $\gamma \models_{\Delta}^{v[y_1 \leftarrow c_1, y_2 \leftarrow c_2]} A(x_1, \dots, x_{\#A}) * \langle y_1.p_1, y_2.p_2 \rangle$, for some components $c_1, c_2 \in \mathbb{C}$, hence there exists a configuration γ' , such that $\gamma' \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. \square

The polynomial reductions from Lemmas 6 and 7 establish the following complexity bounds for the tightness problem:

Theorem 2. $\text{Tight}^{(k,\infty)}[\Delta, A]$ is co-NP-complete, $\text{Tight}^{(\infty,\ell)}[\Delta, A]$ is EXP-complete and $\text{Tight}[\Delta, A]$ is 2EXP.

Proof. Since $\text{Loose}^{(k,\infty)}[\Delta, A]$ is polynomially-reducible to $\text{Sat}^{(k+1,\infty)}[\Delta, A]$, by Theorem 1, we obtain that $\text{Loose}^{(k,\infty)}[\Delta, A]$ is in NP. Moreover, since $\text{Sat}^{(k,\infty)}[\Delta, A]$ is polynomially-reducible to $\text{Loose}^{(k,\infty)}[\Delta, A]$, by Theorem 1, we obtain that $\text{Loose}^{(k,\infty)}[\Delta, A]$ is NP-complete. Because $\text{Tight}^{(k,\infty)}[\Delta, A]$ is the complement of $\text{Loose}^{(k,\infty)}[\Delta, A]$, we obtain that $\text{Tight}^{(k,\infty)}[\Delta, A]$ is co-NP-complete. The rest of the bounds are obtained by the same polynomial reductions and the fact that $\text{Tight}^{(k,\infty)}[\Delta, A]$ is the complement of $\text{Loose}^{(k,\infty)}[\Delta, A]$, for any k and ℓ , either integer constants, or infinity. \square

5 Degree Boundedness

The boundedness problem (Def. 6, point 3) asks for the existence of a bound on the degree (Def. 4) of the models of a sentence $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$. For instance, it is possible to define inductively star topologies, with a central controller connected to an unbounded number of workers stations. Roughly speaking, the $\text{Bnd}[\Delta, A]$ problem has a negative answer if and only if there are increasingly large unfoldings (expansions of a formula by replacement of a predicate atom with one of its definitions) of $A(x_1, \dots, x_{\#A})$ repeating a rule that contains an interaction atom involving a parameter of the rule, which is always bound to the same component. For instance, the rule $\text{Worker}(x) \leftarrow \exists y . \langle x.out, y.in \rangle * [y] * \text{Worker}(x)$ (Example 2) declares an unbounded number of interactions $\langle x.out, y.in \rangle$ involving the component to which x is bound.

Definition 9. Given a predicate A and a sequence $(r_1, i_1), \dots, (r_n, i_n) \in (\Delta \times \mathbb{N})^+$, where r_1 is the rule $A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta$, the unfolding $A(x_1, \dots, x_{\#A}) \xrightarrow{(r_1, i_1) \dots (r_n, i_n)}_{\Delta} \psi$ is inductively defined as (1) $\psi = \phi$ if $n = 1$, and (2) ψ is obtained from ϕ by replacing its i_1 -th predicate atom $B(y_1, \dots, y_{\#B})$ with $\psi_1[x_1/y_1, \dots, x_{\#B}/y_{\#B}]$, where $B(x_1, \dots, x_{\#B}) \xrightarrow{(r_2, i_2) \dots (r_n, i_n)}_{\Delta} \psi_1$ is an unfolding, if $n > 1$.

We show that the $\text{Bnd}[\Delta, A]$ problem can be reduced to the existence of increasingly large unfoldings or, equivalently, a cycle in a finite directed graph, built by a variant of the least fixpoint iteration algorithm used to solve the satisfiability problem (Fig. 3).

Definition 10. Given satisfiable base pairs $t, u \in \text{SatBase}$ and a rule from Δ :

$$r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(z_1^1, \dots, z_{\#B_1}^1) * \dots * B_h(z_1^h, \dots, z_{\#B_h}^h)$$

where ϕ is a quantifier- and predicate-free formula, we write $(A, t) \rightsquigarrow^{(r,i)} (B, u)$ if and only if $B = B_i$ and there exist satisfiable base tuples $t_1, \dots, u = t_i, \dots, t_h \in \text{SatBase}$, such that $t \in (\text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=1}^h t_\ell[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}}$. We define the directed graph with edges labeled by pairs $(r, i) \in \Delta \times \mathbb{N}$:

$$\mathcal{G}(\Delta) \stackrel{\text{def}}{=} (\{\text{def}(\Delta) \times \text{SatBase}\}, \{ \langle (A, t), (r, i), (B, u) \rangle \mid (A, t) \rightsquigarrow^{(r,i)} (B, u) \})$$

input: a SID Δ

output: $\mathcal{G}(\Delta) = (V, E)$

- 1: initially $V := \emptyset, E := \emptyset$
- 2: **for** $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi \in \Delta$, with ϕ quantifier- and predicate-free **do**
- 3: $V := V \cup (\{A\} \times \text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \downarrow_{x_1, \dots, x_{\#A}})$
- 4: **while** V or E still change **do**
- 5: **for** $r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$ **do**
- 6: **if** there exist $(B_1, t_1), \dots, (B_h, t_h) \in V$ **then**
- 7: $X := (\text{Base}(\phi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=1}^h t_\ell[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]) \downarrow_{x_1, \dots, x_{\#A}}$
- 8: $V := V \cup (\{A\} \times X)$
- 9: $E := E \cup \{ \langle (A, t), (r, \ell), (B_\ell, t_\ell) \rangle \mid t \in X, \ell \in [1, h] \}$

Fig. 3: Algorithm for the Construction of $\mathcal{G}(\Delta)$

The graph $\mathcal{G}(\Delta)$ is built by the algorithm in Fig. 3, a slight variation of the classical Kleene iteration algorithm for the computation of the least solution of the constraints of the form (1) (see Fig. 2). A path $(A_1, t_1) \xrightarrow{(r_1, i_1)} (A_2, t_2) \xrightarrow{(r_2, i_2)} \dots \xrightarrow{(r_n, i_n)} (A_n, t_n)$ in $\mathcal{G}(\Delta)$ induces a unique unfolding $A_1(x_1, \dots, x_{\#A_1}) \xrightarrow{(r_1, i_1) \dots (r_n, i_n)}_\Delta \phi$ (Def. 9). Since the vertices of $\mathcal{G}(\Delta)$ are pairs (A, t) , where t is a satisfiable base tuple and the edges of $\mathcal{G}(\Delta)$ reflect the construction of the base tuples from the least solution of the constraints (1), the outcome ϕ of this unfolding is always a satisfiable formula.

Lemma 8. *Given a path $(A_0, t_0) \xrightarrow{(r_1, i_1)} \dots \xrightarrow{(r_n, i_n)} (A_n, t_n)$ in $\mathcal{G}(\Delta)$, where $t_0 = (C^\sharp, I^\sharp, \pi)$, a state map ρ and a store v , such that $(\emptyset, \emptyset, \rho) \models^v \pi$, there exists a configuration (C, I, ρ) , such that $(C, I, \rho) \models_\Delta^v \phi$, where $A_0(x_1, \dots, x_{\#A_0}) \xrightarrow{(r_1, i_1) \dots (r_n, i_n)}_\Delta \phi$ is the unique unfolding corresponding to the path.*

Proof. Let r_1 be the following rule:

$$A_0(x_1, \dots, x_{\#A_0}) \leftarrow \phi, \text{ where } \phi = \exists y_1 \dots \exists y_m . \Psi * \pi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$$

and $\Psi * \pi$ is a quantifier- and predicate-free formula and π is, moreover, pure. The proof goes by induction on the length $n \geq 1$ of the path. For the base case $n = 1$, by Def. 10, the edge $(A_0, t_0) \xrightarrow{(r_1, i_1)} (A_1, t_1)$ implies the existence of base tuples $u_\ell \in \mu \vec{X}. \Delta^\sharp(B_\ell)$, for all $\ell \in [2, h]$, such that $B_{i_1} = A_1, u_{i_1-1} = t_1$ and:

$$t_0 \in \left(\text{Base}(\Psi * \pi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=2}^h u_\ell[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell] \right) \downarrow_{x_1, \dots, x_{\#A}}$$

Let $u_\ell \stackrel{\text{def}}{=} (C_\ell^\sharp, I_\ell^\sharp, \pi_\ell)$, for all $\ell \in [2, h]$ and $\pi' \stackrel{\text{def}}{=} \pi * \bigstar_{\ell=2}^h \pi_\ell$. Since t_0 is satisfiable, there exists a store v' , that agrees with v over $x_1, \dots, x_{\#A_0}$, such that, moreover:

$$v'(x) = v'(y) \text{ only if } x \approx_{\pi'} y, \text{ for all } x, y \in \text{fv}(\Psi * \pi') \cup \bigcup_{\ell=2}^h (C_\ell^\sharp \cup \{z_i \mid \langle z_1, \dots, z_n \rangle \in I_\ell^\sharp(\tau), \tau \in \text{Inter}\}) \quad (\dagger)$$

We define the configurations $(C_1, I_1, \rho), \dots, (C_h, I_h, \rho)$ inductively, as follows:

- $C_1 \stackrel{\text{def}}{=} \{v'(y) \mid [y] \text{ occurs in } \Psi\}$,
- $I_1 \stackrel{\text{def}}{=} \{(v'(z_1), p_1, \dots, v'(z_s), p_s) \mid \langle z_1.p_1, \dots, z_t.p_t \rangle \text{ occurs in } \Psi\}$,
- for all $\ell \in [2, h]$, assuming $C_1, \dots, C_{\ell-1}$ and $I_1, \dots, I_{\ell-1}$ are defined, let:

$$\begin{aligned} \mathcal{D}_\ell &\stackrel{\text{def}}{=} \bigcup_{i=1}^{\ell-1} C_i \cup \bigcup_{i=\ell+1}^h v'(C_i^\#) \\ \mathcal{I}_\ell &\stackrel{\text{def}}{=} \bigcup_{i=1}^{\ell-1} I_i \cup \bigcup_{i=\ell+1}^h v'(I_i^\#) \end{aligned}$$

We prove first that $\mathcal{D}_\ell \cap v'(C_\ell^\#) = \emptyset$ and $\mathcal{I}_\ell \cap v'(I_\ell^\#) = \emptyset$ (we prove only the first point, the second uses a similar reasoning), by induction on $\ell \in [2, h]$. For the base case $\mathcal{D}_2 \cap v'(C_2^\#) \neq \emptyset$, we prove the points below:

- $C_1 \cap v'(C_2^\#) = \emptyset$: suppose, for a contradiction, that there exists $c \in C_1 \cap v'(C_2^\#)$, then $c = v'(y)$, for a component atom $[y]$ from Ψ and $c = v'(x)$, for some $x \in C_2^\#$. By (\dagger) , we obtain $x \approx_{\pi'} y$, contradicting the existence of t_0 .
- $v'(C_i^\#) \cap v'(C_2^\#) = \emptyset$, for some $i \in [3, h]$: suppose, for a contradiction, that there exists $c \in v'(C_i^\#) \cap v'(C_2^\#)$, then $c = v'(x) = v'(y)$, for some $x \in C_i^\#$ and $y \in C_2^\#$. By (\dagger) , we obtain $x \approx_{\pi'} y$, contradicting the existence of t_0 .

We assume that $\mathcal{D}_j \cap v'(C_j^\#) = \emptyset$, for all $j \in [2, \ell-1]$. By Lemma 3, there exist configurations (C_j, I_j, ρ) , such that $C_j \cap \mathcal{D}_j = \emptyset$ (\dagger) and $(C_j, I_j, \rho) \models^{v'} B_j(x_1, \dots, x_{\#B_j})$, for all $j \in [2, \ell-1]$. We prove $\mathcal{D}_\ell \cap v'(C_\ell^\#) = \emptyset$, by showing the following points:

- $C_j \cap v'(C_\ell^\#) = \emptyset$, for all $j \in [1, \ell-1]$: suppose, for a contradiction, that there exists $c \in C_j \cap v'(C_\ell^\#)$, for some $j \in [1, \ell-1]$, then $c \in C_j \cap \mathcal{D}_j$, because $\mathcal{D}_j \subseteq v'(C_\ell^\#)$, in contradiction with $C_j \cap \mathcal{D}_j = \emptyset$ (\dagger) .
- $v'(C_j^\#) \cap v'(C_\ell^\#) = \emptyset$, for all $j \in [2, \ell-1]$: suppose, for a contradiction, that there exists $c \in v'(C_j^\#) \cap v'(C_\ell^\#)$, then $c = v'(x) = v'(y)$, for some $x \in C_j^\#$ and $y \in C_\ell^\#$. By (\dagger) , we obtain $x \approx_{\pi'} y$, contradicting the existence of t_0 .

Consequently, $\mathcal{D}_\ell \cap v'(C_\ell^\#) = \emptyset$, for all $\ell \in [2, h]$. By Lemma 3, there exists a configuration (C_ℓ, I_ℓ, ρ) , such that $C_\ell \cap \mathcal{D}_\ell = \emptyset$ and $(C_\ell, I_\ell, \rho) \models^{v'} B_\ell(x_1, \dots, x_{\#B_\ell})$, for all $\ell \in [2, h]$. We obtain that $C_i \cap C_j = \emptyset$ and $I_i \cap I_j = \emptyset$, for all $1 \leq i < j \leq h$, meaning that the configuration $(C, I, \rho) \stackrel{\text{def}}{=} (C_1, I_1, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$ is defined, which leads to $(C, I, \rho) \models^{v'} \phi$.

For the inductive step $n > 1$, by Def. 10, there exists base tuples $u_\ell \in \mu \vec{X}. \Delta^\#(B_\ell)$, for all $\ell \in [2, h]$, such that $A_1 = B_{i_1}$, $u_{i_1-1} = t_1$ and:

$$t_0 \in \left(\text{Base}(\Psi * \pi, \{x_1, \dots, x_{\#A}\}) \otimes \bigotimes_{\ell=1}^h u_\ell[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell] \right) \downarrow_{x_1, \dots, x_{\#A}}$$

Then there exists a store v' that agrees with v over $x_1, \dots, x_{\#A_0}$ and satisfies (\dagger) . Let $v'' \stackrel{\text{def}}{=} v'[x_1/z_1^{i_1}, \dots, x_{\#A_1}/z_{\#A_1}^{i_1}]$. By the inductive hypothesis, since $(A_1, t_1) \xrightarrow{(r_2, i_2)} \dots \xrightarrow{(r_n, i_n)} (A_n, t_n)$ is a path in $\mathcal{G}(\Delta)$, there exists a configuration (C', I', ρ) , such that $(C', I', \rho) \models^{v''} A_1(z_1^{i_1}, \dots, z_{\#A_1}^{i_1})$, because $(C', I', \rho) \models^{v'} \phi_1$, for the unfolding $A_1(x_1, \dots, x_{\#A_1}) \xrightarrow{(r_2, i_2) \dots (r_n, i_n)}_\Delta \phi_1$. The required configuration is defined as $(C, I, \rho) \stackrel{\text{def}}{=} (C_1, I_1, \rho) \bullet \dots \bullet (C_{i_1-1}, I_{i_1-1}, \rho) \bullet$

$(C', I', \rho) \bullet (C_{i_1+1}, I_{i_1+1}, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$, where $(C_1, I_1, \rho), \dots, (C_{i_1-1}, I_{i_1-1}, \rho)$ and $(C_{i_1+1}, I_{i_1+1}, \rho), \dots, (C_h, I_h, \rho)$ are defined as in the base case, by taking v'' instead of v' and defining, for all $\ell \in [2, h] \setminus \{i_1\}$:

$$\begin{aligned}\mathcal{D}_\ell &\stackrel{\text{def}}{=} C' \cup \bigcup_{i=1}^{\ell-1} C_i \cup \bigcup_{i=\ell+1}^h v'(C_i^\sharp) \\ \mathcal{I}_\ell &\stackrel{\text{def}}{=} I' \cup \bigcup_{i=1}^{\ell-1} I_i \cup \bigcup_{i=\ell+1}^h v'(I_i^\sharp)\end{aligned}$$

The proof of the fact that $C_1, \dots, C_{i_1-1}, C', C_{i_1+1}, \dots, C_h$ and $I_1, \dots, I_{i_1-1}, I', I_{i_1+1}, \dots, I_h$ are pairwise disjoint, respectively, follows by the same argument as in the base case. \square

Lemma 9. *Given an unfolding $A_0(x_1, \dots, x_{\#A_0}) \xrightarrow{(r_1, i_1) \dots (r_n, i_n)}_\Delta \phi$, a configuration (C, I, ρ)*

and a store v , such that $(C, I, \rho) \models_\Delta^v \phi$, then $\mathcal{G}(\Delta)$ has a path $(A_0, (C_0^\sharp, I_0^\sharp, \pi_0)) \rightsquigarrow^{(r_1, i_1)} \dots \rightsquigarrow^{(r_n, i_n)} (A_n, (C_n^\sharp, I_n^\sharp, \pi_n))$, for some $(C_0^\sharp, I_0^\sharp, \pi_0), \dots,$

$(C_n^\sharp, I_n^\sharp, \pi_n) \in \text{SatBase}$, such that $v(C_0^\sharp) \subseteq C_0, v(I_0^\sharp) \subseteq I_0$ and $(\emptyset, \emptyset, \rho) \models^v \pi_0$.

Proof. Let r_1 be the following rule:

$$A_0(x_1, \dots, x_{\#A_0}) \leftarrow \exists y_1 \dots \exists y_m . \Psi * \pi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$$

and $\Psi * \pi$ is a quantifier- and predicate-free formula and π is, moreover, pure. The proof goes by induction on the length $n \geq 1$ of the path. For the base case $n = 1$, we have $\phi = \exists y_1 \dots \exists y_m . \Psi * \pi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, hence there exists a store v' , that agrees with v over $x_1, \dots, x_{\#A_0}$, and configurations $(C_1, I_1, \rho), \dots, (C_h, I_h, \rho)$, such that:

- $(C_1, I_1, \rho) \models^{v'} \Psi * \pi$,
- $(C_\ell, I_\ell, \rho) \models_\Delta^{v'} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for all $\ell \in [2, h]$, and
- $\gamma = (C_1, I_1, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$.

We consider the following base tuples:

- $(\overline{C}_1^\sharp, \overline{I}_1^\sharp, \overline{\pi}_1) \stackrel{\text{def}}{=} \text{Base}(\Psi * \pi, \{x_1, \dots, x_{\#A_0}\})$,
- for all $\ell \in [2, h]$, there exist $(\overline{C}_\ell^\sharp, \overline{I}_\ell^\sharp, \overline{\pi}_\ell) \in \mu \vec{\mathcal{X}} . \Delta^\sharp(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]$, such that $C_\ell \subseteq v'(\overline{C}_\ell^\sharp)$, $I_\ell \subseteq v'(\overline{I}_\ell^\sharp)$ and $(\emptyset, \emptyset, \rho) \models^{v'} \overline{\pi}_\ell$, by Lemma 4.

By similar argument to the one from the proof of Lemma 4 (base case), we show that the composition $\bigotimes_{\ell=1}^h (\overline{C}_\ell^\sharp, \overline{I}_\ell^\sharp, \overline{\pi}_\ell)$ is defined and let $(C_0^\sharp, I_0^\sharp, \pi_0) \stackrel{\text{def}}{=} \left(\bigotimes_{\ell=1}^h (\overline{C}_\ell^\sharp, \overline{I}_\ell^\sharp, \overline{\pi}_\ell) \right) \downarrow_{x_1, \dots, x_{\#A_0}}$.

Moreover, we obtain $v(C_0^\sharp) \subseteq C_0$, $v(I_0^\sharp) \subseteq I_0$ and $(\emptyset, \emptyset, \rho) \models^v \pi_0$, as in the proof of Lemma 4. Then, by Def. 10, $\mathcal{G}(\Delta)$ has an edge $(A_0, (C_0^\sharp, I_0^\sharp, \pi_0)) \rightsquigarrow^{(r_1, i_1)} (B_{i_1-1}, (\overline{C}_{i_1-1}^\sharp, \overline{I}_{i_1-1}^\sharp, \overline{\pi}_{i_1-1}))$.

For the inductive step $n > 1$, let $B_{i_1-1}(x_1, \dots, x_{\#B_{i_1-1}}) \xrightarrow{(r_2, i_2) \dots (r_n, i_n)}_\Delta \phi_1$ be an unfolding, such that ϕ is obtained from $\exists y_1 \dots \exists y_m . \Psi * \pi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, by replacing $B_{i_1-1}(z_1^{i_1-1}, \dots, z_{\#B_{i_1-1}}^{i_1-1})$ with $\phi_1[x_1/z_1^{i_1-1}, \dots, x_{\#B_{i_1-1}}/z_{\#B_{i_1-1}}^{i_1-1}]$. Since $(C, I, \rho) \models_\Delta^v \phi$, there exists a store v' , that agrees with v over $x_1, \dots, x_{\#A_0}$, and configurations $(C_1, I_1, \rho), \dots, (C_h, I_h, \rho)$, where $\gamma = (C_1, I_1, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$ and the following hold:

- $(C_1, I_1, \rho) \models^v \Psi * \pi$,
- $(C_\ell, I_\ell, \rho) \models_\Delta^v B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for all $\ell \in [2, h] \setminus \{i_1 + 1\}$,
- $(C_{i_1-1}, I_{i_1-1}, \rho) \models_\Delta^v \phi_1[x_1/z_1^{i_1-1}, \dots, x_{\#B_{i_1-1}}/z_{\#B_{i_1-1}}^{i_1-1}]$, hence $(C_{i_1-1}, I_{i_1-1}, \rho) \models_\Delta^v B_{i_1-1}(z_1^{i_1-1}, \dots, z_{\#B_{i_1-1}}^{i_1-1})$.

We consider the following base tuples:

- $(\bar{C}_1^\#, \bar{I}_1^\#, \bar{\pi}_1) \stackrel{\text{def}}{=} \text{Base}(\Psi * \pi, \{x_1, \dots, x_{\#A_0}\})$,
- for all $\ell \in [2, h] \setminus \{i_1 + 1\}$, there exist $(\bar{C}_\ell^\#, \bar{I}_\ell^\#, \bar{\pi}_\ell) \in \mu \vec{X}. \Delta^\#(B_\ell)[x_1/z_1^\ell, \dots, x_{\#B_\ell}/z_{\#B_\ell}^\ell]$, such that $C_\ell \subseteq v'(\bar{C}_\ell^\#)$, $I_\ell \subseteq v'(\bar{I}_\ell^\#)$ and $(\emptyset, \emptyset, \rho) \models^v \bar{\pi}_\ell$, by Lemma 4.
- $\mathcal{G}(\Delta)$ has a path $(B_{i_1-1}, (C_1^\#, I_1^\#, \pi_1)) \xrightarrow{(r_2, i_2)} \dots \xrightarrow{(r_n, i_n)} (A_n, (C_n^\#, I_n^\#, \pi_n))$, such that $C_{i_1-1} \subseteq v'(C_1^\#)$, $I_{i_1-1} \subseteq v'(I_1^\#)$ and $(\emptyset, \emptyset, \rho) \models^v \pi_1$, by the inductive hypothesis.

By an argument similar to the one from Lemma 4, the composition $(\bar{C}^\#, \bar{I}^\#, \bar{\pi}) \stackrel{\text{def}}{=} \bigotimes_{\ell=1}^{i_1-2} (\bar{C}_\ell^\#, \bar{I}_\ell^\#, \bar{\pi}_\ell) \otimes (C_1^\#, I_1^\#, \pi_1) \otimes \bigotimes_{\ell=i_1}^h (\bar{C}_\ell^\#, \bar{I}_\ell^\#, \bar{\pi}_\ell)$ is defined and let $(C_0^\#, I_0^\#, \pi_0) \stackrel{\text{def}}{=} (\bar{C}^\#, \bar{I}^\#, \bar{\pi}) \downarrow_{\{x_1, \dots, x_{\#A_0}\}}$. Finally, the conditions $v(C_0^\#) \subseteq C_0$, $v(I_0^\#) \subseteq I_0$ and $(\emptyset, \emptyset, \rho) \models^v \pi_0$ follow from a similar argument to the one used in Lemma 4. \square

An *elementary cycle* of $\mathcal{G}(\Delta)$ is a path from some vertex (B, u) back to itself, such that (B, u) does not occur on the path, except at its endpoints. The cycle is, moreover, *reachable* from (A, t) if and only if there exists a path $(A, t) \xrightarrow{(r_1, i_1)} \dots \xrightarrow{(r_n, i_n)} (B, u)$ in $\mathcal{G}(\Delta)$. We reduce the complement of the $\text{Bnd}[\Delta, A]$ problem, namely the existence of an infinite set of models of $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$ of unbounded degree, to the existence of a reachable elementary cycle in $\mathcal{G}(\Delta')$, where Δ' is obtained from Δ , as described in the following.

First, we consider, for each predicate $B \in \text{def}(\Delta)$, a predicate B' , of arity $\#B + 1$, not in $\text{def}(\Delta)$ i.e., the set of predicates for which there exists a rule in Δ . Second, for each rule $B_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$, where ϕ is a quantifier- and predicate-free formula and $\text{iv}(\phi) \subseteq \text{fv}(\phi)$ denotes the subset of variables occurring in interaction atoms in ϕ , the SID Δ' has the following rules:

$$B'_0(x_1, \dots, x_{\#B_0}, x_{\#B_0+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * \bigstar_{\xi \in \text{iv}(\phi)} x_{\#B_0+1} \neq \xi * \bigstar_{\ell=2}^h B'_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell, x_{\#B_0+1}) \quad (2)$$

$$B'_0(x_1, \dots, x_{\#B_0}, x_{\#B_0+1}) \leftarrow \exists y_1 \dots \exists y_m . \phi * x_{\#B_0+1} = \xi * \bigstar_{\ell=2}^h B'_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell, x_{\#B_0+1}) \quad (3)$$

for each variable $\xi \in \text{iv}(\phi)$, that occurs in an interaction atom in ϕ .

Intuitively, there exists a family of models (with respect to Δ) of $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$ of unbounded degree if and only if these are models of $\exists x_1 \dots \exists x_{\#A+1} . A'(x_1, \dots, x_{\#A+1})$ (with respect to Δ') and the last parameter of each predicate $B' \in \text{def}(\Delta')$ can be mapped, in each of the these models, to a component that occurs in unboundedly many interactions. The latter condition is equivalent to the existence of an elementary cycle, containing a rule of the form (3), that it, moreover, reachable from some vertex (A', t) of $\mathcal{G}(\Delta')$, for some $t \in \text{SatBase}$. This reduction is formalized below:

Lemma 10. *Let A be a predicate and γ be a model of $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$. Then there exists an unfolding $A(x_1, \dots, x_{\#A}) \xrightarrow{w}_\Delta \Psi$ of length $|w| \geq \frac{\log(\delta(\gamma)) - \log \beta_1}{\log \beta_2}$ where β_1*

is the maximal number of components and interaction atoms and β_2 is the maximal number of predicate atoms, occurring in a rule of Δ .

Proof. Let γ be a configuration, v be a store and A be a predicate, such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. We consider the derivation tree T induced by the definition of the \models_{Δ}^v relation. The nodes of T are labelled by a triple $\gamma' \models_{\Delta}^{v'} B(x_1, \dots, x_{\#B})$. We start from the root labelled by $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$ and define the children of a node inductively.

For each node $\gamma' \models_{\Delta}^{v'} B(x_1, \dots, x_{\#B})$, there exists a rule:

$$r : B(x_1, \dots, x_{\#B}) \leftarrow \exists y_1 \dots \exists y_m . \phi * B_1(z_1^1, \dots, z_{\#B_1}^1) * \dots * B_h(z_1^h, \dots, z_{\#B_h}^h)$$

and configurations $\gamma_0, \dots, \gamma_h$, such that $\gamma = \gamma_0 \bullet \dots \bullet \gamma_h$, $\gamma_0 \models_{\Delta}^{v''} \phi$ and $\gamma_i \models_{\Delta}^{v''} B_i(z_1^i, \dots, z_{\#B_i}^i)$ for every $i \in [1, h]$, where ϕ is a predicate-free formula and v'' is a store that agrees with v' over $x_1, \dots, x_{\#B}$. We define $v_{\ell} \stackrel{\text{def}}{=} v''[x_1/z_{\#B_{\ell}}^{\ell}, \dots, x_{\#B_{\ell}}^{\ell}/z_{\#B_{\ell}}^{\ell}]$, for all $\ell \in [1, h]$. Then the node $\gamma' \models_{\Delta}^{v'} B(x_1, \dots, x_{\#B})$ has h children in T , where the ℓ -th child is labelled by $\gamma_{\ell} \models_{\Delta}^{v_{\ell}} B_{\ell}(x_1, \dots, x_{\#B_{\ell}})$, for all $\ell \in [1, h]$. The construction is finite since $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$ has a finite inductive definition.

We now consider the degree of the configurations which occur in T . By Def. 3, we obtain $\delta(\gamma) \leq \delta(\gamma_0) + \sum_{i=1}^h \delta(\gamma_i)$. With each γ_i associated to a child of this node (except γ_0), we obtain that $\delta(\gamma_{\text{init}})$ does not exceed β_1 times the number of nodes in T . Since $h \leq \beta_2$, the height n of T is bound to the degree $\delta(\gamma)$ by the inequality $\delta(\gamma) \leq \beta_1 \times \sum_{k=0}^n \beta_2^k = \beta_1 \times \beta_2^{n+1} - 1$, leading to:

$$n + 1 \geq \frac{\log \delta(\gamma) - \log \beta_1}{\log \beta_2}$$

Finally, with T of height n , there exists a branch in T (starting from the root) of length exactly $n + 1$. Yet each branch of T corresponds to an unfolding $A(x_1, \dots, x_{\#A}) \xRightarrow{w}_{\Delta} \Psi$, with w obtained by concatenating for every node (γ, v, A) of the branch (from root to leaf) the couple (r, i) consisting of:

- the rule $r \in \Delta$ used to unfold $A(x_1, \dots, x_{\#A})$, and
- the position i of this node among its brothers in T (take $i = 1$ for the root).

This unfolding has the length required, which concludes the lemma. \square

Lemma 11. *There exists an infinite sequence of configurations $\gamma_1, \gamma_2, \dots$ such that $\gamma_i \models_{\Delta} \exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$ and $\delta(\gamma_i) < \delta(\gamma_{i+1})$, for all $i \geq 1$ if and only if $\mathcal{G}(\Delta')$ has an elementary cycle containing a rule (3), reachable from a node (A', t) , for $t \in \text{SatBase}$.*

Proof. “ \Rightarrow ” Let v_1, v_2, \dots be stores such that $\gamma_i \models_{\Delta}^{v_i} A(x_1, \dots, x_{\#A})$, for all $i \geq 1$. By Lemma 10, there exists unfoldings $A(x_1, \dots, x_{\#A}) \xRightarrow{w_i}_{\Delta'} \phi_i$ of lengths $|w_1| < |w_2| < \dots$, such that $\gamma_i \models_{\Delta}^{v_i} \phi_i$, for all $i \geq 1$. For each configuration γ_i , let $d_i \in \mathbb{C}$ be a component, such that $\delta(\gamma_i) = \|\{(c_1, p_1, \dots, c_n, p_n) \mid d_i = c_j, j \in [1, n]\}\|$. By induction on $|w_i| \geq 1$, we build unfoldings $A'(x_1, \dots, x_{\#A}, x_{\#A+1}) \xRightarrow{w'_i}_{\Delta'} \phi'_i$ that bind $x_{\#A+1}$ to all variables bound to d_i , using rules of type (3). By Lemma 9, w'_1, w'_2, \dots are labels of paths from $\mathcal{G}(\Delta')$, that

start in $(A, t_1), (A, t_2), \dots$, respectively. Since $\mathcal{G}(\Delta')$ is finite, we can chose an infinite subsequence of paths that start in the same node of $\mathcal{G}(\Delta')$ and repeat the same vertex, with a rule of type (3) in between.

“ \Leftarrow ” Let $(A', t) \xrightarrow{(r'_1, i_1)} \dots \xrightarrow{(r'_n, i_n)} (B_n, t_n) \xrightarrow{(r'_{n+1}, i_{n+1})} \dots \xrightarrow{(r'_{n+p}, i_{n+p})} (B_n, t_n)$ be a path in $\mathcal{G}(\Delta')$, such that one of the rules $r'_{n+1}, \dots, r'_{n+p}$ is of the form (3) and let $w'_i \stackrel{\text{def}}{=} (r'_1, i_1) \dots (r'_n, i_n) [(r'_{n+1}, i_{n+1}) \dots (r'_{n+p}, i_{n+p})]^i$, for all $i \geq 1$. By Lemma 8, there exist unfoldings $A'(x_1, \dots, x_{\#A+1}) \xrightarrow{w'_i}_{\Delta'} \phi'_i$, stores v_i and configurations γ_i , such that $\gamma_i \models^{v_i} \phi'_i$. We define:

$$\delta_i \stackrel{\text{def}}{=} \|\{(c_1, p_1, \dots, c_n, p_n) \in I_i \mid v(x_{\#A+1}) = c_j, j \in [1, n]\}\|, \text{ for all } i \geq 1$$

where $\gamma_i \stackrel{\text{def}}{=} (C_i, I_i, p_i)$. Since $\gamma_i \models^{v_i} \phi'_i$ and one of the rules $r'_{n+1}, \dots, r'_{n+p}$ is of type (3), the sequence $\delta_1, \delta_2, \dots$ is strictly increasing. Moreover, we have $\delta_i \leq \delta(\gamma_i)$, for all $i \geq 1$, hence there exists a sequence of integers $1 \leq i_1 < i_2 < \dots$ such that $\delta(\gamma_{i_j}) < \delta(\gamma_{i_{j+1}})$, for all $j \geq 1$. \square

The complexity result below uses a similar argument on the maximal size of (hence the number of) base tuples as in Theorem 1, leading to similar complexity gaps:

Theorem 3. $\text{Bnd}^{(k, \infty)}[\Delta, A]$ is in co-NP, $\text{Bnd}^{(\infty, \ell)}[\Delta, A]$ is in EXP and $\text{Bnd}[\Delta, A]$ is in 2EXP.

Proof. Lemma 11 shows the reduction of the complement of $\text{Bnd}[\Delta, A]$ to the existence of a reachable cycle in the graph $\mathcal{G}(\Delta')$, where Δ' is constructed from Δ in polynomial time. Moreover, we have $\text{arity}(\Delta') = \text{arity}(\Delta) + 1$ and $\text{intersize}(\Delta') = \text{intersize}(\Delta)$. We distinguish the three cases below:

- $k < \infty, \ell = \infty$: in this case, we can define a non-deterministic algorithm as follows. We guess the solution $(\langle W_1, \dots, W_K \rangle, W_{K+1}, \langle i_1, i_2, \dots, i_n \rangle)$, where:
 - $\langle W_1, \dots, W_K \rangle$ defines an acyclic witness for a satisfiable least solution of A' in Δ' constructed as in the proof of Thm. 1;
 - $W_{K+1} = (T_{K+1}, r_{K+1}, t_{K+1,1}, \dots, e_{K+1, h_{K+1}})$ is similar to a regular entry W_i , that is, contains a base tuple T_{K+1} , an index r_{K+1} of a rule of Δ' and indices $e_{K+1,1}, \dots, e_{K+1, h_{K+1}} \in \{1, \dots, K\}$ such that T_{K+1} is computed correctly by applying the rule r_{K+1} from base tuples $T_{e_{K+1,1}}, \dots, T_{e_{K+1, h_{K+1}}}$ as explained in the proof of Thm. 1;
 - $\langle i_1, i_2, \dots, i_n \rangle$ defines an acyclic path starting at the initial node in the directed acyclic graph defined by W , that is, $1 = i_1 < i_2 < \dots < i_n \leq K$ and moreover $i_{j+1} \in \{e_{i_j,1}, \dots, e_{i_j, h_{i_j}}\}$ for all $j \in \{1, 2, \dots, n-1\}$;
 - the path $\langle i_1, i_2, \dots, i_n \rangle$ can be closed into a witness reachable cycle from i_1 by using W_{K+1} that is, whenever (i) rules r_{i_n} and r_K define the same predicate, and moreover $T_{i_n} = T_{K+1}$, (ii) the intersection $X = \{e_{K+1,1}, \dots, e_{K+1, h_{K+1}}\} \cap \{i_1, i_2, \dots, i_n\} \neq \emptyset$, (iii) if $i_j = \min X$, that is, the cycle starts at i_j then at least one of the rules used along the cycle $r_{i_j}, r_{i_{j+1}}, \dots, r_{i_{n-1}}, r_{K+1}$ is of the form (3).

The solution is of linear size $O(\text{size}(\Delta'))$ by the same arguments as in the proof of Thm. 1. Therefore, it can be guessed in polynomial time, and moreover checked in polynomial time following the conditions above. This implies the membership of the complement problem in NP, henceforth $\text{Bnd}^{(k,\infty)}[\Delta, A]$ is in co-NP.

- $k = \infty, \ell < \infty$: in this case, using the algorithm from Fig. 3, the graph $\mathcal{G}(\Delta')$ is constructed in time $2^{\text{poly}(\text{size}(\Delta'))}$ as previously explained in the proof of Theorem 1. Finding a reachable cycle with the additional properties required by Lemma 11 can be done in two additional steps, respectively, first building the SCCs decomposition of $\mathcal{G}(\Delta')$ and then checking reachability of SCCs containing edges derived from rules of form (3) from SCCs containing vertices (A', t) . Both steps can be done in linear time in the size of $\mathcal{G}(\Delta')$ i.e., using Tarjan algorithm for SCC decomposition and standard graph traversals. Therefore, the overall time complexity remains $2^{\text{poly}(\text{size}(\Delta'))}$, and as such $\text{Bnd}^{(\infty,\ell)}[\Delta, A]$ is in EXP.
- $k = \infty, \ell = \infty$: following the same argument as in the previous point and noticing that the graph $\mathcal{G}(\Delta')$ is constructed in time $2^{\text{poly}(\text{size}(\Delta'))}$ we conclude that $\text{Bnd}[\Delta, A]$ is in 2EXP. \square

Moreover, the construction of $\mathcal{G}(\Delta')$ allows to prove the following cut-off result:

Proposition 1. *Let γ be a configuration and v be a store, such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. If $\text{Bnd}^{(k,\ell)}[\Delta, A]$ then (1) $\delta(\gamma) = \text{poly}(\text{size}(\Delta))$ if $k < \infty, \ell = \infty$, (2) $\delta(\gamma) = 2^{\text{poly}(\text{size}(\Delta))}$ if $k = \infty, \ell < \infty$ and (3) $\delta(\gamma) = 2^{2^{\text{poly}(\text{size}(\Delta))}}$ if $k = \infty, \ell = \infty$.*

Proof. First, we show that in all cases, the degree is bounded by $2^{B^*} \cdot L \cdot I$ where B^* is the maximal length of a satisfiable base tuple in Δ' , L is the number of predicates in Δ' and I is the maximal number of interactions defined in a rule in Δ' . The maximal length B^* of a satisfiable base tuples has been considered in the proof of Thm. 1 to derive an upper bound on the number of distinct satisfiable base tuples for a SID. Then, $2^{B^*} \cdot L$ represents a bound on the number of nodes in the graph $\mathcal{G}(\Delta')$ as for every predicate there will be at most 2^{B^*} satisfiable base tuples associated to it. Meantime, this value also represents a bound on the longest acyclic path in $\mathcal{G}(\Delta')$. We are interested on acyclic paths because cycles in $\mathcal{G}(\Delta')$ are guaranteed to never connect (use in interactions) the extra variable introduced in Δ' (otherwise the system would not be of bounded degree). But then, along the acyclic paths, at most I interactions are defined at each step, henceforth, the bound of $2^{B^*} \cdot L \cdot I$ on the number on total interactions that could involve the extra variable.

Second, let observe that both L and I are the same in Δ' and in Δ and equal to $O(\text{size}(\Delta))$. Moreover, it was shown in the proof of Thm. 1 that $B^* = 2\alpha + 2\alpha^2 + p^{\min(\alpha,\beta)}\alpha^{\min(\alpha,\beta)}$ for $\alpha = \text{arity}(\Delta') = \text{arity}(\Delta) + 1$ and $\beta = \text{intersize}(\Delta') = \text{intersize}(\Delta)$, $p = \|\mathcal{P}\|$ the number of ports. Henceforth, we distinguished the three cases, respectively (i) $B^* = O(1)$ if $k < \infty, \ell = \infty$, (ii) $B^* = \text{poly}(\text{size}(\Delta))$ if $k = \infty, \ell < \infty$ and (iii) $B^* = 2^{\text{poly}(\text{size}(\Delta))}$ if $k = \infty, \ell = \infty$. By using the above in the expression $2^{B^*} \cdot L \cdot I$ we obtain the values of the bound as stated in the Proposition. \square

6 Entailment

This section is concerned with the entailment problem $\text{Entl}[\Delta, A, B]$, that asks whether $\gamma \models_{\Delta}^v \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, for every configuration γ and store v , such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. For instance, the proof from Fig. 1 (c) relies on the following entailments, that occur as the side conditions of the Hoare logic rule of consequence:

$$\begin{aligned} \text{ring}_{h,t}(y) &\models_{\Delta} \exists x \exists z. [y] @ H * \langle y.out, z.in \rangle * \text{chain}_{h-1,t}(z, x) * \langle x.out, y.in \rangle \\ [z] @ H * \langle z.out, x.in \rangle * \text{chain}_{h-1,t}(x, y) * \langle y.out, z.in \rangle &\models_{\Delta} \text{ring}_{h,t}(z) \end{aligned}$$

By introducing two fresh predicates A_1 and A_2 , defined by the rules:

$$A_1(x_1) \leftarrow \exists y \exists z. [x_1] @ H * \langle x_1.out, z.in \rangle * \text{chain}_{h-1,t}(z, y) * \langle y.out, x_1.in \rangle \quad (4)$$

$$A_2(x_1, x_2) \leftarrow \exists z. [x_1] @ H * \langle x_1.out, z.in \rangle * \text{chain}_{h-1,t}(z, x_2) * \langle x_2.out, x_1.in \rangle \quad (5)$$

the above entailments are equivalent to $\text{Entl}[\Delta, \text{ring}_{h,t}, A_1]$ and $\text{Entl}[\Delta, A_2, \text{ring}_{h,t}]$, respectively, where Δ consists of the rules (4) and (5), together with the rules that define the $\text{ring}_{h,t}$ and $\text{chain}_{h,t}$ predicates (§1.1).

We show that the entailment problem is undecidable, in general (Thm. 4), and recover a decidable fragment, by means of three syntactic conditions, typically met in our examples. These conditions use the following notion of *profile*:

Definition 11. *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 time $\text{poly}(\text{size}(\Delta))$.

Definition 12. *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 quantifier- and predicate-free formula, is said to be:*

1. *progressing 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 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) if and only if, for every disequation $x \neq y$ from ϕ , we have $\{x, y\} \cap \{x_i \mid i \in \lambda_{\Delta}(A)\} \neq \emptyset$.*

A SID Δ is progressing, connected and e-restricted if and only if each rule in Δ is progressing, connected and e-restricted, respectively.

For example, the SID consisting of the rules from §1.1, together with rules (4) and (5) is progressing, connected and e-restricted. For a configuration $\gamma = (C, I, \rho)$, let:

$$\text{nodes}(\gamma) \stackrel{\text{def}}{=} C \cup \{c_i \mid (c_1, p_1, \dots, c_n, p_n) \in I, i \in [1, n]\}$$

be the set of (possibly absent) components that occur in γ .

Lemma 12. *Given a progressing SID Δ and a predicate $A \in \text{def}(\Delta)$, for any configuration $\gamma = (C, I, \rho)$ and store v , such that $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$, we have $\{v(x_1), \dots, v(x_{\#A})\} \subseteq \text{nodes}(\gamma) = C$.*

Proof. We proceed by fixpoint induction on the definition of $\gamma \models_{\Delta}^v A(x_1, \dots, x_{\#A})$. By definition, there exists a progressing rule

$$r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . [x_1] * \Psi * \bigstar_{\ell=1}^h B_{\ell}(z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell})$$

a store v' and configurations $\gamma_0, \dots, \gamma_h$ such that:

- $\gamma = \gamma_0 \bullet \dots \bullet \gamma_h$,
- $v(x_i) = v'(x_i)$ for all $i \in [1, \#A]$,
- $\gamma_0 \models_{\Delta}^{v'} [x_1] * \Psi$, and
- $\gamma_{\ell} \models_{\Delta}^{v'} B_{\ell}(z_1^{\ell}, \dots, z_{\#B_{\ell}}^{\ell})$ for all $\ell \in [1, h]$.

For $1 \leq \ell \leq h$, let $v_{\ell}(x_i) = v'(z_i^{\ell})$ for $1 \leq i \leq \#B_{\ell}$. Now apply the induction hypothesis on the derivation of $\gamma_{\ell} \models_{\Delta}^{v'} B_{\ell}(x_1, \dots, x_{\#B_{\ell}})$ to obtain that $\{v_{\ell}(x_1), \dots, v_{\ell}(x_{\#B_{\ell}})\} \subseteq \text{nodes}(\gamma_{\ell})$. Since r is progressing, we have:

$$\begin{aligned} \{v(x_1), \dots, v(x_{\#A})\} &\subseteq \{v'(x_1), \dots, v'(x_{\#A})\} \cup \{v'(y_1), \dots, v'(y_m)\} \\ &= \{v'(x_1)\} \cup \bigcup_{\ell=1}^h \{v'(z_1^{\ell}), \dots, v'(z_{\#B_{\ell}}^{\ell})\} = \{v'(x_1)\} \cup \bigcup_{\ell=1}^h \{v_{\ell}(x_1), \dots, v_{\ell}(x_{\#B_{\ell}})\} \\ &\subseteq \text{nodes}(\gamma_0) \cup \bigcup_{\ell=1}^h \text{nodes}(\gamma_{\ell}) = \text{nodes}(\gamma) \quad \square \end{aligned}$$

We recall that $\text{def}_{\Delta}(A)$ is the set of rules from Δ that define A and denote by $\text{def}_{\Delta}^*(A)$ the least superset of $\text{def}_{\Delta}(A)$ containing the rules that define a predicate from a rule in $\text{def}_{\Delta}^*(A)$. The following result shows that the entailment problem becomes undecidable as soon as the connectivity condition is even slightly lifted:

Theorem 4. *$\text{Entl}[\Delta, A, B]$ is undecidable, even when Δ is progressing and e -restricted, and only the rules in $\text{def}_{\Delta}^*(A)$ are connected (the rules in $\text{def}_{\Delta}^*(B)$ may be disconnected).*

Proof. By a reduction from the known undecidable problem of universality of context-free languages [3]. A context-free grammar $G = \langle N, T, S, \Delta \rangle$ consists of a finite set N of nonterminals, a finite set T of terminals, a start symbol $S \in N$ and a finite set Δ of productions of the form $A \rightarrow w$, where $A \in N$ and $w \in (N \cup T)^*$. Given finite strings $u, v \in (N \cup T)^*$, the step relation $u \Rightarrow v$ replaces a nonterminal A of u by the right-hand side w of a production $A \rightarrow w$ and \Rightarrow^* denotes the reflexive and transitive closure of \Rightarrow . The language of G is the set $\mathcal{L}(G)$ of finite strings $w \in T^*$, such that $s \Rightarrow^* w$. The problem $T^* \subseteq \mathcal{L}(G)$ is known as the universality problem, known to be undecidable. Moreover, we assume w.l.o.g. that:

- $T = \{0, 1\}$, because every terminal can be encoded as a binary string,
- $\mathcal{L}(G)$ does not contain the empty string ϵ , because computing a grammar G' such that $\mathcal{L}(G') = \mathcal{L}(G) \cap T^+$ is possible and, moreover, we can reduce from the modified universality problem $T^+ \subseteq \mathcal{L}(G')$ instead of the original $T^* \subseteq \mathcal{L}(G)$,

- G is in Greibach normal form, i.e. it contains only production rules of the form $B_0 \rightarrow bB_1 \dots B_n$, where $B_0, \dots, B_n \in N$, for some $n \geq 0$ and $b \in T$.

Let $\mathcal{P} = \{p_0, p_1\}$ be a set of ports. For each nonterminal $B_0 \in N$, we have a predicate B_0 or arity two and a rule $B_0(x_1, x_2) \leftarrow \exists y_1 \dots \exists y_n . [x_1] * \langle x_1.p_a, y_1.p_a \rangle * B_1(y_1, y_2) * \dots * B_n(y_n, x_2)$, for each rule $B_0 \rightarrow bB_1 \dots B_n$ of G . Moreover, we consider the rules $A(x_1, x_2) \leftarrow \exists z . \langle x_1.p_a, z.p_a \rangle * A(z, x_2)$ and $A(x_1, x_2) \leftarrow \langle x_1.p_a, x_2.p_a \rangle$, for all $a \in \{0, 1\}$. Let Δ be the SID containing the above rules. It is easy to check that the SID is progressing and e-restricted and that, moreover, the rules from $\text{def}_\Delta^*(A)$ are connected. Finally, $A(x_1, x_2) \models_\Delta B(x_1, x_2)$ if and only if $T^+ \subseteq \mathcal{L}(G)$. \square

On the positive side, we prove that $\text{Entl}[\Delta, A, B]$ is decidable, if Δ is progressing, connected and e-restricted, assuming further that $\text{Bnd}[\Delta, A]$ has a positive answer. In this case, the bound on the degree of the models of $A(x_1, \dots, x_{\#A})$ is effectively computable, using the algorithm from Fig. 3 (see Prop. 1 for a cut-off result) and denote by \mathfrak{B} this bound, throughout this section.

The proof uses a reduction of $\text{Entl}[\Delta, A, B]$ to a similar problem for SL, showed to be decidable [20]. We recall the definition of SL, interpreted over heaps $h : \mathbb{C} \xrightarrow{\text{fin}} \mathbb{C}^{\mathfrak{K}}$, introduced in §2.3. SL rules are denoted as $\bar{A}(x_1, \dots, x_{\#(\bar{A})}) \leftarrow \phi$, where ϕ is a SL formula, such that $\text{fv}(\phi) \subseteq \{x_1, \dots, x_{\#(\bar{A})}\}$ and SL SIDs are denoted as $\bar{\Delta}$. The profile $\lambda_{\bar{\Delta}}$ is defined for SL same as for CL (Def. 11).

Definition 13. A SL rule $\bar{A}(x_1, \dots, x_{\#(\bar{A})}) \leftarrow \phi$ from a SID $\bar{\Delta}$ is said to be:

1. progressing if and only if $\phi = \exists t_1 \dots \exists t_m . x_1 \mapsto (y_1, \dots, y_{\mathfrak{K}}) * \psi$, where ψ contains only predicate and equality atoms,
2. connected if and only if $z_1 \in \{x_i \mid i \in \lambda_{\bar{\Delta}}(\bar{A})\} \cup \{y_1, \dots, y_{\mathfrak{K}}\}$, for every predicate atom $\bar{B}(z_1, \dots, z_{\#(\bar{B})})$ from ϕ .

Note that the definitions of progressing and connected rules are different for SL, compared to CL (Def. 12); in the rest of this section, we rely on the context to distinguish progressing (connected) SL rules from progressing (connected) CL rules. Moreover, e-restricted rules are defined in the same way for CL and SL (point 3 of Def. 12). A tight upper bound on the complexity of the entailment problem between SL formulæ, interpreted by progressing, connected and e-restricted SIDs, is given below:

Theorem 5 ([20]). The SL entailment problem is in $2^{poly(\text{width}(\bar{\Delta}) \cdot \log \text{size}(\bar{\Delta}))}$, for progressing, connected and e-restricted SIDs.

The reduction of $\text{Entl}[\Delta, A, B]$ to SL entailments is based on the idea of viewing a configuration as a logical structure (hypergraph), represented by an indirected *Gaifman graph*, in which every tuple from a relation (hyperedge) becomes a clique [26]. In a similar vein, we encode a configuration, of degree at most \mathfrak{B} , by a heap of degree \mathfrak{K} (Def. 14), such that \mathfrak{K} is defined using the following integer function:

$$\text{pos}(i, j, k) \stackrel{\text{def}}{=} 1 + \mathfrak{B} \cdot \sum_{\ell=1}^{j-1} |\tau_\ell| + i \cdot |\tau_j| + k$$




Fig. 4: Gaifman Heap for a Chain Configuration

where $\text{Inter} \stackrel{\text{def}}{=} \{\tau_1, \dots, \tau_M\}$ is the set of interaction types and $Q \stackrel{\text{def}}{=} \{q_1, \dots, q_N\}$ is the set of states of the behavior $\mathbb{B} = (\mathcal{P}, Q, \rightarrow)$ (§2). Here $i \in [0, \mathfrak{B} - 1]$ denotes an interaction of type $j \in [1, M]$ and $k \in [0, N - 1]$ denotes a state. We use M and N throughout the rest of this section, to denote the number of interaction types and states, respectively.

For a set I of interactions, let $\text{Tuples}_I^j(c) \stackrel{\text{def}}{=} \{(c_1, \dots, c_n) \mid (c_1, p_1, \dots, c_n, p_n) \in I, \tau_j = \langle p_1, \dots, p_n \rangle, c \in \{c_1, \dots, c_n\}\}$ be the tuples of components from an interaction of type τ_j from I , that contain a given component c .

Definition 14. Given a configuration $\gamma = (C, I, \rho)$, such that $\delta(\gamma) \leq \mathfrak{B}$, a Gaifman heap for γ is a heap $h : \mathbb{C} \rightarrow_{\text{fm}} \mathbb{C}^{\mathfrak{R}}$, where $\mathfrak{R} \stackrel{\text{def}}{=} \text{pos}(0, M + 1, N)$, $\text{dom}(h) = \text{nodes}(\gamma)$ and, for all $c_0 \in \text{dom}(h)$, such that $h(c_0) = \langle c_1, \dots, c_{\mathfrak{R}} \rangle$, the following hold:

1. $c_1 = c_0$ if and only if $c_0 \in C$,
2. for all $j \in [1, M]$, $\text{Tuples}_I^j(c) = \{c_1, \dots, c_s\}$ if and only if there exist integers $0 \leq k_1 < \dots < k_s < \mathfrak{B}$, such that $\langle h(c_0) \rangle_{\text{inter}(k_i, j)} = c_i$, for all $i \in [1, s]$, where $\text{inter}(i, j) \stackrel{\text{def}}{=} [\text{pos}(i - 1, j, 0), \text{pos}(i, j, 0)]$ are the entries of the i -th interaction of type τ_j in $h(c_0)$,
3. for all $k \in [1, N]$, we have $\langle h(c_0) \rangle_{\text{state}(k)} = c_0$ if and only if $\rho(c_0) = q_k$, where the entry $\text{state}(k) \stackrel{\text{def}}{=} \text{pos}(0, M + 1, k - 1)$ in $h(c_0)$ corresponds to the state $q_k \in Q$.

We denote by $\mathbb{G}(\gamma)$ the set of Gaifman heaps for γ .

Intuitively, if h is a Gaifman heap for γ and $c_0 \in \text{dom}(h)$, then the first entry of $h(c_0)$ indicates whether c_0 is present (condition 1 of Def. 14), the next $\mathfrak{B} \cdot \sum_{j=1}^M |\tau_j|$ entries are used to encode the interactions of each type τ_j (condition 2 of Def. 14), whereas the last N entries are used to represent the state of the component (condition 3 of Definition 14). Note that the encoding of configurations by Gaifman heaps is not unique: two Gaifman heaps for the same configuration may differ in the order of the tuples from the encoding of an interaction type and the choice of the unconstrained entries from $h(c_0)$, for each $c_0 \in \text{dom}(h)$. On the other hand, if two configurations have the same Gaifman heap encoding, they must be the same configuration.

Example 4. Fig. 4 (b) shows a Gaifman heap for the configuration in Fig. 4 (a), where each component belongs to at most 2 interactions of type $\langle \text{out}, \text{in} \rangle$. ■

We say that a configuration γ' is a *subconfiguration* of γ , denoted $\gamma' \sqsubseteq \gamma$ if and only if $\gamma = \gamma' \bullet \gamma''$, for some configuration γ'' . The following lemma builds Gaifman heaps for subconfigurations:

Lemma 13. Given configurations γ and γ' , such that $\gamma' \sqsubseteq \gamma$, if $h \in \mathbb{G}(\gamma)$, then $h' \in \mathbb{G}(\gamma')$, where $\text{dom}(h') = \text{dom}(h) \cap \text{nodes}(\gamma')$ and $h'(c) = h(c)$, for all $c \in \text{dom}(h')$.

Proof. We have $\text{dom}(h') = \text{dom}(h) \cap \text{nodes}(\gamma') = \text{nodes}(\gamma) \cap \text{nodes}(\gamma') = \text{nodes}(\gamma')$, because $\text{dom}(h) = \text{nodes}(\gamma) \supseteq \text{nodes}(\gamma')$. The points (1-3) of Def. 14 are by easy inspection. \square

We build a SL SID $\bar{\Delta}$ that generates the Gaifman heaps of the models of the predicate atoms occurring in a progressing CL SID Δ . The construction associates to each variable x , that occurs free or bound in a rule from Δ , a unique \mathfrak{K} -tuple of variables $\eta(x) \in \mathbb{V}^{\mathfrak{K}}$, that represents the image of the store value $v(x)$ in a Gaifman heap h i.e., $h(v(x)) = v(\eta(x))$. Moreover, we consider, for each predicate symbol $A \in \text{def}(\Delta)$, an annotated predicate symbol \bar{A}_ι of arity $\# \bar{A}_\iota = (\mathfrak{K} + 1) \cdot \#A$, where $\iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$ is a map associating each parameter $i \in [1, \#A]$ and each interaction type τ_j , for $j \in [1, M]$, a set of integers $\iota(i, j)$ denoting the positions of the encodings of the interactions of type τ_j , involving the value of x_i , in the models of $\bar{A}_\iota(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A}))$ (point 2 of Def. 14). Then $\bar{\Delta}$ contains rules of the form:

$$\begin{aligned} \bar{A}_\iota(x_1, \dots, x_{\#(A)}, \eta(x_1), \dots, \eta(x_{\#(A)})) \leftarrow \\ \exists y_1 \dots \exists y_m \exists \eta(y_1) \dots \exists \eta(y_m) . \bar{\Psi} * \pi * \bigstar_{\ell=1}^h \bar{B}_\iota^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell, \eta(z_1^\ell), \dots, \eta(z_{\#(B^\ell)}^\ell)) \end{aligned} \quad (6)$$

for which Δ has a *stem rule* $A(x_1, \dots, x_{\#(A)}) \leftarrow \exists y_1 \dots \exists y_m . \Psi * \pi * \bigstar_{\ell=1}^h B^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell)$, where $\Psi * \pi$ is a quantifier- and predicate-free formula and π is the conjunction of equalities and disequalities from $\Psi * \pi$. However, not all rules (6) are considered in $\bar{\Delta}$, but only the ones meeting the following condition:

Definition 15. A rule of the form (6) is well-formed if and only if, for each $i \in [1, \#A]$ and each $j \in [1, M]$, there exists a set of integers $Y_{i,j} \subseteq [0, \mathfrak{B} - 1]$, such that:

- $\|Y_{i,j}\| = \|I_{\Psi, \pi}^j(x_i)\|$, where $I_{\Psi, \pi}^j(x)$ is the set of interaction atoms $\langle z_1.p_1, \dots, z_n.p_n \rangle$ from Ψ of type $\tau_j = \langle p_1, \dots, p_n \rangle$, such that $z_s \approx_\pi x$, for some $s \in [1, n]$,
- $Y_{i,j} \subseteq \iota(i, j)$ and $\iota(i, j) \setminus Y_{i,j} = Z_j(x_i)$, where $Z_j(x) \stackrel{\text{def}}{=} \bigcup_{\ell=1}^h \bigcup_{k=1}^{\#B^\ell} \{\iota^\ell(k, j) \mid x \approx_\pi z_k^\ell\}$ is the set of positions used to encode the interactions of type τ_j involving the store value of the parameter x , in the sub-configuration corresponding to an atom $B^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell)$, for some $\ell \in [1, h]$.

We denote by $\bar{\Delta}$ the set of well-formed rules (6), such that, moreover:

$$\begin{aligned} \bar{\Psi} &\stackrel{\text{def}}{=} x_1 \mapsto \eta(x_1) * \bigstar_{x \in \text{fv}(\Psi)} \text{CompStates}_\Psi(x) * \bigstar_{i=1}^{\#A} \text{InterAtoms}_\Psi(x_i), \text{ where:} \\ \text{CompStates}_\Psi(x) &\stackrel{\text{def}}{=} \bigstar_{[x] \text{ occurs in } \Psi} \langle \eta(x) \rangle_1 = x * \bigstar_{x @ q_k \text{ occurs in } \Psi} \langle \eta(x) \rangle_{\text{state}(k)} = x \\ \text{InterAtoms}_\Psi(x_i) &\stackrel{\text{def}}{=} \bigstar_{j=1}^M \bigstar_{p=1}^{r_j} \langle \eta(x_i) \rangle_{\text{inter}(j, k_p^j)} = \mathbf{x}_p^j \text{ and } \{k_1^j, \dots, k_{r_j}^j\} \stackrel{\text{def}}{=} \iota(i, j) \setminus Z_j(x_i) \end{aligned}$$

Here for two tuples of variables $\mathbf{x} = \langle x_1, \dots, x_k \rangle$ and $\mathbf{y} = \langle y_1, \dots, y_k \rangle$, we denote by $\mathbf{x} = \mathbf{y}$ the formula $\bigstar_{i=1}^k x_i = y_i$. Intuitively, the SL formula $\text{CompStates}_\Psi(x)$ realizes the encoding of the component and state atoms from Ψ , in the sense of points (1) and (3) from Def. 14, whereas the formula $\text{InterAtoms}_\Psi(x_i)$ realizes the encodings of the interactions involving a parameter x_i in the stem rule (point 2 of Def. 14). In particular, the definition of $\text{InterAtoms}_\Psi(x_i)$ uses the fact that the rule is well-formed.

Lemma 14. *Let Δ be a progressing SID and $A \in \text{def}(\Delta)$ be a predicate, such that $\gamma \models^v A(x_1, \dots, x_{\#A})$, for some configuration $\gamma = (C, I, \rho)$ and store v . Then, for each heap $h \in \mathbb{G}(\gamma)$, there exists a map $\iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$ and a store \bar{v} , such that the following hold:*

1. $\bar{v}(x_i) = v(x_i) \in \text{dom}(h)$ and $\bar{v}(\eta(x_i)) = h(v(x_i))$, $\forall i \in [1, \#A]$,
2. $\text{Tuples}_{\bar{v}}^j(\bar{v}(x_i)) = \{ \langle \bar{v}(\eta(x_i)) \rangle_{\text{inter}(j,k)} \mid k \in \iota(i, j) \}$, $\forall i \in [1, \#A] \forall j \in [1, M]$,
3. $h \models_{\bar{\Delta}}^{\bar{v}} \bar{A}_\iota(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#(A)}))$.

Proof. By induction on the definition of $\gamma \models^v A(x_1, \dots, x_{\#A})$, assume that $\gamma \models^v \exists y_1 \dots \exists y_m \cdot \Psi * \pi * \bigstar_{\ell=1}^h B^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell)$, where $A(x_1, \dots, x_{\#(A)}) \leftarrow \exists y_1 \dots \exists y_m \cdot \Psi * \pi * \bigstar_{\ell=1}^h B^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell)$ is a rule from Δ , such that $\Psi * \pi$ is quantifier- and predicate-free and π is the conjunction of equalities and disequalities from $\Psi * \pi$. Then there exists a store v' , that agrees with v over $x_1, \dots, x_{\#A}$, and configurations $\gamma_0 = (C_0, I_0, \rho), \dots, \gamma_h = (C_h, I_h, \rho)$, such that:

- $\gamma_0 \models_{\bar{\Delta}}^{v'} \Psi * \pi$,
- $\gamma_\ell \models_{\bar{\Delta}}^{v'} B_\ell(z_1^\ell, \dots, z_{\#(B_\ell)}^\ell)$, for all $\ell \in [1, h]$, and
- $\gamma = \gamma_0 \bullet \dots \bullet \gamma_h$.

We define the heaps h_0, \dots, h_h , as follows:

- for each $\ell \in [1, h]$, let $\text{dom}(h_\ell) = \text{nodes}(\gamma_\ell)$, $h_\ell(c) = h(c)$, for all $c \in \text{dom}(h_\ell)$,
- $h_0 \stackrel{\text{def}}{=} h \setminus (\bigcup_{\ell=1}^h h_\ell)$.

By Lemma 13, we obtain that $h_\ell \in \mathbb{G}(\gamma_\ell)$, for all $\ell \in [1, h]$. We define $h \stackrel{\text{def}}{=} h_0 \cup \dots \cup h_h$ and prove that this is indeed a heap, by showing $\text{dom}(h_i) \cap \text{dom}(h_j) = \emptyset$, for all $0 \leq i < j \leq h$. If $i = 0$, we have $\text{dom}(h_i) \cap \text{dom}(h_j) = \emptyset$, by the definition of h_i . Else, suppose, for a contradiction, that $c \in \text{dom}(h_i) \cap \text{dom}(h_j)$, for some $1 \leq i < j \leq h$. Then $c \in \text{nodes}(\gamma_i) \cap \text{nodes}(\gamma_j)$. Since $\gamma_\ell \models \exists x_1 \dots \exists x_{\#B_\ell} \cdot B(x_1, \dots, x_{\#B_\ell})$, by Lemma 12, we obtain $c \in C_i \cap C_j$, which contradicts the fact that $\gamma_i \bullet \gamma_j$ is defined. Next, we apply the inductive hypothesis to find stores \bar{v}_ℓ and maps ι^ℓ such that, for $\ell \in [1, h]$, we have:

- $\bar{v}(z_i^\ell) = v'(z_i^\ell) \in \text{dom}(h_\ell)$ and $h_\ell(v'(z_i^\ell)) = \bar{v}(\eta(z_i^\ell))$, $\forall i \in [1, \#B_\ell]$,
- $\text{Tuples}_{\bar{v}_\ell}^j(\bar{v}(z_i^\ell)) = \{ \langle \bar{v}_\ell(\eta(z_i^\ell)) \rangle_{\text{inter}(j,k)} \mid k \in \iota^\ell(i, j) \}$, $\forall i \in [1, \#(B_\ell)] \forall j \in [1, M]$, and
- $h_\ell \models_{\bar{\Delta}}^{\bar{v}_\ell} \bar{B}_{\iota^\ell}(z_1^\ell, \dots, z_{\#B_\ell}^\ell, \eta(z_1^\ell), \dots, \eta(z_{\#(B_\ell)}^\ell))$.

First, for each $i \in [1, \#A]$ and each $j \in [1, M]$, we define $\iota(i, j) \stackrel{\text{def}}{=} \{k_1, \dots, k_{s_i}\} \cup \bigcup_{\ell=1}^h \bigcup_{k=1}^{\#B_\ell} \{ \iota^\ell(k, j) \mid x_i \approx_\pi z_k^\ell \}$, where:

- $\text{Tuples}_{\bar{v}}^j(v(x_i)) = \{c_1, \dots, c_{s_i}\}$, and
- $0 \leq k_1^{i,j} < \dots < k_{s_i}^{i,j} < \mathfrak{B}$ are integers, such that $\langle h(v(x_i)) \rangle_{\text{inter}(j, k_{\ell}^{i,j})} = c_\ell$, $\forall \ell \in [1, s_i]$; the existence of these integers is stated by point (2) of Def. 14, relative to $v(x_i)$.

Second, we define the store \bar{v} as follows:

- $\bar{v}(x_1) = v(x_1)$ and $\bar{v}(\eta(x_1)) \stackrel{\text{def}}{=} h(v(x_1))$,
- $\bar{v}(z_i^\ell) \stackrel{\text{def}}{=} v'(z_i^\ell)$, $\forall \ell \in [1, h] \forall i \in [1, \#B_\ell]$,
- $\bar{v}(\eta(z_i^\ell)) \stackrel{\text{def}}{=} h(v'(z_i^\ell))$, $\forall \ell \in [1, h] \forall i \in [1, \#B_\ell]$,
- \bar{v} is arbitrary everywhere else.

The points (1) and (2) of the statement follow from the definitions of \bar{v} and ι , respectively. To prove point (3), suppose, for a contradiction, that $k \in \{k_1^{i,j}, \dots, k_{s_i}^{i,j}\} \cap \iota^\ell(t, j) \neq$

\emptyset , for some $i \in [1, \#A]$, $j \in [1, M]$, $t \in [1, \#B_\ell]$ and $\ell \in [1, h]$, such that $x_i \approx_\pi z_t^\ell$. Then there exists a tuple of components $\mathbf{c} \in \text{Tuples}_{I_0}^j(\mathbf{v}(x_i))$, such that $\mathbf{c} = \langle \bar{\mathbf{v}}_\ell(\eta(z_t^\ell)) \rangle_{\text{inter}(j,t)} \in \text{Tuples}_{I_\ell}^j(\bar{\mathbf{v}}(z_t^\ell))$. Hence $I_0 \cap I_\ell \neq \emptyset$, which contradicts the fact that the composition $\gamma_0 \bullet \gamma_\ell$ is defined. Hence, the rule:

$$\begin{aligned} & \bar{A}_1(x_1, \dots, x_{\#(A)}, \eta(x_1), \dots, \eta(x_{\#(A)})) \leftarrow \\ & \exists y_1 \dots \exists y_m \exists \eta(y_1) \dots \exists \eta(y_m) . \bar{\Psi} * \pi * \bigstar_{\ell=1}^h \bar{B}_{1^\ell}(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell, \eta(z_1^\ell), \dots, \eta(z_{\#(B^\ell)}^\ell)) \end{aligned}$$

is well-formed and thus belongs to $\bar{\Delta}$. To obtain $h \Vdash_{\bar{\Delta}} \bar{A}_1(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#(A)}))$, by the definition of:

$$\bar{\Psi} \stackrel{\text{def}}{=} x_1 \mapsto \eta(x_1) * \bigstar_{x \in \text{fv}(\Psi)} \text{CompStates}_\Psi(x) * \bigstar_{i=1}^{\#A} \text{InterAtoms}_\Psi(x_i)$$

it is sufficient to prove the following points:

- $h_0 \Vdash^{\bar{\mathbf{v}}} x_1 \mapsto \eta(x_1)$: by the definition $\bar{\mathbf{v}}$, we have $\bar{\mathbf{v}}(\eta(x_1)) = h(\mathbf{v}(x_1))$, hence it is sufficient to prove that $\text{dom}(h_0) = \{\mathbf{v}(x_1)\}$. “ \subseteq ” Let $c \in \text{dom}(h_0)$ be a component. By the definition of $h_0 = h \setminus \bigcup_{\ell=1}^h h_\ell$, we have $\text{dom}(h_0) = \text{dom}(h) \setminus \bigcup_{\ell=1}^h \text{dom}(h_\ell) = \text{nodes}(\gamma) \setminus \bigcup_{\ell=1}^h \text{nodes}(\gamma_\ell) = \text{nodes}(\gamma_0)$, because $h \in \mathbb{G}(\gamma)$ and $h_\ell \in \mathbb{G}(\gamma_\ell)$, for all $\ell \in [1, h]$. Since $\gamma_0 \models_{\bar{\Delta}}^{\mathbf{v}'} \Psi$, we have $c = \mathbf{v}'(x)$, for some $x \in \text{fv}(\Psi)$. Suppose, for a contradiction, that x and x_1 are not the same variable, then $x \in \{z_1^\ell, \dots, z_{\#B_\ell}^\ell\}$, for some $\ell \in [1, h]$, because Δ is progressing (Def. 12). By Lemma 12, we obtain $c \in \text{nodes}(\gamma_\ell)$, contradiction. Then $c = \mathbf{v}'(x_1) = \mathbf{v}(x_1)$. “ \supseteq ” Because Δ is progressing, $\Psi = [x_1] * \Phi$, hence $\mathbf{v}(x_1) = \mathbf{v}'(x_1) \in \text{nodes}(\gamma_0) = \text{dom}(h_0)$, because $\gamma_0 \models_{\bar{\Delta}}^{\mathbf{v}'} \Psi$.
- $\emptyset \Vdash^{\bar{\mathbf{v}}} \text{CompStates}_\Psi(x)$, for each $x \in \text{fv}(\Psi)$: by the definition of $\bar{\mathbf{v}}$, $h \in \mathbb{G}(\gamma)$ and points 1 and 3 of Def. 14.
- $\emptyset \Vdash^{\bar{\mathbf{v}}} \text{InterAtoms}_\Psi(x_i)$, for each $i \in [1, \#A]$: by the definition of $\bar{\mathbf{v}}$, $h \in \mathbb{G}(\gamma)$, definition of $\mathbf{t}(i, j)$, for all $i \in [1, \#A]$ and $j \in [1, M]$ and point 2 of Def. 14.
- $\emptyset \Vdash^{\bar{\mathbf{v}}} \pi$: because $(\emptyset, \emptyset, \rho) \models^{\mathbf{v}'} \pi$ and $\bar{\mathbf{v}}$ agrees with \mathbf{v}' over $\text{fv}(\pi)$. \square

Lemma 15. *Let Δ be a progressing SID and $A \in \text{def}(\Delta)$ be a predicate, such that $h \Vdash_{\bar{\Delta}} \bar{A}_1(x_1, \dots, x_{\#(A)}, \eta(x_1), \dots, \eta(x_{\#(A)}))$, for a map $\mathbf{v} : [1, \#(A)] \times [1, m] \rightarrow 2^{[0, \mathfrak{B}-1]}$ and a store $\bar{\mathbf{v}}$. Then, the following hold:*

1. $\bar{\mathbf{v}}(x_i) \in \text{dom}(h)$ and $h(\bar{\mathbf{v}}(x_i)) = \bar{\mathbf{v}}(\eta(x_i))$, for all $i \in [1, \#A]$,
2. there exists a configuration γ , such that $h \in \mathbb{G}(\gamma)$ and $\gamma \models_{\bar{\Delta}}^{\bar{\mathbf{v}}} A(x_1, \dots, x_{\#A})$.

Proof. By fixpoint induction on the definition of $h \Vdash_{\bar{\Delta}} \bar{A}_1(x_1, \dots, x_{\#(A)}, \eta(x_1), \dots, \eta(x_{\#(A)}))$. Consider the following well-formed rule from $\bar{\Delta}$:

$$\begin{aligned} & \bar{A}_1(x_1, \dots, x_{\#(A)}, \eta(x_1), \dots, \eta(x_{\#(A)})) \leftarrow \\ & \exists y_1 \dots \exists y_m \exists \eta(y_1) \dots \exists \eta(y_m) . \bar{\Psi} * \pi * \bigstar_{\ell=1}^h \bar{B}_{1^\ell}(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell, \eta(z_1^\ell), \dots, \eta(z_{\#(B^\ell)}^\ell)) \end{aligned}$$

such that $h \Vdash_{\Delta}^{\bar{v}'} \bar{\psi} * \pi * \bigstar_{\ell=1}^h \bar{B}_{\ell}^{\ell}(z_1^{\ell}, \dots, z_{\#(B_{\ell})}^{\ell}, \eta(z_1^{\ell}), \dots, \eta(z_{\#(B_{\ell})}^{\ell}))$, where \bar{v}' is a store that agrees with v over $x_1, \dots, x_{\#(A)}$ and $\eta(x_1), \dots, \eta(x_{\#(A)})$.

(1) If $i = 1$ then $x_1 \mapsto \eta(x_1)$ is a subformula of $\bar{\psi}$, thus $\bar{v}(x_1) = \bar{v}'(x_1) \in \text{dom}(h)$ and $h(\bar{v}(x_1)) = \bar{v}'(\eta(x_1)) = \bar{v}(\eta(x_1))$. Otherwise, because Δ is progressing, $x_i \in \{z_1^{\ell}, \dots, z_{\#(B_{\ell})}^{\ell}\}$, for some $\ell \in [1, h]$ and point (1) follows from the inductive hypothesis.

(2) There exist heaps h_0, \dots, h_h , such that the following hold:

- $\text{dom}(h_i) \cap \text{dom}(h_j) = \emptyset$, for all $0 \leq i < j \leq h$ and $h = h_0 \cup \dots \cup h_h$,
- $h_0 \Vdash_{\Delta}^{\bar{v}'} \bar{\psi} * \pi$,
- $h_{\ell} \Vdash_{\Delta}^{\bar{v}'} \bar{B}_{\ell}^{\ell}(y_1^{\ell}, \dots, y_{\#(B_{\ell})}^{\ell}, \eta(y_1^{\ell}), \dots, \eta(y_{\#(B_{\ell})}^{\ell}))$, for all $\ell \in [1, h]$.

By the inductive hypothesis, there exist configurations $\gamma_1 = (C_1, I_1, \rho_1), \dots, \gamma_h = (C_h, I_h, \rho_h)$, such that $h_{\ell} \in \mathbb{G}(\gamma_{\ell})$ and $\gamma_{\ell} \models_{\Delta}^{\bar{v}} B_{\ell}(z_1^{\ell}, \dots, z_{\#(B_{\ell})}^{\ell})$, for all $\ell \in [1, h]$. We define the configuration $\gamma_0 = (C_0, I_0, \rho_0)$, as follows:

- $C_0 \stackrel{\text{def}}{=} \{\bar{v}(x_1)\}$,
- $I_0 \stackrel{\text{def}}{=} \{(\bar{v}'(z_1), p_1, \dots, \bar{v}'(z_n), p_n) \mid \langle z_1.p_1, \dots, z_n.p_n \rangle \text{ occurs in } \psi\}$,
- $\rho_0(\bar{v}'(z)) \stackrel{\text{def}}{=} q_k$ if and only if $z@q_k$ occurs in ψ , otherwise ρ_0 is arbitrary.

Moreover, we define the state map ρ as $\rho(c) \stackrel{\text{def}}{=} \rho_{\ell}(c)$ if $c \in \text{dom}(h_{\ell})$, for all $\ell \in [0, h]$ and $\rho(c)$ is arbitrary, for $c \notin \bigcup_{\ell=0}^h \text{dom}(h_{\ell})$. Since $\text{dom}(h_0), \dots, \text{dom}(h_h)$ are pairwise disjoint, ρ is properly defined. First, we prove that the composition $(C_0, I_0, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$ is defined, namely that, for all $0 \leq i < j \leq h$, we have:

- $C_i \cap C_j = \emptyset$: If $i = 0$ then either $C_0 = \emptyset$, in which case we are done, or $C_0 = \{\bar{v}(x_1)\} = \{\bar{v}'(x_1)\}$. By the definition of $\bar{\psi} = x_1 \mapsto \eta(x_1) * \phi$ and $h_0 \Vdash_{\Delta}^{\bar{v}'} \bar{\psi}$, we obtain $\bar{v}'(x_1) \in \text{dom}(h_0)$. Since $\text{dom}(h_0) \cap \text{dom}(h_j) = \emptyset$, we obtain $\bar{v}(x_1) \notin \text{dom}(h_j)$. Since $h_j \in \mathbb{G}(\gamma_j)$, we obtain $C_j \subseteq \text{nodes}(\gamma_j) = \text{dom}(h_j)$, thus $C_i \cap C_j = \emptyset$. Else $i > 0$ and, since $h_i \in \mathbb{G}(\gamma_i)$ and $h_j \in \mathbb{G}(\gamma_j)$, we obtain $C_i \subseteq \text{nodes}(\gamma_i) = \text{dom}(h_i)$ and $C_j \subseteq \text{nodes}(\gamma_j) = \text{dom}(h_j)$. But $\text{dom}(h_i) \cap \text{dom}(h_j) = \emptyset$, leading to $C_i \cap C_j = \emptyset$.
- $I_i \cap I_j = \emptyset$: If $i = 0$, by the definition of I_0 , each interaction from I_0 is of the form $(\bar{v}'(z_1), p_1, \dots, \bar{v}'(z_n), p_n)$, such that $\langle z_1.p_1, \dots, z_n.p_n \rangle$ is an interaction atom occurring in ψ . Since, moreover, Δ is progressing, we have $x_1 \in \{z_1, \dots, z_n\}$, hence $\bar{v}(x_1) \in \{\bar{v}'(z_1), \dots, \bar{v}'(z_n)\}$. Let $(c_1, p_1, \dots, c_n, p_n) \in I_j$ be an interaction. Since $h_j \in \mathbb{G}(\gamma_j)$, we obtain $\{c_1, \dots, c_n\} \subseteq \text{nodes}(\gamma_j) = \text{dom}(h_j)$, hence $\{c_1, \dots, c_n\} \cap \text{dom}(h_0) = \{c_1, \dots, c_n\} \cap \{\bar{v}(x_1)\} = \emptyset$, leading to $I_i \cap I_j = \emptyset$, because the choices of $(\bar{v}'(z_1), p_1, \dots, \bar{v}'(z_n), p_n) \in I_i$ and $(c_1, p_1, \dots, c_n, p_n) \in I_j$ are arbitrary. Else, $i > 0$ and let $(c_1^i, p_1, \dots, c_n^i, p_n) \in I_i$, $(c_1^j, p_1, \dots, c_n^j, p_n) \in I_j$ be two interactions of the same type. Since $h_i \in \mathbb{G}(\gamma_i)$ and $h_j \in \mathbb{G}(\gamma_j)$, we have $\{c_1^i, \dots, c_n^i\} \subseteq \text{nodes}(\gamma_i) = \text{dom}(h_i)$ and $\{c_1^j, \dots, c_n^j\} \subseteq \text{nodes}(\gamma_j) = \text{dom}(h_j)$, respectively. Since $\text{dom}(h_i) \cap \text{dom}(h_j) = \emptyset$, we obtain $\{c_1^i, \dots, c_n^i\} \cap \{c_1^j, \dots, c_n^j\} = \emptyset$. Since the choices of $(c_1^i, p_1, \dots, c_n^i, p_n) \in I_i$ and $(c_1^j, p_1, \dots, c_n^j, p_n) \in I_j$ are arbitrary, we obtain $I_i \cap I_j = \emptyset$.

Consequently, we define $\gamma \stackrel{\text{def}}{=} (C_0, I_0, \rho) \bullet \dots \bullet (C_h, I_h, \rho)$ and conclude by proving the following points:

- $h = \mathbb{G}(\gamma)$: we prove that $\text{dom}(h) = \bigcup_{\ell=0}^h \text{dom}(h_\ell) = \text{nodes}(\gamma)$, as required by Def. 14. The conditions (1-3) for $c_0 = \bar{v}(x_1)$ are by the definition of $\bar{\Psi}$; for $c_0 \in \text{dom}(h_\ell)$ these conditions follow from $h_\ell \in \mathbb{G}(\gamma_\ell)$. “ \subseteq ” Let $c \in \text{dom}(h)$ be a component. If $c \in \text{dom}(h_0) = \{\bar{v}x_1\}$, then $c \in C_0 \subseteq \text{nodes}(\gamma_0) \subseteq \text{nodes}(\gamma)$. Else $c \in \text{dom}(h_\ell)$, for some $\ell \in [1, h]$, then $c \in \text{nodes}(\gamma_\ell) \subseteq \text{nodes}(\gamma)$, because $h_\ell \in \mathbb{G}(\gamma_\ell)$. “ \supseteq ” Let $c \in \text{nodes}(\gamma) = \bigcup_{\ell=0}^h \text{nodes}(\gamma_\ell)$ be a component. If $c \in \text{nodes}(\gamma_0)$ then either $c \in C_0$ or c occurs in some interaction from I_0 . If $c \in C_0$ then $c = \bar{v}(x_1) \in \text{dom}(h_0) \subseteq \text{dom}(h)$, by the definition of C_0 . Else there exists an interaction $(c_1, p_1, \dots, c_n, p_n) \in I_0$, such that $c \in \{c_1, \dots, c_n\}$. In this case $c = \bar{v}'(z)$, for some variable z that occurs in an interaction atom from Ψ . Since Δ is progressing, $z \in \{z_1^\ell, \dots, z_{\#B_\ell}^\ell\}$, for some $\ell \in [1, h]$. Because $\gamma_\ell \models_{\Delta}^{\bar{v}'} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, we obtain $c = \bar{v}'(z) \in \text{nodes}(\gamma_\ell)$, by Lemma 12, and $c \in \text{dom}(h_\ell) \subseteq \text{dom}(h)$, because $h_\ell \in \mathbb{G}(\gamma_\ell)$. If $c \in \text{nodes}(\gamma_\ell)$, for some $\ell \in [1, h]$, we have $c \in \text{dom}(h_\ell) \subseteq \text{dom}(h)$, because $h_\ell \in \mathbb{G}(\gamma_\ell)$.
- $\gamma \models_{\Delta}^{\bar{v}} A(x_1, \dots, x_{\#A})$: Let the stem of the above rule from Δ be:

$$A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m . \Psi * \pi * *_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$$

Since $(C_\ell, I_\ell, \rho_\ell) \models_{\Delta}^{\bar{v}'} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$ and ρ agrees with ρ_ℓ over $\text{nodes}(\gamma_\ell)$, it follows that $(C_\ell, I_\ell, \rho) \models_{\Delta}^{\bar{v}'} B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, for all $\ell \in [1, h]$. Moreover, $(C_0, I_0, \rho) \models_{\Delta}^{\bar{v}'} \Psi$ by definition, and $(\emptyset, \emptyset, \rho) \models_{\Delta}^{\bar{v}'} \pi$, because $\emptyset \Vdash^{\bar{v}'} \pi$. Altogether, we obtain $\gamma \models_{\Delta}^{\bar{v}'} \Psi * \pi * *_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell)$, leading to $\gamma \models_{\Delta}^{\bar{v}} A(x_1, \dots, x_{\#A})$. \square

We state below the main result of this section on the complexity of the entailment problem. The upper bounds follow from a many-one reduction of $\text{Entl}[\Delta, A, B]$ to the SL entailment $\bar{A}_\iota(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A})) \Vdash_{\bar{\Delta}} \exists x_{\#B+1} \dots \exists x_{\#B} \exists \eta(x_{\#B+1}) \dots \exists \eta(x_{\#B}) . \bar{B}_{\iota'}(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B}))$, in combination with the upper bound provided by Theorem 5, for SL entailments. If $k < \infty$, the complexity is tight for CL, whereas gaps occur for $k = \infty, \ell < \infty$ and $k = \infty, \ell = \infty$, due to the cut-off on the degree bound (Prop. 1), which impacts the size of $\bar{\Delta}$ and time needed to generate it from Δ .

Theorem 6. *If Δ is progressing, connected and e-restricted and, moreover, $\text{Bnd}[\Delta, A]$ has a positive answer, $\text{Entl}^{k, \ell}[\Delta, A, B]$ is in 2EXP, $\text{Entl}^{\infty, \ell}[\Delta, A, B]$ is in $3\text{EXP} \cap 2\text{EXP-hard}$, and $\text{Entl}[\Delta, A, B]$ is in $4\text{EXP} \cap 2\text{EXP-hard}$.*

Proof. The proof consists of three parts. (1) We reduce $\text{Entl}[\Delta, A, B]$ to an equivalent SL entailment problem, for a progressing, connected and e-restricted SID. (2) This reduction provides upper bounds for $\text{Entl}^{k, \ell}[\Delta, A, B]$, in the cases $k, \ell < \infty$, $k = \infty, \ell < \infty$ and $k = \ell = \infty$, respectively. (3) We give a lower bound for $\text{Entl}[\Delta, A, B]$, by reduction from SL entailment.

- (1) We prove that, for each map $\iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$ there exists a map $\iota' : [1, \#B] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$, such that:

$$\begin{aligned} A(x_1, \dots, x_{\#A}) \models_{\Delta} \quad \exists x_{\#B+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B}) &\iff \\ \bar{A}_\iota(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A})) \Vdash_{\bar{\Delta}} \quad \exists x_{\#B+1} \dots \exists x_{\#B} \exists \eta(x_{\#B+1}) \dots \exists \eta(x_{\#B}) . & \\ \bar{B}_{\iota'}(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B})) & \end{aligned}$$

“ \Rightarrow ” Let h be a heap and \bar{v} be a store, such that $h \models_{\bar{\Delta}} \bar{A}_1(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A}))$. By Lemma 15, we have $h(\bar{v}(x_i)) = \bar{v}(\eta(x_i))$, for all $i \in [1, \#A]$ and, moreover, there exists a configuration γ , such that $h \in \mathbb{G}(\gamma)$ and $\gamma \models_{\bar{\Delta}} A(x_1, \dots, x_{\#A})$. By the hypothesis, we obtain $\gamma \models_{\bar{\Delta}} \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$, hence there exists a store \bar{v}' , that agrees with \bar{v} over $x_1, \dots, x_{\#A}$, such that $\gamma \models_{\bar{\Delta}}^{\bar{v}'} B(x_1, \dots, x_{\#B})$. By Lemma 14, there exists a store \bar{v}'' that agrees with \bar{v}' over $x_1, \dots, x_{\#B}$, such that $h(\bar{v}''(x_i)) = \bar{v}''(\eta(x_i))$, for all $i \in [1, \#B]$ and $h \models_{\bar{\Delta}}^{\bar{v}''} \bar{B}_{1'}(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B}))$, for some map $\iota' : [1, \#B] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$, because $h \in \mathbb{G}(\gamma)$. Hence, \bar{v}'' agrees with \bar{v} over $x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(\#A)$, thus we obtain $h \models_{\bar{\Delta}} \exists x_{\#A+1} \dots \exists x_{\#B} . \bar{B}_{1'}(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B}))$.

“ \Leftarrow ” Let γ be a configuration, v be a store such that $\gamma \models_{\bar{\Delta}}^v A(x_1, \dots, x_{\#A})$ and $h \in \mathbb{G}(\gamma)$ be a heap. Clearly, such a heap exists, for any given configuration, by Def. 14. By Lemma 14, there exists a map $\iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$ and a store \bar{v} , that agrees with v over $x_1, \dots, x_{\#A}$, such that $\bar{v}(\eta(x_i)) = h(v(x_i))$, for all $i \in [1, \#A]$ and $h \models_{\bar{\Delta}} \bar{A}_1(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A}))$. By the hypothesis, we have $h \models_{\bar{\Delta}}^{\bar{v}'} \bar{B}_{1'}(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B}))$, for some map $\iota' : [1, \#B] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]}$ and a store \bar{v}' that agrees with \bar{v} over $x_1, \dots, x_{\#A}$ and $\eta(x_1), \dots, \eta(x_{\#A})$. By Lemma 15, we have $\bar{v}'(\eta(x_i)) = h(\bar{v}(x_i)) = \bar{v}(\eta(x_i))$, for all $i \in [1, \#A]$ and there exists a configuration γ' , such that $h \in \mathbb{G}(\gamma')$ and $\gamma' \models_{\bar{\Delta}}^{\bar{v}'} B(x_1, \dots, x_{\#B})$. Since $h \in \mathbb{G}(\gamma) \cap \mathbb{G}(\gamma')$, by Def. 14, we obtain $\gamma = \gamma'$, hence $\gamma \models_{\bar{\Delta}}^{\bar{v}'} B(x_1, \dots, x_{\#B})$. Since, moreover \bar{v}', \bar{v} and v all agree over $x_1, \dots, x_{\#A}$, we obtain $\gamma \models_{\bar{\Delta}}^v \exists x_{\#A+1} \dots \exists x_{\#B} . B(x_1, \dots, x_{\#B})$.

Since Δ is progressing and connected, $\bar{\Delta}$ is progressing and connected as well. Moreover, $\bar{\Delta}$ is e-restricted, because Δ is e-restricted and the construction of $\bar{\Delta}$ only introduces equalities, not disequalities.

(2) The upper bound relies on the result of [19, Theorem 32], that gives a $2^{2^{\text{poly}(\text{width}(\bar{\Delta}) \cdot \log \text{size}(\bar{\Delta}))}}$ upper bound for SL entailments. Note that the number of variables in each rule from $\bar{\Delta}$ is the number of variables in its stem rule multiplied by $\mathfrak{K} + 1$, hence $\text{width}(\bar{\Delta}) \leq \text{width}(\Delta) \cdot (\mathfrak{K} + 1) = \text{width}(\Delta) \cdot \mathfrak{B} \cdot \text{intersize}(\Delta) \cdot 2^{O(\text{intersize}(\Delta))}$, because $\mathfrak{K} = \text{pos}(0, M + 1, N) = 1 + \mathfrak{B} \cdot \sum_{\ell=1}^M |\tau_{\ell}| + N = \mathfrak{B} \cdot \text{intersize}(\Delta) \cdot 2^{O(\text{intersize}(\Delta))}$. The time needed to build $\bar{\Delta}$ and its size are bounded as follows:

$$\begin{aligned} \text{size}(\bar{\Delta}) &\leq \|\bar{\Delta}\| \cdot \text{width}(\bar{\Delta}), \text{ since there are } 2^{\mathfrak{B} \cdot M \cdot \text{arity}(\Delta)} \text{ maps } \iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \mathfrak{B}-1]} \\ &\leq 2^{\mathfrak{B} \cdot M \cdot \text{arity}(\Delta)} \cdot \|\Delta\| \cdot \text{width}(\bar{\Delta}) \\ &= 2^{\mathfrak{B} \cdot 2^{\text{intersize}(\Delta)} \cdot \text{arity}(\Delta)} \cdot \|\Delta\| \cdot \mathfrak{B} \cdot \text{intersize}(\Delta) \cdot 2^{O(\text{intersize}(\Delta))} \\ &= \text{size}(\Delta) \cdot 2^{\text{poly}(\mathfrak{B} \cdot 2^{\text{intersize}(\Delta)} \cdot \text{arity}(\Delta))} \end{aligned}$$

By Prop. 1, we consider the following cases:

- if $k, \ell < \infty$ then $\mathfrak{B} = \text{poly}(\text{size}(\Delta))$, thus $\text{size}(\bar{\Delta}) = 2^{2^{O(\text{size}(\text{size}(\Delta)))}}$
- if $k = \infty$ and $\ell < \infty$ then $\mathfrak{B} = 2^{\text{poly}(\text{size}(\Delta))}$, thus $\text{size}(\bar{\Delta}) = 2^{2^{2^{O(\text{size}(\text{size}(\Delta)))}}}$
- if $k = \infty$ and $\ell = \infty$ then $\mathfrak{B} = 2^{\text{poly}(\text{size}(\Delta))}$, thus $\text{size}(\bar{\Delta}) = 2^{2^{2^{2^{O(\text{size}(\text{size}(\Delta)))}}}}$

(3) The 2EXP-hard lower bound for $\text{Entl}^{\infty, \ell}[\Delta, A, B]$ and $\text{Entl}[\Delta, A, B]$ is obtained by reduction from the SL entailment problem $\bar{A}(x_1, \dots, x_k) \Vdash_{\bar{\Delta}} \bar{B}(x_1, \dots, x_k)$, where $\bar{\Delta}$ is a progressing and connected SID, with no disequalities [18, Theorem 18]. Note that the maximum arity of Δ cannot be bounded to a constant, in order to obtain 2EXP-hardness of the SL entailment problem, hence the lower bound does not apply to $\text{Entl}^{k, \ell}[\Delta, A, B]$. The idea of the reduction is to encode each SL atomic proposition of the form $x \mapsto (y_1, \dots, y_{\bar{R}})$ by the formula $[x] * \langle x.p_0, \dots, y_{\bar{R}}.p_{\bar{R}} \rangle$. Then each model h of a SL predicate atom $\bar{A}(x_1, \dots, x_{\#(\bar{A})})$ is represented by a configuration $\gamma = (C, I, \rho)$, such that $C = \text{dom}(h)$ and $I = \{(c_0, p_0, \dots, c_{\bar{R}}, p_{\bar{R}}) \mid h(c_0) = \langle c_1, \dots, c_{\bar{R}} \rangle\}$. Since $\bar{\Delta}$ is progressing and connected, the CL SID Δ , obtained from the reduction, is progressing and connected. Since, moreover, the reduction does not introduce disequalities, Δ is trivially e-restricted. Because the reduction takes polynomial time, we obtain a 2EXP-hard lower bound. \square

7 Conclusions and Future Work

We study the satisfiability and entailment problems in a logic used to write proofs of correctness for dynamically reconfigurable distributed systems. The logic views the components and interactions from the network as resources and reasons also about the local states of the components. We reuse existing techniques for Separation Logic [40], showing that our configuration logic is more expressive than SL, fact which is confirmed by a number of complexity gaps. Closing up these gaps and finding tight complexity classes in the more general cases is considered for future work. In particular, we aim at lifting the boundedness assumption on the degree of the configurations that must be considered to check the validity of entailments.

References

1. E. Ahrens, M. Bozga, R. Iosif, and J. Katoen. Local reasoning about parameterized reconfigurable distributed systems. *CoRR*, abs/2107.05253, 2021.
2. F. Arbab. Reo: A channel-based coordination model for component composition. *Mathematical. Structures in Comp. Sci.*, 14(3):329–366, June 2004.
3. Y. Bar-Hillel, M. Perles, and E. Shamir. On formal properties of simple phrase structure grammars. *Sprachtypologie und Universalienforschung*, 14:143–172, 1961.
4. A. Basu, M. Bozga, and J. Sifakis. Modeling heterogeneous real-time components in BIP. In *Fourth IEEE International Conference on Software Engineering and Formal Methods (SEFM 2006)*, pages 3–12. IEEE Computer Society, 2006.
5. R. Bloem, S. Jacobs, A. Khalimov, I. Konnov, S. Rubin, H. Veith, and J. Widder. *Decidability of Parameterized Verification*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2015.
6. M. Bozga, R. Iosif, and J. Sifakis. Verification of component-based systems with recursive architectures. *CoRR*, abs/2112.08292, 2021.
7. J. Bradbury, J. Cordy, J. Dingel, and M. Wermelinger. A survey of self-management in dynamic software architecture specifications. In *Proceedings of the 1st ACM SIGSOFT workshop on Self-managed systems*, pages 28–33. ACM, 2004.
8. S. Brookes and P. W. O’Hearn. Concurrent separation logic. *ACM SIGLOG News*, 3(3):47–65, Aug. 2016.

9. J. Brotherston, C. Fuhs, J. A. N. Pérez, and N. Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In *CSL-LICS*, pages 25:1–25:10. ACM, 2014.
10. A. Bucchiarone and J. P. Galeotti. Dynamic software architectures verification using dynalloy. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.*, 10, 2008.
11. A. Butting, R. Heim, O. Kautz, J. O. Ringert, B. Rumpe, and A. Wortmann. A classification of dynamic reconfiguration in component and connector architecture description. In *Proceedings of MODELS 2017 Satellite Event: Workshops (ModComp)*, volume 2019 of *CEUR Workshop Proceedings*, pages 10–16. CEUR-WS.org, 2017.
12. C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10-12 July 2007, Wrocław, Poland, Proceedings*, pages 366–378. IEEE Computer Society, 2007.
13. E. Cavalcante, T. V. Batista, and F. Oquendo. Supporting dynamic software architectures: From architectural description to implementation. In L. Bass, P. Lago, and P. Kruchten, editors, *12th Working IEEE/IFIP Conference on Software Architecture, WICSA 2015*, pages 31–40. IEEE Computer Society, 2015.
14. D. Clarke. A basic logic for reasoning about connector reconfiguration. *Fundam. Inf.*, 82(4):361–390, Feb. 2008.
15. T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang. Views: Compositional reasoning for concurrent programs. *SIGPLAN Not.*, 48(1):287–300, Jan. 2013.
16. T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP 2010 – Object-Oriented Programming*, pages 504–528, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
17. J. Dormoy, O. Kouchnarenko, and A. Lanoix. Using temporal logic for dynamic reconfigurations of components. In L. S. Barbosa and M. Lumpe, editors, *Formal Aspects of Component Software - 7th International Workshop, FACS 2010*, volume 6921 of *Lecture Notes in Computer Science*, pages 200–217. Springer, 2010.
18. M. Echenim, R. Iosif, and N. Peltier. Entailment checking in separation logic with inductive definitions is 2-exptime hard. In *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*, volume 73 of *EPiC Series in Computing*, pages 191–211. EasyChair, 2020.
19. M. Echenim, R. Iosif, and N. Peltier. Decidable Entailments in Separation Logic with Inductive Definitions: Beyond Establishment. In C. Baier and J. Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, volume 183 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
20. M. Echenim, R. Iosif, and N. Peltier. Unifying decidable entailments in separation logic with inductive definitions. In *28th International Conference on Automated Deduction (CADE’21), Proceedings, to appear*, 2021.
21. R. El-Ballouli, S. Bensalem, M. Bozga, and J. Sifakis. Programming dynamic reconfigurable systems. *International Journal on Software Tools for Technology Transfer*, January 2021.
22. A. El-Hokayem, M. Bozga, and J. Sifakis. A temporal configuration logic for dynamic reconfigurable systems. In C. Hung, J. Hong, A. Bechini, and E. Song, editors, *SAC ’21: The 36th ACM/SIGAPP Symposium on Applied Computing, Virtual Event, Republic of Korea, March 22-26, 2021*, pages 1419–1428. ACM, 2021.
23. F. Farka, A. Nanevski, A. Banerjee, G. A. Delbianco, and I. Fábregas. On algebraic abstractions for concurrent separation logics. *Proc. ACM Program. Lang.*, 5(POPL), Jan. 2021.
24. X. Feng, R. Ferreira, and Z. Shao. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *Programming Languages and Systems*, pages 173–188. Springer Berlin Heidelberg, 2007.

25. K. Foerster and S. Schmid. Survey of reconfigurable data center networks: Enablers, algorithms, complexity. *SIGACT News*, 50(2):62–79, 2019.
26. H. Gaifman. On local and non-local properties. *Studies in logic and the foundations of mathematics*, 107:105–135, 1982.
27. D. Hirsch, P. Inverardi, and U. Montanari. Graph grammars and constraint solving for software architecture styles. In *Proceedings of the Third International Workshop on Software Architecture*, ISAW '98, page 69–72, New York, NY, USA, 1998. Association for Computing Machinery.
28. C. Jansen, J. Katelaan, C. Matheja, T. Noll, and F. Zuleger. Unified reasoning about robustness properties of symbolic-heap separation logic. In *ESOP*, volume 10201 of *Lecture Notes in Computer Science*, pages 611–638. Springer, 2017.
29. C. B. Jones. *Developing methods for computer programs including a notion of interference*. PhD thesis, University of Oxford, UK, 1981.
30. I. V. Konnov, T. Kotek, Q. Wang, H. Veith, S. Bliudze, and J. Sifakis. Parameterized systems in BIP: design and model checking. In *27th International Conference on Concurrency Theory, CONCUR 2016*, volume 59 of *LIPIcs*, pages 30:1–30:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
31. C. Krause, Z. Maraïkar, A. Lazovik, and F. Arbab. Modeling dynamic reconfigurations in reo using high-level replacement systems. *Science of Computer Programming*, 76:23–36, 01 2011.
32. A. Lanoix, J. Dormoy, and O. Kouchnarenko. Combining proof and model-checking to validate reconfigurable architectures. *Electron. Notes Theor. Comput. Sci.*, 279(2):43–57, 2011.
33. D. Le Metayer. Describing software architecture styles using graph grammars. *IEEE Transactions on Software Engineering*, 24(7):521–533, 1998.
34. J. Magee and J. Kramer. Dynamic structure in software architectures. In *ACM SIGSOFT Software Engineering Notes*, volume 21(6), pages 3–14. ACM, 1996.
35. A. Mavridou, E. Baranov, S. Bliudze, and J. Sifakis. Configuration logics: Modeling architecture styles. *J. Log. Algebr. Meth. Program.*, 86(1):2–29, 2017.
36. M. Noormohammadpour and C. S. Raghavendra. Datacenter traffic control: Understanding techniques and tradeoffs. *IEEE Commun. Surv. Tutorials*, 20(2):1492–1525, 2018.
37. P. W. O'Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007.
38. P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bull. Symb. Log.*, 5(2):215–244, 1999.
39. S. Owicki and D. Gries. *An Axiomatic Proof Technique for Parallel Programs*, pages 130–152. Springer New York, New York, NY, 1978.
40. J. 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.
41. Z. Shtadler and O. Grumberg. Network grammars, communication behaviors and automatic verification. In J. Sifakis, editor, *Automatic Verification Methods for Finite State Systems, International Workshop*, volume 407 of *LNCS*, pages 151–165. Springer, 1989.
42. G. Taentzer, M. Goedicke, and T. Meyer. Dynamic change management by distributed graph transformation: Towards configurable distributed systems. In *International Workshop on Theory and Application of Graph Transformations*, pages 179–193. Springer, 1998.
43. V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR 2007 – Concurrency Theory*, pages 256–271. Springer Berlin Heidelberg, 2007.
44. M. Wermelinger. Towards a chemical model for software architecture reconfiguration. *IEE Proceedings-Software*, 145(5):130–136, 1998.

45. M. Wermelinger and J. L. Fiadeiro. A graph transformation approach to software architecture reconfiguration. *Sci. Comput. Program.*, 44(2):133–155, 2002.

A Proofs